FINOVIS RESEARCH

EXECUTION & RISK · RP-001

From Model to Market

Execution Risk in Systematic Trading

March 2026

finovis.ai

Executive Summary

Systematic investment strategies are typically evaluated through the quality of their models, signals and theoretical portfolio construction. Yet a quantitative model does not trade markets directly. Its economic outcome depends on how intended exposures are translated into actual transactions under real-world conditions.

This paper examines that transition through a simple research question:

> *How and why can the move from a quantitative model to real-market execution materially alter the economic outcome of a systematic trading strategy?*

The central thesis is that execution should not be treated merely as a fixed deduction from theoretical returns or as a final operational step after strategy design. Execution is a state-dependent process through which investment intentions encounter liquidity, time, order size, market structure, infrastructure and changing market conditions.

The evidence reviewed across market-microstructure research, institutional execution frameworks and regulatory sources supports the analytical framework developed in RP-001 and several related conclusions.

First, theoretical and realized outcomes can diverge through more than explicit transaction costs. Bid-ask spread, slippage relative to a defined benchmark, interaction with available liquidity, delay, partial implementation and unexecuted quantity can all influence the realized portfolio. Implementation shortfall provides a framework for measuring this economic gap, but measurement does not by itself establish its precise causal origin.

Second, execution conditions depend on both the order and the market environment. Liquidity is multidimensional, executable quantity is finite, and the economic significance of order size depends on market activity, volatility and execution horizon. A model that scales mathematically does not therefore imply that its implementation can scale proportionally.

Third, execution logic determines how risk is distributed between price, timing, liquidity demand and completion. Different execution methods represent different ways of allocating these trade-offs rather than universally superior or inferior solutions.

Fourth, infrastructure defines part of the feasible execution environment. Brokers, venues, liquidity relationships, connectivity, market-data systems and risk controls can influence what is observable, accessible, permitted and ultimately executable. The same intended trade can therefore encounter different opportunity sets and execution paths depending on the infrastructure through which it is implemented.

Within the analytical framework developed in RP-001, these observations have implications for systematic strategy design. Relevant considerations include aligning signal horizon with feasible implementation horizon; evaluating turnover against the expected economic value of the portfolio changes it generates; incorporating implementation requirements into position sizing and capacity analysis; and making material execution assumptions explicit in backtests intended to

inform real-market deployment. Realized execution evidence can then inform transaction-cost analysis, monitoring and subsequent research without being treated as automatically causal or predictive.

RP-001 therefore distinguishes between the intended strategy represented by model outputs and target portfolios, and the realized strategy produced through execution, infrastructure and market interaction.

The principal analytical conclusion is:

> *Execution is not merely the final operational step of systematic trading. It is part of the process that determines what the strategy becomes economically.*

For CIOs, portfolio managers and quantitative investment professionals, the framework therefore adds a second question to the evaluation of systematic strategies:

> *Does the model identify attractive opportunities?*

> *Can those opportunities be translated into the intended exposures, at the required scale and under market conditions consistent with the strategy's economic objective?*

Key Findings

1. A model does not trade markets. An execution process does.

Systematic strategies generate target exposures, but realized economic outcomes depend on how those intentions are translated into executable orders and actual positions.

2. Execution costs and implementation outcomes are state-dependent rather than universally fixed.

Spread, slippage, interaction with available liquidity, delay and opportunity cost can vary with order characteristics, liquidity, volatility, execution horizon and market structure.

3. The liquidity relevant to a strategy depends partly on how, when and at what scale it trades; capacity is strategy-specific.

Within the RP-001 framework, market liquidity is not itself a property of the strategy, but the liquidity economically relevant to an execution process depends on factors including order size, timing, participation and market access. Model scalability therefore does not necessarily imply implementation scalability.

4. Execution risk includes price, time and quantity dimensions.

A strategy can diverge from its intended implementation through adverse execution prices, delays, partial fills or non-completion; infrastructure or control events such as order rejection can create similar downstream implementation effects.

5. Infrastructure shapes the feasible execution set.

Brokers, venues, liquidity relationships, market-access systems and risk controls influence what is observable, accessible, permitted and ultimately executable.

6. Execution measurement requires an appropriate benchmark.

Implementation shortfall, arrival price, VWAP, TWAP and closing-price benchmarks answer different economic questions. Measurement can identify a gap without necessarily establishing its precise causal origin.

7. Within the RP-001 framework, realized execution evidence can inform subsequent strategy research and implementation decisions.

This may include execution logic, position sizing, capacity, portfolio construction and implementation assumptions in backtesting.

Table of Contents

1. The Model-to-Market Problem

A systematic trading strategy begins as an abstraction. A model processes information, applies a set of rules and produces an investment decision: an instrument should be bought or sold, an existing position should be adjusted, or a particular exposure should be established.

At that point, however, the intended change in economic exposure has not necessarily been created.

The model has identified what the strategy intends to own or trade. The market determines whether, when, at what price and in what quantity that intention can actually be implemented.

This distinction is fundamental. Theoretical models and backtests may abstract from, simplify or approximate features that become economically relevant in live implementation. Depending on the methodology, prices may be treated as immediately observable, position changes as instantaneous, available quantity as unconstrained, or trading as having no effect on market prices.

Real trading does not share those assumptions by default. Orders must be generated and transmitted, and may need to be accepted, routed or otherwise executed against available liquidity or counterparties. Execution takes time. Prices can change during that time. Orders may be partially filled or remain unexecuted. The act of trading can itself interact with available liquidity and market prices.

The resulting difference between an intended portfolio and an implemented portfolio has long been recognized in investment research. Perold's formulation of *implementation shortfall* explicitly distinguished the economics of a paper portfolio from the economics of implementing investment decisions in actual markets. More recent empirical research has likewise examined whether returns that appear attainable in theoretical trading strategies remain economically available once real-world implementation is considered (Perold, 1988; Patton & Weller, 2020).

For systematic investing, this creates a central analytical problem:

The output of a quantitative model is not the same thing as the economic outcome of the strategy that reaches the market.

Understanding the distance between those two states is the starting point for understanding execution risk.

1.1 Theoretical Performance and Realized Performance

A theoretical strategy typically evaluates a sequence of investment decisions against historical or simulated market data. Depending on the methodology, those decisions may be represented as target weights, desired positions, trade lists or changes in exposure.

Such analysis answers an important question:

If the strategy's intended decisions could be implemented under the assumptions of the model, what economic outcome would those decisions have produced?

It does not necessarily answer a second and different question:

What outcome could have been realized by executing those decisions in an actual market?

The distinction is not merely semantic.

A paper portfolio can treat the investment decision and the acquisition of the corresponding position as effectively simultaneous. Real implementation introduces a transition between the two. During that transition, the desired trade must encounter actual prices, available quantities, market participants, execution rules and technological infrastructure.

Perold's implementation-shortfall framework was built precisely around this distinction between investment decisions on paper and their implementation in practice. In that framework, the relevant economic comparison is not limited to commissions or other directly observable trading charges. The comparison extends to the difference between the portfolio that could conceptually have been established at the decision point and the portfolio that was actually implemented (Perold, 1988).

The importance of this distinction has not disappeared with the automation of trading. Patton and Weller, for example, develop an empirical approach to examine the gap between the on-paper profitability of financial-market anomalies and implementation outcomes inferred from the realized returns of asset managers exposed to those factors (Patton & Weller, 2020).

Their analysis is specific to the strategies, market setting and methodology studied and should not be interpreted as a universal estimate of implementation costs. It nevertheless reinforces the broader point that theoretical profitability and implementable profitability are separate empirical questions.

For a systematic strategy, realized performance can therefore deviate from theoretical performance even when the underlying model behaves exactly as designed.

The source of the deviation may lie elsewhere: in the process through which the model's intended exposure is converted into actual market exposure.

This paper refers to that process as the model-to-market transition.

1.2 Why the Gap Matters in Systematic Trading

Systematic trading can create the impression that investment decisions and their implementation form one continuous automated process. In operational terms, they may indeed occur within an integrated technology stack. Economically, however, the two functions remain distinguishable.

European regulatory standards for algorithmic trading make this distinction explicitly. Commission Delegated Regulation (EU) 2017/589 differentiates an investment decision algorithm, which determines which financial instrument should be purchased or sold, from an order execution algorithm, which optimizes the generation and submission of orders after the investment decision has been taken (European Commission, 2017).

The regulatory distinction is jurisdiction-specific and is not used here as a universal definition of systematic investing. It is analytically useful because it exposes a separation that can otherwise remain hidden inside an automated process:

deciding what exposure is desirable is not the same problem as determining how that exposure should be obtained.

Consider a model that generates a target position at time t_0. The target describes the exposure the strategy intends to hold. Before that intended state becomes a realized position, several additional decisions or events may intervene:

- the signal must be translated into an executable quantity;
- an execution method must be selected;
- one or more orders must be generated;
- those orders must reach the relevant broker, venue, liquidity source or counterparty;
- available liquidity must be encountered;
- orders must be filled, partially filled, cancelled, rejected or otherwise resolved;
- and the resulting position must ultimately correspond, to some degree, with the intended exposure.

Each stage creates the possibility of divergence.

The divergence may concern price: the eventual execution price differs from the reference price associated with the model decision.

It may concern time: the position is established later than the model implicitly or explicitly assumed.

It may concern quantity: only part of the intended position is obtained.

Or it may concern the path of exposure: a position that the model treats as an instantaneous transition may in reality be accumulated or liquidated gradually.

These differences matter because a trading strategy earns returns on its realized positions, not on its intended positions.

A model may therefore remain statistically or logically unchanged while the economic characteristics of the implemented strategy change. Holding period, effective entry price, realized position size, turnover, exposure during the transition period and ultimately portfolio return can all depend on the execution process.

This observation does not imply that execution necessarily worsens outcomes. Nor does it imply that every difference between model and market is attributable to execution. Favorable price movements may improve an execution relative to a chosen benchmark, while adverse movements may occur independently of the order itself.

The important point is more neutral:

Once implementation begins, the strategy becomes exposed to a new set of economic mechanisms that do not exist in the model signal alone.

1.3 From Signal to Realized Position

To analyze these mechanisms systematically, this paper uses the following Model-to-Market Transmission Framework.

Figure 1. Model-to-Market Transmission

MODEL SIGNAL → TARGET EXPOSURE → EXECUTION DECISION → ORDER GENERATION

↓

BROKER / VENUE / LIQUIDITY INTERACTION → FILLS → REALIZED EXPOSURE → REALIZED ECONOMIC OUTCOME

Figure 1 is a conceptual analytical framework developed for this paper. It does not represent a standardized industry taxonomy or a specific proprietary trading architecture.

The framework separates the investment process into three broad states.

The first is intention.

A model produces information about the exposure the strategy seeks to obtain. Depending on the strategy, this might take the form of a directional signal, a target portfolio weight, a desired risk exposure or a change relative to an existing position.

The second is implementation.

The intended exposure must be transformed into instructions capable of interacting with the market. This requires decisions about execution that may include order quantity, timing, urgency, order type, routing or venue selection. The subsequent orders encounter real market conditions rather than the assumptions of the model.

The third is realization.

Executed quantities create actual positions. Those positions — including any difference between intended and executed quantity, price or timing — become the exposures from which the strategy's subsequent economic result is generated.

This distinction produces an important analytical shift.

Execution is often described as the final operational stage of an investment process. The Model-to-Market Framework instead treats it as a transmission layer. Its function is to convert one economic object — an intended exposure — into another — a realized exposure.

The two objects need not be identical.

A strategy that intended to acquire 100 units but acquired 60 has not simply paid a different transaction cost. It temporarily or permanently holds a different position.

A strategy that intended to establish exposure at one point in time but completes the transition later has experienced a different exposure path.

A strategy that achieves the full intended quantity but at materially different prices may reach the desired final position while generating a different economic entry point.

These examples are illustrative rather than empirical. Their purpose is to show that implementation differences cannot always be reduced to a single line item labelled "transaction costs."

The model-to-market gap can alter both the cost of reaching a position and the position that is actually reached.

That distinction will become important throughout this paper.

1.4 Four Distinct Sources of Risk

Not every difference between theoretical and realized outcomes should be called execution risk.

A useful analysis requires separating several mechanisms that can occur simultaneously but have different economic causes.

This paper therefore applies a four-part analytical classification:

Table 1. Model-to-Market Risk Taxonomy

Risk Category	Core Question	Illustrative Source of Divergence
Model Risk	Was the intended investment decision itself valid?	Specification error, unstable relationship, incorrect assumptions or inappropriate model structure
Execution Risk	Was the intended exposure obtained at the intended price, time and quantity?	Slippage relative to a defined reference price, delayed execution, partial execution or non-completion
Market Risk	How did market prices move while the strategy held, acquired or disposed of exposure?	Adverse or favourable price movements that may occur independently of whether execution operated as intended
Infrastructure / Operational Risk	Did the systems and market-access infrastructure function as required?	Connectivity failure, erroneous orders, system capacity constraints or malfunctioning controls

The classification used in this paper is an analytical framework rather than a universally standardized risk taxonomy.

MODEL RISK	EXECUTION RISK	MARKET RISK	INFRASTRUCTURE / OPERATIONAL RISK

Figure 2. Four-Layer Risk Framework. RP-001 analytical taxonomy separating the principal causal risk lenses used throughout the paper. The categories are analytically distinct but their economic consequences may interact.

Model Risk

Model risk originates before execution.

If a quantitative relationship is misspecified, if its assumptions cease to hold, or if a model responds incorrectly to the information it receives, execution cannot repair the underlying investment decision. Perfect implementation of an invalid signal remains perfect implementation of an invalid signal.

Model risk therefore concerns the quality of what the strategy intends to do.

Execution Risk

Execution risk begins once an intended exposure must be translated into actual market transactions.

It concerns whether the position implied by the investment decision can be obtained at the intended time, in the intended quantity and under economically acceptable execution conditions.

The use of *intended price* in this context requires care. A model may not specify a single executable price, and different execution benchmarks answer different analytical questions. Chapter 6 will therefore treat benchmark selection explicitly rather than assume that execution quality can be measured against one universal price reference.

Execution risk concerns the quality of how the intended trade becomes a position.

Market Risk

Market risk must remain conceptually separate from execution risk because prices can move even when execution is functioning exactly as intended.

Suppose a strategy begins buying and the market subsequently rises. Part of the price change might be related to the strategy's own interaction with liquidity; part might reflect information or trading activity entirely unrelated to the order. Distinguishing those effects can be methodologically difficult.

Conversely, a perfectly executed position may immediately lose value because the market moves against it after execution.

Market risk therefore concerns what happens to prices while exposure exists or is being established, rather than whether the execution mechanism itself performed correctly.

This distinction is causal rather than purely measurement-based. A transaction-cost or implementation-shortfall measure may capture price changes occurring during the execution period even when those changes cannot be attributed to the order itself. Accordingly, an observed execution shortfall should not automatically be interpreted as evidence of an execution failure.

Measurement identifies the economic difference between a benchmark and the realized outcome; causal attribution asks why that difference occurred.

The two questions are related, but they are not identical.

Infrastructure and Operational Risk

Execution does not occur directly between an abstract model and an abstract market. It operates through technological and institutional infrastructure.

Trading systems must process market data, generate and transmit orders, interact with brokers, trading venues, liquidity providers or counterparties, apply controls and record the resulting activity.

The MiFID II framework for algorithmic trading, for example, places explicit emphasis on system resilience, capacity, trading thresholds and limits, testing, monitoring and business-continuity arrangements (European Parliament & Council of the European Union, 2014; European Commission, 2017).

These regulatory requirements apply within their relevant legal scope and are not presented here as requirements applicable to every systematic trader. Their analytical significance is broader: they demonstrate that the infrastructure through which orders reach markets represents a distinct layer of risk rather than a frictionless conduit.

Infrastructure and operational risk concern whether the machinery required to implement the intended trade functions as expected.

The four categories are distinct, but they are not independent.

A rapidly decaying model signal may increase the urgency of execution. Greater urgency may alter the execution method. The chosen execution method may interact differently with available liquidity. Infrastructure constraints may limit which orders can be submitted. Market prices may move during the resulting delay.

The realized outcome is therefore produced by interaction across multiple layers rather than by the model alone.

This leads to the central analytical proposition of this paper:

A quantitative model produces an intended trade. The market produces an executed trade. The difference between the two is an economic variable.

The proposition should not be interpreted as a claim that execution dominates model quality, or that execution necessarily reduces strategy performance. Rather, it establishes the scope of the research problem.

A systematic strategy becomes economically real only when its intended exposures are implemented in actual markets. Evaluating such a strategy therefore requires understanding not only how investment decisions are generated, but how those decisions survive the transition from model to market.

The following chapters examine the mechanisms that determine that transition: transaction costs and spreads, market impact, liquidity and order size, execution timing and order logic, infrastructure, incomplete execution, and the methods used to measure the resulting gap between intention and realization.

2. The Economics of Execution

Once an investment decision enters the market, implementation acquires an economic cost.

Some components are directly observable. A broker may charge a commission. An exchange or trading venue may levy a fee. A transaction may be subject to taxes, stamp duties or other explicit charges.

Other costs are less visible.

A trade may cross the bid-ask spread. The available price may change before an order is executed. A sufficiently large or urgent order may interact with available liquidity and influence the prices at which subsequent quantities can be traded. An investment decision may remain unimplemented while the market moves. Part of an intended order may never be executed at all.

These effects differ in mechanism, timing and observability. Treating them as a single generic concept of “slippage” obscures economically important distinctions.

Institutional transaction-cost analysis therefore commonly separates explicit costs from implicit costs. Explicit costs may include broker commissions, exchange or venue fees, transaction taxes, stamp duties and other directly identifiable trading charges, depending on the market and institutional arrangement. Implicit costs include effects associated with spread, execution prices, market impact, delay and unexecuted quantities (CFA Institute, 2026b).

Perold's implementation-shortfall framework provides a broader economic perspective by comparing the outcome associated with an investment decision on paper with the outcome achieved through actual implementation (Perold, 1988).

The central argument of this chapter is that execution cost should not be understood as a single deduction applied after a model has generated its return.

Instead, the economics of execution arise from several distinct mechanisms between decision and realized position.

Understanding those mechanisms separately is necessary before examining how liquidity, order size, volatility and execution urgency affect them.

2.1 Explicit and Implicit Transaction Costs

The most visible costs of trading are explicit.

They arise directly from executing a transaction and may include broker commissions, exchange or venue fees, transaction taxes, stamp duties and other directly identifiable trading charges, depending on the market and institutional arrangement.

These costs are generally observable after execution and can often be attributed to individual transactions with relatively little ambiguity.

Implicit costs are different.

They arise from the interaction between the investment decision, the execution process and the market. They may not appear as a separate accounting charge, but they nevertheless affect the economic outcome of implementation.

Examples include:

- crossing the bid-ask spread;
- executing at prices different from an appropriate reference price;
- price effects associated with the order's interaction with available liquidity;
- price changes occurring while execution is delayed;
- and the economic consequence of intended quantities that remain unexecuted.

The distinction matters because a strategy can have very low explicit trading charges while still experiencing meaningful implementation costs.

Conversely, minimizing one visible component does not necessarily minimize total execution cost.

A trader could, for example, seek to avoid crossing the spread by submitting a passive limit order. If the order is not executed and the market subsequently moves away, the saving in immediate spread cost may be offset — potentially by a much larger amount — by delay or non-completion costs.

This is not an argument in favour of aggressive execution. It illustrates a more general principle:

The economic cost of execution is determined by the complete implementation path, not by any single observable charge.

This distinction is particularly important in systematic strategies because execution decisions may occur repeatedly and at scale. Even where each individual friction appears small relative to portfolio value, the relevant question is how those frictions interact with turnover, position changes and the expected economic value of the underlying signals.

Those relationships are strategy-specific. RP-001 therefore does not assume a universal transaction-cost level or a fixed basis-point deduction applicable across strategies or markets.

2.2 The Bid-Ask Spread

The bid-ask spread is one of the most fundamental costs associated with immediacy.

At a given point in time, a market participant seeking to sell immediately will generally interact with available buying interest, while a participant seeking to buy immediately will interact with available selling interest. The difference between the highest quoted bid and the lowest quoted ask constitutes the quoted bid-ask spread.

For an investor demanding immediate liquidity, the spread creates an implicit economic cost.

Relative to the contemporaneous midpoint, a simple purchase at the best ask or sale at the best bid reflects approximately one half of the quoted spread, before accounting for factors such as price improvement, order size or additional depth consumed during execution.

The spread is therefore closely connected both to transaction cost and to market liquidity. Research on market liquidity commonly considers the bid-ask spread alongside other measures such as market depth and price impact. The spread is informative, but it represents only one dimension of liquidity rather than a complete measure of executable market capacity.

It is also necessary to distinguish the quoted spread from the spread actually experienced through execution.

Quoted prices describe the best available bid and ask under a particular market structure and at a particular point in time. They do not necessarily describe the average price at which a larger order can be completed.

The effective spread is conventionally measured as twice the signed distance between the transaction price and the contemporaneous midpoint. It therefore captures the realized transaction price relative to the prevailing quote rather than relying only on displayed bid and ask prices (CFA Institute, 2026b).

This measure can be useful for evaluating individual executions, particularly where trades occur close to the quoted market.

Its interpretation becomes more difficult for large orders executed in multiple parts over time. In those cases, a single contemporaneous midpoint may not capture the broader economic process through which the entire position was implemented.

The economic implication is straightforward:

A visible spread describes the price of accessing a particular layer of liquidity; it does not necessarily describe the complete cost of implementing a position.

This distinction becomes increasingly important once order size, market depth and execution duration are introduced in Chapter 3.

2.3 Slippage: A Definition Problem

The term *slippage* is widely used in trading, but it is often used imprecisely.

Depending on context, it may refer to almost any difference between an expected and an executed price. In some discussions, market impact is called slippage. In others, the spread is included. Delay may also be incorporated.

Such usage can be operationally convenient, but it creates problems for analytical research because the same economic effect can be counted more than once.

For purposes of RP-001, the term is therefore used in a deliberately narrow sense:

Slippage is the difference between the realized execution price and an explicitly defined reference price.

For purposes of this paper, any reported slippage must therefore identify the reference price against which it is measured.

Depending on the analytical question, that reference price might be the price associated with the investment decision, the market price when an order reached the market, a contemporaneous midpoint or another appropriately defined benchmark.

Using an illustrative adverse-direction sign convention, price slippage for an execution i can be expressed as:

$$S_i = d_i \, (P_i^{exec} - P_i^{ref})$$

where:

- P_i^{exec} is the realized execution price;
- P_i^{ref} is the defined reference price; and
- d_i = +1 for a purchase and d_i = -1 for a sale.

Under this convention, a positive value represents execution at a less favourable price than the chosen reference.

The equation is intentionally simple. It defines a measurement convention rather than a model of why the price difference occurred.

That distinction is essential.

Suppose a buy order is executed above its reference price. The observed slippage may reflect several different mechanisms:

- the spread crossed by the order;
- market prices moving while execution was pending;
- the order's own interaction with liquidity;
- the execution method selected;
- or some combination of these factors.

The measured price difference alone does not establish which mechanism caused it.

This is why RP-001 does not use slippage as an umbrella term for all execution costs.

If an analysis measures slippage relative to a benchmark and separately attributes spread, impact or delay, the researcher must ensure that the same price effect is not counted twice.

The purpose of defining slippage narrowly is therefore not terminological. It is methodological:

Measurement should preserve the distinction between an observed price difference and the economic mechanism that produced it.

2.4 Market Impact

Market impact differs fundamentally from a fixed transaction fee.

It arises from the interaction between an order and the market through which that order is executed.

A sufficiently demanding order may consume available liquidity at one price and subsequently interact with liquidity at less favourable prices. As a result, the process of implementing a position can influence the prices available for its own completion.

Institutional trading research provides substantial evidence that transaction costs vary with the characteristics and difficulty of a trade. Using order-level institutional equity transactions, Keim and Madhavan found that costs differed with factors including investment style, order-submission strategy and trade difficulty (Keim & Madhavan, 1997).

Their findings are specific to the institutional equity sample studied, but they demonstrate why market impact cannot generally be represented as a fixed cost independent of the trade being executed.

Theoretical execution research formalizes the same basic problem.

Almgren and Chriss model portfolio liquidation as a trade-off between transaction costs associated with market impact and the uncertainty created by exposing a position to market-price movements during the execution horizon (Almgren & Chriss, 2001).

Within execution models such as Almgren–Chriss, temporary impact represents an execution-related price effect associated with immediate trading activity, while permanent impact represents a component assumed to persist in the price process after execution (Almgren & Chriss, 2001).

These are modelling constructs. Separating corresponding effects empirically is considerably more difficult.

If a price rises during the execution of a buy order, the observed movement cannot automatically be attributed to the order. Other investors may be trading simultaneously. New information may reach the market. The strategy itself may have generated its signal because prices had already begun to move.

Market impact is therefore partly a causal concept.

Observed execution prices can be measured directly. The portion of the price movement caused by the order itself generally requires additional assumptions, data or modelling.

This creates an important distinction:

Market impact is not synonymous with every adverse price movement observed during execution.

Execution outcomes are jointly shaped by the characteristics of the order and the market conditions under which it is implemented.

An order does not simply encounter a price; it interacts with the amount and type of liquidity available through the relevant market structure.

Chapter 3 will examine the variables that make this interaction economically important, particularly order size, liquidity, participation rate, execution duration and volatility.

2.5 Delay and Opportunity Costs

Not all implementation costs arise from the prices at which completed trades occur.

Time itself can have economic consequences.

Once an investment decision has been made, submission of the corresponding order may still be delayed by order preparation, execution scheduling, operational procedures, controls or a deliberate decision to wait.

If the market moves during that period, the achievable economics of the original decision may change.

Implementation-shortfall analysis refers to this form of divergence as delay cost when the market moves between the relevant investment decision point and the beginning of execution.

The precise terminology and decomposition depend on the methodology used, but the underlying economic principle is broader: postponing implementation leaves the strategy exposed to changes in the terms at which the intended position can subsequently be established.

Delay is not necessarily a failure.

Waiting may be intentional. A trader may delay execution to reduce immediate liquidity demand, avoid unfavourable market conditions or follow a specified execution schedule.

The economically relevant question is therefore not whether a delay occurred, but what consequences arose from the chosen timing relative to the investment objective.

A second category emerges when intended trades are not completed.

A limit order may remain unfilled. An execution schedule may end before the entire quantity is traded. Market conditions may move sufficiently far that completing the original instruction is no longer considered appropriate. Operational or risk controls may prevent further execution.

Whatever the mechanism, the strategy may finish the implementation period with less exposure than the original investment decision specified.

This gives rise to opportunity cost.

In implementation-shortfall analysis, opportunity cost captures the economic effect associated with the portion of an investment decision that was not executed.

The concept is important because an unexecuted trade does not generate a transaction record from which a conventional trading cost can be observed. Yet the difference between holding the intended position and holding the actual position can still have an economic consequence.

This returns to the distinction introduced in Chapter 1:

Execution risk is quantity risk as well as price risk.

A strategy that pays a favourable price for 50% of its intended position has not necessarily achieved a better implementation than a strategy that pays a somewhat less favourable price but obtains the economically required exposure.

Whether it has done so depends on the investment objective and the subsequent market path.

That is precisely why transaction-cost measurement must extend beyond the prices of trades that happened to occur.

2.6 Implementation Shortfall

The individual components discussed above create a measurement problem.

A commission can be observed.

A spread can be estimated.

Execution prices can be compared with a benchmark.

But none of these measures alone answers the broader economic question:

How did the actual implementation of an investment decision compare with the outcome of implementing that decision in a paper portfolio?

Perold's implementation-shortfall framework was designed to address this problem (Perold, 1988).

At its most general level, implementation shortfall compares the value associated with a hypothetical or paper implementation of an investment decision with the value produced by its actual implementation.

Conceptually, the framework can be expressed as:

$$\text{Implementation Shortfall} = \text{Paper Portfolio Outcome} - \text{Actual Portfolio Outcome}$$

with the exact sign convention depending on the analytical presentation.

The value of the framework lies less in the arithmetic than in the scope of the comparison.

Instead of measuring only commissions or completed-trade price differences, implementation shortfall seeks to capture the broader economic cost of moving from investment decision to realized portfolio.

For exposition, the total economic shortfall can be viewed as comprising explicit trading costs together with implicit delay, execution and opportunity-cost components: (CFA Institute, 2026b; Perold, 1988)

$$\text{Implementation Shortfall} \approx \text{Explicit Costs} + \text{Delay Cost} + \text{Execution Cost} + \text{Opportunity Cost}$$

The execution-cost component should not be interpreted as a pure causal estimate of market impact.

It may reflect both price pressure associated with the order and market movements occurring during the execution interval.

The decomposition is therefore an analytical framework rather than a claim that every component can always be observed or causally isolated without ambiguity.

Implementation shortfall provides a more complete economic perspective than a narrow comparison of fill prices because it retains the investment decision as the reference point.

That matters because execution quality cannot be assessed solely by asking whether a trader achieved a good price relative to other trades occurring later in the day.

If a valuable investment decision was delayed, only partially implemented or abandoned while the market moved, a benchmark that considers only completed trades may omit economically important information.

At the same time, implementation shortfall should not be interpreted as a perfect causal decomposition of execution quality.

As established in Chapter 1, measurement and causality are different questions.

A shortfall measured during an execution interval may incorporate general market movements that the strategy did not cause. Conversely, apparently favourable execution relative to a benchmark does not necessarily establish that the execution process was optimal; the benchmark itself may simply have moved favourably.

Implementation shortfall therefore measures an economic gap.

Explaining why that gap arose requires further analysis.

This distinction will become central in Chapter 6, where implementation shortfall is considered alongside alternative execution benchmarks and transaction-cost analysis.

Table 2. Components of the Model-to-Market Gap

Component	Economic Mechanism	Primary Observable	Typical Analytical Reference
Explicit Costs	Direct charges associated with trading	Commissions, fees, taxes	Currency amount / basis points
Bid-Ask Spread	Cost associated with accessing immediately available quoted liquidity	Bid, ask, transaction price	Midpoint / effective spread
Slippage	Difference between execution price and a defined reference price	Fill price and benchmark price	Decision, arrival or other stated reference
Market Impact	Price effects attributable to the interaction between an order and market liquidity	Not directly observable as a causal quantity; inferred from order data, quotes, fills and price paths	Model- and benchmark-dependent
Delay	Change in achievable economics between decision and execution	Decision and execution timestamps, market prices	Decision price / arrival price
Opportunity Cost	Economic consequence of intended quantity that remains unexecuted	Intended vs. executed quantity	Paper-position comparison

The categories are analytically distinct but may overlap in observed price data. Care is therefore required when decomposing total implementation cost to avoid double counting.

Chapter Synthesis — Execution Cost Is a Process, Not a Constant

The distinction between explicit costs, spread, slippage, market impact, delay and opportunity cost leads to a broader conclusion.

Execution cost is frequently introduced into strategy evaluation as if it were a constant parameter:

Gross Model Return − Assumed Transaction Cost = Net Return.

Such an approximation can be useful in preliminary analysis. It should not, however, be confused with a complete model of implementation.

The evidence reviewed in this chapter indicates that different execution-cost components arise through different mechanisms.

Some are fixed or directly observable. Others depend on the prices and liquidity available when the strategy trades. Some arise because implementation takes time. Others emerge because the intended position is never fully obtained.

Institutional transaction-cost research further demonstrates that trading costs can vary with the characteristics and difficulty of the trade rather than remaining constant across transactions.

From a systematic-strategy perspective, the implication is significant.

Execution costs are not merely deducted from the outcome of a strategy; they are generated by the process through which the strategy is implemented.

This is an analytical conclusion rather than a claim that every systematic strategy requires a complex execution model.

The appropriate level of sophistication depends on strategy horizon, turnover, trade size, market structure and the economic importance of execution relative to expected returns.

But once execution conditions materially affect the price, timing or quantity at which target positions are established, a constant cost assumption can conceal economically relevant variation.

The next chapter examines where that variation comes from.

In particular, it considers the relationship between liquidity, order size, execution duration, participation and volatility — and the fundamental trade-off between obtaining liquidity quickly and reducing the market footprint of execution.

3. Liquidity, Size and the Cost of Immediacy

Execution does not occur against an unlimited quantity available at a single observable price.

The prices displayed or observed in a market describe only part of the environment in which a position must be implemented. The quantity available at those prices matters. So does the amount of liquidity that may become available over time, the size of the intended order relative to normal market activity, the speed with which execution is required and the market conditions prevailing while the order is being completed.

These variables are closely connected.

A small order in a deep market may be executable with little observable interaction with prices. The same nominal order in a less liquid instrument may require a different execution path. A large order can potentially be spread across a longer period, reducing immediate demand for liquidity, but doing so leaves the strategy exposed to market movements while the position remains incomplete.

Execution therefore presents an allocation problem across both quantity and time.

The central proposition of this chapter is that liquidity cannot be treated as a static property of an instrument and execution cost cannot be evaluated independently of the scale and urgency of the trade.

Instead, the relevant economic question is:

How much exposure is the strategy attempting to obtain, relative to the liquidity available, and over what period must that transition occur?

3.1 Liquidity Is Multidimensional

Market liquidity is often described informally by saying that an instrument is either "liquid" or "illiquid."

For execution analysis, that distinction is too coarse.

A market can display a narrow bid-ask spread but provide relatively little quantity at the quoted prices. Another may display greater depth but require more time to execute a substantial order. A market may appear liquid under ordinary conditions while becoming significantly more difficult to trade during periods of stress.

Liquidity is therefore better understood through several related dimensions rather than a single statistic.

Institutional and central-bank research commonly distinguishes among measures including bid-ask spread, market depth and price impact (CFA Institute, 2026b; Federal Reserve Board, 2019).

The spread describes the difference between immediately available buying and selling prices.

In order-book markets, quoted depth describes the quantity displayed at or around prevailing bid and ask prices. In dealer, request-for-quote or fragmented OTC markets, comparable liquidity may not be observable through a consolidated order book.

Price-impact measures seek to describe the relationship between trading activity and subsequent or contemporaneous price response.

These measures capture different aspects of the same execution problem.

A narrow spread does not guarantee that a large quantity can be executed near the quoted price.

High displayed depth does not guarantee that the same depth will remain available as an order progresses.

And historically low estimated price impact does not guarantee that equivalent conditions will prevail when volatility, market participation or liquidity provision changes.

A price-impact liquidity measure should also not be confused with a causal estimate of the market impact generated by one specific order.

The former describes an empirical relationship between trading and price response. The latter requires an attribution of causality.

From an execution perspective, liquidity therefore concerns not simply the existence of a tradable price, but the ability to transfer the required amount of risk within the required period without generating economically significant price concessions.

This framing is intentionally broader than any single liquidity metric.

Liquidity itself is not directly observable as one universal quantity. It must be inferred from characteristics such as prices, spreads, available quantities, trading activity and the market's response to order flow.

This creates the first important principle of the chapter:

A quoted price is not the same thing as executable capacity.

That distinction becomes increasingly important as order size increases.

3.2 Order Size and Market Depth

Consider an illustrative order-book market in which the best ask displays 100 units.

A strategy intending to purchase 20 units may be able to obtain its desired quantity entirely at that price.

A strategy intending to purchase 1,000 units faces a different problem.

Even if both strategies observe the same quoted ask when the investment decision is made, the larger order cannot assume that the displayed price applies to the entire intended quantity. It may consume the quantity available at the best ask and subsequently interact with additional liquidity at different prices.

The example is illustrative, but the economic mechanism is fundamental.

Order size must be considered relative to available liquidity rather than in isolation.

Institutional execution practice therefore commonly evaluates order size relative to measures of market activity. In equity markets, percentage of average daily volume is frequently used as a standardized indication of the scale of an order relative to normal trading activity.

Percentage of average daily volume is therefore primarily a normalization measure for order size, not a direct measure of market liquidity.

Average daily volume does not describe the precise liquidity available at a particular moment. Trading volume can vary substantially within a day. Displayed volume may not represent all available liquidity. In OTC or fragmented markets, comprehensive market-wide volume information may not be observable in the same way as in centralized exchange environments.

Nevertheless, relative order size is economically important because the demand represented by an order must ultimately be absorbed by the market.

Kyle and Obizhaeva provide broader empirical evidence that transaction costs and the size of risk transfers have systematic relationships with observable market characteristics including dollar trading volume and volatility. Their empirical tests use more than 400,000 portfolio-transition orders (Kyle & Obizhaeva, 2016).

Their results support the narrower proposition that the economic significance of order size depends on the characteristics of the market in which risk is transferred.

This evidence does not imply that one universal function describes transaction costs across every market.

It does reinforce a more modest conclusion:

The same nominal trade size can represent very different execution demands in markets with different levels of activity, volatility and capacity to absorb risk.

Order size is therefore not merely a portfolio-construction variable.

Once a target position reaches the execution layer, size becomes a market-microstructure variable as well.

3.3 Participation Rate

One way to describe the intensity of an execution is to compare the quantity being traded with the amount of relevant market activity occurring during the same period.

This relationship is commonly expressed through a participation rate.

Conceptually:

Participation Rate = Executed Quantity / Observed or Estimated Relevant Market Volume over the Defined Interval

The relevant volume universe may be market-wide, venue-specific or estimated, depending on market structure and data availability.

Participation rate provides a useful bridge between absolute order size and the liquidity environment.

A 10,000-unit execution may represent a very small fraction of trading in one instrument and a dominant share of observed or estimated activity in another. The economic significance of the order is therefore better understood when its size is considered relative to the market activity available to absorb it.

Higher participation mechanically represents a larger share of observed or estimated trading activity.

Whether this results in greater market impact depends on the execution method, available liquidity and prevailing market conditions.

Lower participation distributes the desired position across a larger amount of market activity, but typically requires more time if the total intended quantity remains unchanged.

This produces an important relationship:

Participation rate connects execution size with execution duration.

For a given target quantity, reducing participation generally requires allowing more market activity — and therefore potentially more time — to pass before the order can be completed.

Increasing participation may permit faster completion, but requires the strategy to account for a larger fraction of trading activity during the execution period.

This relationship should not be interpreted as a universal rule that lower participation always produces lower transaction costs.

Liquidity supply can respond dynamically. Trading activity is not constant. Information conditions change. Some orders may benefit from faster completion, while others may be implemented more efficiently over a longer horizon.

Participation is therefore best understood as an execution-control variable, not as an independent measure of execution quality.

Its importance lies in the trade-off it creates between the intensity of market interaction and the time required to obtain the desired position.

3.4 Execution Duration and the Cost of Immediacy

A strategy cannot generally minimize both market interaction and execution time without constraint.

This creates the central trade-off of execution.

If a position is implemented over a shorter horizon, the strategy reduces the period during which the intended exposure remains incomplete.

That can be valuable.

Market prices may move while execution is underway. A signal may lose economic relevance over time. Risk exposures may remain different from their intended levels until the transition is complete. In such circumstances, completing the order more quickly can reduce exposure to timing and non-completion risk.

A shorter execution horizon may, however, require more intensive interaction with immediately available liquidity.

For a sufficiently large order, that can increase the quantity consumed over a shorter period and potentially increase the price concession required to complete the trade.

A longer execution horizon reverses the balance.

By distributing an order over more time and potentially more market activity, the strategy may reduce the intensity of its immediate liquidity demand.

But the position remains incomplete for longer.

The strategy therefore continues to bear uncertainty regarding the prices at which the remaining quantity will eventually be executed.

This trade-off forms the theoretical core of the optimal-execution problem introduced by Almgren and Chriss: execution involves balancing expected transaction costs associated with market impact against the risk created by market-price uncertainty during the execution horizon (Almgren & Chriss, 2001).

More recent research extends the analysis by allowing liquidity supply itself to respond to large orders.

Capponi, Menkveld and Zhang model the execution of large orders in the presence of strategic competitive market makers (Capponi et al., 2025).

Their framework establishes a joint relationship between order size, execution duration and resulting price effects, rather than treating execution cost as a function of order size alone.

Their model concerns a particular market setting and should not be interpreted as a universal description of all execution environments.

Its broader implication is nevertheless important:

The economic effect of a large order cannot necessarily be understood without considering how long the market is given to absorb it.

Figure 3. The Cost of Immediacy

SHORTER EXECUTION HORIZON	TRADE-OFF	LONGER EXECUTION HORIZON
Potential benefits: • shorter time to completion; • reduced exposure to price movements before completion; • lower non-completion risk; • faster convergence toward intended exposure. Potential costs: • greater immediate liquidity demand; • potentially greater price impact; • potentially larger execution-price concessions.	EXECUTION TRADE-OFF	Potential benefits: • liquidity demand distributed across more time; • potentially lower immediate market footprint; • greater opportunity for liquidity replenishment. Potential costs: • longer exposure to market-price uncertainty; • greater timing risk; • greater possibility that the original signal changes or decays; • longer period of divergence between intended and realized exposure.

Figure 3 is conceptual. It does not imply a universal monotonic relationship between execution horizon and transaction cost, nor does it identify an optimal execution horizon for any particular strategy or market. Execution urgency is one factor that may influence the appropriate horizon.

The important concept is not that execution should always be fast or always be slow.

It is that immediacy has an economic price, and the value of immediacy depends on the characteristics of both the strategy and the market.

3.5 Volatility and Changing Market Conditions

The execution problem would be considerably simpler if liquidity conditions were stable.

They are not.

Bid-ask spreads, available depth, trading volume and price volatility can change over time. Conditions observed when an execution plan is designed may differ from those encountered when the order actually reaches the market.

Volatility is particularly important because it affects both sides of the execution trade-off.

First, greater price volatility increases the uncertainty associated with leaving a position unexecuted over time. If prices are changing rapidly, extending the execution horizon can expose the remaining quantity to larger market movements.

Second, periods of elevated volatility are frequently associated with weaker liquidity conditions.

Federal Reserve research has documented relationships between higher volatility and deterioration in liquidity measures in markets including U.S. Treasuries and equity futures (Federal Reserve Board, 2019).

One plausible mechanism is that greater uncertainty increases the risk faced by liquidity providers, leading them to supply less liquidity or demand greater compensation for doing so.

This mechanism should not be interpreted as a universal causal law. The relationship between volatility and liquidity depends on market structure, participants and prevailing conditions.

The relationship is also visible in FX markets.

During the severe market disruption of March 2020, the BIS Markets Committee documented a sharp rise in currency-market volatility and a deterioration in FX bid-ask spreads, with the deterioration particularly pronounced for larger orders (BIS Markets Committee, 2020).

Importantly, the BIS also observed that overall FX market functioning was not significantly impaired (BIS Markets Committee, 2020).

Deteriorating execution conditions and outright market dysfunction are therefore not equivalent.

The March 2020 episode should not be interpreted as a calibration for normal market conditions or as evidence that every increase in volatility produces an identical liquidity response.

Its value is conceptual.

It demonstrates that an execution process calibrated to ordinary conditions may encounter a materially different environment during a market shock.

This creates a state-dependent execution problem.

An order size that represents modest liquidity demand in normal markets may become economically significant when available liquidity contracts.

An execution horizon that appears appropriate under ordinary volatility may expose the strategy to substantially greater timing risk when prices move more rapidly.

And an execution rule based on historical average volume may become less informative when actual trading patterns diverge sharply from those averages.

Market conditions therefore matter not simply because prices change.

They matter because the terms under which exposure can be transferred can change at the same time.

From the analytical perspective adopted in this paper:

Execution capacity should therefore be treated as conditional on market state.

3.6 Strategy Capacity

The interaction between order size, liquidity, participation, execution duration and volatility leads to a broader question:

How much capital can a systematic strategy deploy before implementation materially changes its economics?

This is the problem of strategy capacity.

For purposes of this paper, strategy capacity refers to the scale of capital or exposure that can be implemented while maintaining execution conditions consistent with the strategy's economic objective.

This is an analytical definition used in RP-001 rather than a standardized industry measure.

Capacity is sometimes discussed as if it were a single assets-under-management threshold.

Such a number can be misleading without reference to how the strategy trades.

Two strategies managing the same amount of capital may impose very different demands on markets.

One may hold positions for long periods and trade gradually in highly liquid instruments.

Another may rebalance frequently, concentrate trading into shorter windows or seek exposure in instruments with less available liquidity.

Their assets under management may be identical while their execution requirements are not.

From the Model-to-Market perspective developed in this paper, strategy capacity should therefore be understood as an implementation constraint, not merely a fund-size statistic.

This is an analytical synthesis rather than an established universal definition.

The reasoning is supported by several strands of execution research.

Kyle and Obizhaeva show empirically that transaction costs are related to the scale of risk transfers, trading volume and volatility (Kyle & Obizhaeva, 2016).

Capponi, Menkveld and Zhang show theoretically that large-order economics can depend jointly on order size and execution duration when liquidity supply responds endogenously (Capponi et al., 2025).

Institutional execution frameworks similarly treat order size, liquidity profile, volatility and market conditions as relevant inputs when selecting an appropriate trade strategy.

Taken together, this evidence suggests that capacity cannot be assessed independently of the mechanism through which positions are established and unwound.

For a systematic strategy, relevant considerations may include:

- total capital deployed;
- position concentration;
- turnover;
- frequency of rebalancing;
- typical order size;
- order size relative to available market activity;
- liquidity of the instruments traded;
- required execution horizon;
- urgency associated with the underlying signals;

- and the behaviour of liquidity under changing market conditions.

These variables interact.

Increasing capital can increase order size.

Larger orders may require longer execution horizons or higher participation.

Longer execution horizons increase exposure to timing uncertainty.

Higher participation increases the strategy's share of observed or estimated trading activity and may alter the way the order interacts with available liquidity.

The resulting execution characteristics can, in turn, alter the realized economics of the strategy.

This produces an important analytical distinction between model scalability and implementation scalability.

A quantitative model can often be evaluated at a larger notional size simply by multiplying its theoretical positions.

Markets cannot necessarily be scaled in the same way.

The assumption that a position can be multiplied without changing the prices, quantities or time required for implementation is itself an execution assumption.

For that reason, a strategy should not be described as economically scalable solely because its model can generate proportionally larger theoretical positions.

A more demanding question is required:

Can the intended exposure be increased while preserving implementation conditions that remain consistent with the strategy's economic objective?

There is no universal answer.

Capacity depends on strategy design, instruments, market structure, trading horizon and prevailing liquidity conditions.

It can also change over time.

Accordingly, RP-001 does not propose a universal capacity formula or a fixed relationship between assets under management and execution cost.

The analytical conclusion is narrower:

Strategy capacity cannot be evaluated independently of execution.

Chapter Synthesis — Liquidity Is a Constraint on Translation

Chapter 1 introduced execution as the transmission layer between intended and realized exposure.

Chapter 2 established that the economic cost of this transmission arises through several distinct mechanisms rather than through a single fixed deduction.

This chapter adds a further constraint:

Real-market implementation cannot generally assume unlimited executable quantity at a fixed price or on a predetermined timetable.

Liquidity has multiple dimensions.

Order size matters relative to the market in which the order is implemented.

Participation determines how large a share of relevant market activity the strategy represents during execution.

Execution duration creates a trade-off between immediacy and timing risk.

Volatility can alter both market-price uncertainty and the conditions under which liquidity is supplied.

These relationships make execution inherently state-dependent.

They also explain why the economics of a systematic strategy can change as its scale changes, even if the underlying model does not.

Execution must therefore answer two related questions:

How much of the intended exposure can the market absorb under the relevant conditions?

and

Over what horizon should the strategy attempt to obtain it?

The next chapter moves from these constraints to the decisions made within them.

It examines Execution Logic: Choosing How to Trade — including market and limit orders, passive and aggressive execution, order slicing, participation-based strategies, routing, latency, partial fills and rejected orders.

4. Execution Logic: Choosing How to Trade

Once a systematic strategy has determined the exposure it seeks to obtain, execution still requires a series of choices.

The strategy must determine how aggressively to interact with available liquidity, whether to prioritize execution probability or price control, whether a larger instruction should be divided into smaller orders, how execution should respond to observed market activity, and — where multiple trading channels are available — how orders should be distributed across them.

These choices do not alter the underlying investment signal.

They determine how that signal attempts to become a realized position.

This distinction is important because there is rarely one execution method that dominates across all objectives and market conditions. A method designed to obtain exposure quickly may accept greater uncertainty around execution price. A method designed to control price may accept a greater probability of incomplete execution. An algorithm designed to participate gradually in market activity may reduce the concentration of trading over time while leaving the strategy exposed to market movements for longer.

The appropriate execution logic therefore depends on the economic objective of the trade.

Relevant considerations include:

- the urgency of the investment decision;
- the expected persistence or decay of the underlying signal;
- order size;
- available liquidity;
- volatility;
- acceptable execution-price uncertainty;
- tolerance for partial or non-execution;
- market structure;
- and the infrastructure through which orders reach the market.

The central proposition of this chapter is therefore not that one execution technique is superior.

It is that:

Execution logic determines which combination of price risk, timing risk, liquidity demand and completion risk the strategy accepts in attempting to obtain its intended exposure.

4.1 Market Orders and Limit Orders

In conventional order-driven markets, market and limit orders provide a useful illustration of the trade-off between immediacy and price control.

Exact order definitions and handling rules depend on the relevant market, broker and venue.

A market order generally prioritizes prompt execution at available prices without imposing a limit price (U.S. Securities and Exchange Commission, 2011).

Its principal economic advantage is immediacy.

The order does not specify a maximum purchase price or minimum sale price that must be satisfied before execution can occur. It therefore generally places greater priority on obtaining exposure promptly than on controlling the exact execution price.

That does not mean that a market order guarantees a particular price.

The price visible when the order is submitted may no longer be available by the time execution occurs. The displayed quantity may also be insufficient for the entire order. A larger market order may therefore execute across several price levels or through a sequence of fills.

Market orders generally offer greater immediacy and execution probability than non-marketable limit orders, but neither a particular price nor complete execution should be treated as universally guaranteed.

A limit order imposes a price constraint (U.S. Securities and Exchange Commission, 2011).

A buy limit order specifies the maximum price the investor is willing to pay. A sell limit order specifies the minimum acceptable sale price.

This provides greater control over the prices at which execution may occur.

But that control comes with another form of uncertainty:

the order may not execute.

If available market prices never satisfy the limit condition — or if insufficient quantity is available at the specified price before the order expires or is cancelled — some or all of the intended exposure can remain unimplemented.

The distinction can therefore be summarized as a tension between price control and completion risk.

This should not be reduced to a simple rule that market orders are appropriate for urgent trades and limit orders for non-urgent trades.

Actual order selection can be more complex.

A sufficiently aggressive limit order may be immediately marketable.

A non-marketable limit order may rest in the market and wait for execution.

A market order may encounter deep liquidity and execute with little observable price dispersion.

A passive limit order may obtain favourable pricing but expose the strategy to substantial non-execution risk.

The relevant question is not the label attached to the order.

It is:

Which uncertainty is the strategy willing to accept in exchange for controlling another?

Market and limit orders therefore represent two basic mechanisms through which execution logic allocates risk between price and completion.

4.2 Passive and Aggressive Execution

A related distinction concerns how an order interacts with available liquidity.

In an order-book environment, a passive order typically rests in the market and waits for another participant to trade against it.

An aggressive or marketable order instead interacts with liquidity that is already available.

Passive versus aggressive execution describes liquidity interaction, not simply order type.

The distinction is therefore not identical to market order versus limit order.

A sufficiently aggressive limit order can be immediately marketable, while a non-marketable limit order may rest passively in the book.

The distinction is economically important.

Passive execution can provide greater control over the price at which a transaction occurs and may avoid immediately crossing the quoted spread.

But passive execution introduces uncertainty.

The market may never trade at the required price.

Other orders may have priority.

Only part of the desired quantity may execute.

And the market may move away before the intended position is obtained.

Passive orders may therefore face several forms of execution risk, including non-execution and the possibility that execution occurs primarily when market conditions become less favourable to the resting order.

Aggressive execution increases the probability of obtaining exposure promptly by interacting with currently available liquidity.

This can reduce the time during which the strategy remains away from its target position.

But the strategy accepts the prices and quantities currently available and may consume liquidity across more than one price level.

The choice therefore extends the trade-off introduced in Chapter 3:

greater immediacy generally requires greater willingness to interact with existing liquidity, while greater patience increases exposure to timing and non-completion risk.

Again, neither approach is inherently superior.

The economic value of passive execution depends on the value of waiting.

The economic value of aggressive execution depends on the value of obtaining the position now.

A rapidly decaying signal may make non-execution particularly costly. A less time-sensitive portfolio adjustment may permit more patient implementation.

Execution logic should therefore be connected to the objective of the investment decision rather than evaluated in isolation.

4.3 Order Slicing and Scheduling

Large investment decisions do not necessarily enter the market as single orders.

A larger parent order may be divided into multiple smaller child orders that are submitted over time, across venues or according to changing market conditions.

This process is commonly referred to as order slicing.

The economic motivation follows directly from Chapter 3.

Submitting the entire desired quantity at once may require substantial immediate liquidity and can expose the order to less favourable prices beyond the quantity available at the best quotes.

Slicing allows the execution process to distribute that demand.

But slicing does not eliminate the execution problem.

It changes its structure.

Once an order is divided, execution logic must determine:

- the size of individual child orders;
- when they should be submitted;
- whether they should provide or demand liquidity;
- how execution should respond to changing market conditions;
- and when the remaining quantity should be completed more aggressively.

Different algorithms implement these decisions in different ways.

Some follow relatively predetermined schedules.

Others respond to observed volume or activity.

Others seek available liquidity across trading venues or liquidity sources.

Others adjust aggressiveness according to the difference between actual execution progress and a desired benchmark or completion objective.

Institutional execution frameworks therefore classify algorithms according to their principal objective rather than treating algorithmic execution as one homogeneous technique.

The BIS describes execution algorithms in the FX market as tools designed to execute a predefined quantity according to user instructions, frequently by dividing a larger order into smaller transactions and distributing them across time and available liquidity sources (BIS Markets Committee, 2020).

An execution algorithm does not determine whether the underlying investment decision is economically valid.

It implements an execution objective according to a defined set of instructions or rules.

Automation therefore changes how execution decisions are applied; it does not eliminate the distinction between investment decision and execution decision.

Order slicing also creates an additional exposure.

Until the final child order has been completed, the realized position remains different from the intended final position.

The strategy therefore trades one form of risk — concentrated immediate liquidity demand — against another — continued uncertainty over the prices and quantities associated with the remaining execution.

4.4 Participation-Based Execution

Chapter 3 introduced participation rate as the share of relevant market activity represented by an execution.

Participation-based execution turns that relationship into an execution rule.

A percentage-of-volume, or participation-based, algorithm seeks to trade at approximately a specified proportion of a defined reference measure of market activity over time.

Conceptually:

Participation Measure = Strategy Executed Quantity / Defined Reference Market Activity over the Relevant Interval

The reference activity measure must be defined for the specific execution methodology.

Depending on market structure, it may rely on observed transaction volume, venue-specific activity or estimates of relevant liquidity rather than a directly observable market-wide volume total.

In transparent exchange markets, the reference may be based on reported transactions.

In fragmented or OTC markets, the relevant activity measure may instead be venue-specific, partially observed or estimated from available transaction and liquidity information.

The BIS notes that FX participation algorithms can use indicators such as recently observed transaction volume, estimated liquidity or order-book information when determining execution pace (BIS Markets Committee, 2020).

Participation-based execution has an intuitive advantage.

It links execution pace to market conditions represented by the chosen activity measure.

When the execution algorithm observes or estimates greater relevant market activity or available liquidity, its rules may permit a faster execution pace.

When those indicators decline, the permitted pace may slow.

This can prevent a fixed schedule from demanding the same quantity of liquidity in materially different market environments.

But participation-based execution is not equivalent to minimizing transaction costs.

A target participation rate is still an execution choice.

A higher target may complete the order more quickly but make the strategy a larger component of relevant market activity.

A lower target may distribute execution more broadly over market activity but extend the execution horizon.

Moreover, observed or estimated market activity does not measure the economic value of the underlying investment signal.

A strategy whose expected alpha is decaying rapidly may have an economic reason to depart from a mechanically low participation target.

Conversely, a low-urgency trade may not require rapid completion merely because current market activity is elevated.

Participation-based execution therefore solves a specific problem:

how execution pace should relate to a defined measure of market activity.

It does not by itself determine how urgently the investment objective should be implemented.

4.5 Routing and Venue Selection

Execution logic also determines *where* an order interacts with the market.

In some instruments, trading is concentrated on a relatively small number of venues.

In others, liquidity is fragmented.

Equity orders may interact with exchanges and other multilateral trading venues.

Foreign exchange is primarily an OTC market in which liquidity is distributed across dealer platforms, electronic communication networks, single-dealer venues, aggregators and other trading channels.

The relevant market therefore may not present one consolidated source of executable liquidity.

This creates a routing and liquidity-access problem.

In market structures with multiple execution venues, *smart order routing* can be used to direct orders among eligible venues according to predefined criteria.

In fragmented OTC markets such as FX, execution algorithms may instead navigate multiple liquidity sources, dealer platforms and electronic venues according to the protocols available in that market.

The BIS identifies fragmentation as an important characteristic of electronic FX execution and describes execution algorithms as distributing smaller pieces of parent orders across different venues and liquidity pools (BIS Markets Committee, 2020).

Routing or liquidity-source selection can therefore affect the set of prices and quantities that an execution process encounters.

But venue selection should not be reduced to a search for the displayed best price.

A quoted price is only one relevant attribute of an execution opportunity.

Available quantity, likelihood and speed of execution, protocol design and other market-specific characteristics can also affect the realized outcome.

A trading venue or liquidity source can therefore affect not only what price is quoted, but whether and how an instruction becomes an actual fill.

The precise mechanisms are market-specific.

RP-001 does not assume that one venue structure, routing rule or liquidity-access model is optimal across asset classes.

The narrower proposition is sufficient:

Where liquidity is fragmented, execution quality can depend partly on how the execution process identifies and accesses that liquidity.

The infrastructure implications of this proposition are examined more fully in Chapter 5.

4.6 Latency and Execution Timing

Execution requires time.

That statement is obvious, but the economic meaning of latency depends heavily on the strategy being executed.

The term *latency* is often associated with high-frequency trading, where differences measured in milliseconds or smaller intervals can influence how quickly a participant observes and responds to market events.

Hasbrouck and Saar, for example, define low-latency trading activity in the millisecond environment and study its relationship with market quality in NASDAQ equities (Hasbrouck & Saar, 2013).

That literature is important, but it should not be generalized to all systematic trading.

For many strategies, the economically relevant latency may be measured in seconds, minutes or longer rather than milliseconds.

For analytical purposes, RP-001 separates latency into four stages.

This classification is used to identify where time enters the model-to-market process and is not presented as a standardized industry taxonomy.

Signal latency

The interval between the market or informational event relevant to the strategy and its incorporation into the model.

Decision latency

The interval between model evaluation and the generation of an actionable investment decision.

Order-generation and transmission latency

The interval required to translate that decision into an order and transmit it through the relevant infrastructure.

Market-access and execution latency

The interval between order submission and the resulting acknowledgement, market interaction and execution.

The purpose of this classification is to show why a generic statement that "lower latency is better" is incomplete.

Latency matters economically only when the delay is significant relative to the horizon over which the relevant information, price or liquidity condition changes.

For a strategy attempting to react to short-lived order-book information, milliseconds may matter materially.

For a strategy implemented over a materially longer horizon, millisecond-level network improvements may be economically secondary to variables such as execution horizon, participation and liquidity access.

Conversely, apparently small operational delays can matter when the underlying signal decays rapidly or when markets are moving unusually quickly.

This produces the relevant principle:

The economic value of speed is conditional on the time scale of the strategy and the market process it seeks to exploit or manage.

Latency should therefore be treated as an execution variable, not as a universal measure of technological quality.

4.7 Partial Fills and Rejected Orders

Execution is not binary.

An instruction does not always progress directly from "order submitted" to "desired position obtained."

A large order may be filled in several transactions.

A limit order may execute only partially before the market moves away.

An execution schedule may end while some intended quantity remains outstanding.

A venue may cancel or expire an order under its rules.

And in some market-access environments, an order may be rejected before it reaches the market at all.

These outcomes should be distinguished.

A partial fill means that some quantity has been executed but the full desired quantity has not yet been obtained.

A non-fill means that an order or remaining quantity did not execute under the conditions imposed on it.

A rejection means that an order was not accepted for further processing or market entry by a relevant system, broker, venue or control layer.

The economic consequences differ.

Partial execution creates a position, but not necessarily the intended position.

Non-execution leaves the strategy without the targeted exposure.

Rejection can reveal an infrastructure, control or market-access constraint rather than a lack of available market liquidity.

A rejection should not automatically be classified as execution risk in the causal taxonomy introduced in Chapter 1.

Its immediate cause may lie in the infrastructure, operational or risk-control layer; its economic consequence can nevertheless appear as execution delay, non-completion or a difference between intended and realized exposure.

The U.S. SEC's Market Access Rule provides a useful jurisdiction-specific example of the latter distinction.

Rule 15c3-5 requires broker-dealers with market access to maintain pre-trade controls designed, among other things, to prevent orders that exceed specified credit, capital, price or size parameters from entering an exchange or alternative trading system (U.S. Securities and Exchange Commission, 2010).

Such controls are not evidence that rejected orders are inherently undesirable.

Their purpose is risk management.

Nor does the U.S. rule establish the architecture used in every jurisdiction or market.

Its analytical relevance is narrower:

an executable investment intention can encounter constraints before it interacts with market liquidity.

Partial fills and rejections therefore reinforce the quantity-risk problem introduced earlier in this paper.

Suppose a strategy intends to obtain an exposure Q^*, but execution results in a realized quantity Q^R.

Using signed quantities consistently, the residual quantity can be expressed conceptually as:

$$Q^{residual} = Q^* - Q^R$$

A non-zero residual is not merely an operational statistic.

It represents a difference between intended and realized exposure.

The execution process must therefore determine what happens next.

The strategy may:

- continue attempting execution;
- alter its aggressiveness;
- modify or cancel the remaining instruction;
- reassess whether the original signal remains valid;
- or accept the incomplete position.

The correct response depends on the investment objective and cannot be determined from fill rate alone.

This is why fill probability and execution price must be considered jointly.

A method that achieves a highly favourable execution price but leaves a strategically important position largely unfilled may not have produced a superior economic outcome.

Likewise, maximizing fill probability at any price may be inconsistent with the strategy's objective.

Execution logic must determine the acceptable balance between the two.

Table 3. Order-Execution Trade-offs

Execution Choice	Typical Execution Objective	Potential Advantage	Principal Exposure
Market / Marketable Order	Obtain exposure promptly	Higher immediacy and execution probability	Execution-price uncertainty; potential depth consumption; completion still subject to market and control constraints
Passive Limit Order	Control acceptable execution price	Price control; potential liquidity provision	Partial or non-execution
Shorter Execution Horizon	Converge rapidly toward target exposure	Lower timing and completion risk	Greater immediate liquidity demand
Longer Execution Horizon	Distribute liquidity demand	Lower concentration of execution	Greater exposure to price movements before completion
Higher Participation	Complete more rapidly relative to the chosen activity measure	Faster progress toward target quantity	Larger share of relevant observed or estimated market activity
Lower Participation	Spread execution across more market activity	Reduced concentration of activity	Longer execution horizon
Routing / Liquidity Seeking	Access executable liquidity across available sources	Broader opportunity set	Venue-, protocol- and routing-dependent execution uncertainty

The table describes economic trade-offs rather than rankings. No execution choice is universally superior, and actual outcomes depend on order characteristics, market structure and prevailing conditions.

Chapter Synthesis — Execution Logic Allocates Risk

Chapter 3 established that a systematic strategy faces constraints imposed by liquidity, size, volatility and time.

Chapter 4 shows that execution logic determines how the strategy chooses to operate within those constraints.

A marketable order prioritizes immediacy but accepts greater uncertainty over execution price.

A passive order increases price control while accepting non-execution risk.

Order slicing distributes execution across time but creates an incomplete exposure path while the parent order remains unfinished.

Participation-based algorithms connect execution pace to a defined measure of market activity but do not determine the economic urgency of the investment decision.

Routing and liquidity-source selection determine which available execution opportunities the process can attempt to access.

Latency influences how quickly information, decisions and orders move through that process, but its economic significance depends on strategy horizon and market structure.

Partial fills and rejected orders can leave realized exposure different from intended exposure even when the underlying model signal remains unchanged.

These are not independent operational details.

Together, they determine the execution policy through which a quantitative investment decision interacts with real markets.

Within the analytical framework of this paper:

Execution logic is a mechanism for allocating risk between price, time, liquidity and completion.

There is therefore no universally optimal way to execute an investment decision independently of the decision itself.

Execution logic must be evaluated relative to the purpose, urgency and constraints of the intended trade.

The next chapter moves one layer further outward.

It examines The Infrastructure Layer — the brokers, trading venues, liquidity sources, market-access systems, controls and technological architecture through which execution logic must operate before an intended order can become a realized market transaction.

5. The Infrastructure Layer

An execution decision does not interact with the market directly.

Between an investment decision and a realized transaction lies an infrastructure composed of software, hardware, communication networks, market data, brokers, trading venues, liquidity providers, order-management systems, risk controls and post-trade processes.

Much of this infrastructure is invisible when it works as intended.

An order is generated, transmitted, acknowledged and executed. The resulting fill can appear to be a direct consequence of the trading decision.

Economically, however, that apparent continuity can be misleading.

The infrastructure determines which markets can be reached, which order types can be submitted, which liquidity sources can be accessed, how quickly information and orders can move, which controls an instruction must pass, how execution requests are handled and whether the resulting transactions are ultimately accepted.

Regulatory frameworks for algorithmic trading recognize this explicitly. Commission Delegated Regulation (EU) 2017/589 defines the production environment of an algorithmic trading system broadly enough to include software, hardware, order routing, market data, databases, risk controls, data capture, analysis systems and post-trade processing. It also requires systems to interact correctly with trading-venue or direct-market-access infrastructure and emphasizes testing, resilience and controls (European Commission, 2017).

These requirements are specific to the relevant European regulatory framework and are not presented here as universal obligations.

Their analytical significance is broader:

Execution infrastructure is not a neutral conduit. It defines part of the environment within which an execution strategy can operate.

This chapter examines that environment.

5.1 The Execution Architecture

The model-to-market process described in Chapter 1 can be represented conceptually as a chain:

INVESTMENT DECISION

→ ORDER GENERATION

→ EXECUTION SYSTEM

→ MARKET ACCESS

→ BROKER / VENUE / LIQUIDITY SOURCE

→ EXECUTION

→ CONFIRMATION AND RECORDING

Each stage may contain several technological or institutional components.

A production execution environment can include:

- portfolio-management or decision systems;
- order-management systems;
- execution-management systems;
- market-data feeds;
- risk-control systems;
- broker connectivity;
- venue or liquidity-provider interfaces;
- communication networks;
- databases;
- monitoring and logging systems;
- and post-trade processing.

The precise architecture varies significantly across firms and markets.

A proprietary trading firm with direct exchange membership may connect to a venue differently from an asset manager accessing markets through a broker.

A systematic FX strategy may encounter dealer liquidity, electronic platforms and aggregation technology rather than a centralized exchange order book.

An institutional equity execution workflow may involve an investment manager, broker algorithm, smart order router and multiple trading venues.

The purpose of distinguishing these components is not to prescribe a particular architecture.

It is to recognize that an investment decision must pass through an execution stack, and each layer can impose capabilities, constraints and dependencies on the process.

The European regulatory framework provides a useful institutional example. Its definition of an algorithmic trading production environment includes not only the algorithm itself, but also the software and hardware used by traders, order routing, market data, dependent databases, risk-control systems, data capture, analytical systems and post-trade processing.

That breadth reflects an important principle:

The economic implementation of an algorithm cannot be evaluated entirely independently of the environment in which the algorithm operates.

A correctly specified execution rule can still produce an unintended outcome if the surrounding infrastructure does not process its instructions as expected.

Infrastructure is therefore not simply a technology concern.

It is a distinct risk layer within model-to-market analysis, and failures or constraints at that layer can subsequently affect execution outcomes.

5.2 Trading Venues and Liquidity Sources

A trading instruction ultimately requires executable interest on the other side of the transaction.

Market infrastructure determines how that executable interest is accessed.

Depending on market structure, this may occur through:

- an exchange or multilateral matching process;
- a broker;
- bilateral dealer liquidity;
- an electronic platform;
- or another eligible execution arrangement.

In exchange-based markets, orders may interact through centralized or venue-specific matching systems.

In markets with multiple exchanges or trading venues, executable liquidity may be distributed across several order books.

In dealer markets, liquidity may instead be supplied bilaterally or through electronic platforms.

In OTC FX, liquidity can be distributed across dealer platforms, multidealer venues, electronic communication networks, aggregators and proprietary liquidity pools.

The execution process must therefore interact not only with a price, but with a specific mechanism through which that price can potentially become executable.

This distinction matters because two quoted execution opportunities need not be economically identical.

They may differ in:

- accessible quantity;

- execution protocol;
- response time;
- likelihood of acceptance;
- order-handling rules;
- available order types;
- counterparty structure;
- and the way the venue or liquidity source responds as market conditions change.

Chapter 4 considered routing and venue selection as execution decisions.

The infrastructure perspective asks a different question:

What execution opportunities does the available market-access architecture make reachable in the first place?

The distinction can be significant.

An execution algorithm cannot route to a venue to which the relevant infrastructure is not connected.

It cannot use an order type that the broker or venue does not support.

It cannot consume liquidity that is not visible or accessible through its execution arrangements.

Within the framework of this paper, infrastructure therefore partly defines the feasible execution set.

This does not imply that connecting to more venues necessarily improves execution.

Additional connectivity can increase complexity, data requirements, operational dependencies and the number of market protocols that must be understood and monitored.

The relevant objective is not maximal connectivity.

It is access to an execution environment appropriate to the strategy and the markets in which it operates.

5.3 Fragmented and OTC Markets

The infrastructure question becomes particularly important when liquidity is fragmented.

The foreign-exchange market provides a useful example.

Unlike a market organized around one consolidated central limit order book, electronic FX liquidity is distributed across multiple venues, dealer relationships and liquidity pools.

The BIS Markets Committee describes the FX market as highly fragmented and documents how the increasing use of execution algorithms has developed alongside this fragmentation (BIS Markets Committee, 2020).

This structure creates both opportunity and complexity.

A sophisticated execution process can access multiple liquidity sources and divide orders across venues over time.

That can broaden the set of available execution opportunities.

But fragmentation also means that there may be no single observable market state.

Different platforms may display different prices.

Different liquidity providers may offer different quantities.

A liquidity pool accessible to one participant may not be available to another.

Some activity may occur inside dealers' proprietary pools and therefore be less visible to the wider market.

The BIS notes that increased internalisation and fragmented liquidity can reduce the visibility of transactions and complicate the measurement of liquidity (BIS Markets Committee, 2020).

This has an important implication for systematic execution.

In a fragmented market, the observable and accessible market environment of a particular execution system is partly determined by its connectivity and liquidity relationships.

This does not mean that market prices are arbitrary.

It means that the observable and executable opportunity set may vary between participants.

Participants seeking to trade the same currency pair at approximately the same time may therefore encounter different combinations of dealers, venues, pricing streams, available quantities and execution protocols.

Infrastructure can consequently affect:

- the liquidity visible to the strategy;
- the liquidity accessible to the strategy;
- the sequence in which that liquidity is reached;
- and the information available for evaluating execution.

The BIS also identifies a broader risk-transfer effect (BIS Markets Committee, 2020).

FX execution algorithms can shift execution risk from dealers toward end users, requiring users themselves to manage a greater share of timing and implementation risk.

Automation can therefore alter who bears execution risk, rather than eliminating the risk itself.

This is a recurring theme of the model-to-market framework:

Technology can transform the location and management of execution risk without making that risk disappear.

5.4 Pre-Trade Risk Controls

Before an order interacts with market liquidity, it may have to pass through one or more risk-control layers.

These controls serve an important purpose.

Automated trading systems can generate and transmit orders at speeds and volumes that make purely manual intervention insufficient as the primary means of preventing erroneous or excessive orders.

In regulated algorithmic and electronic market-access environments, pre-trade controls can therefore form a mandatory part of the order path.

EU RTS 6 and U.S. SEC Rule 15c3-5 provide two jurisdiction-specific examples.

Commission Delegated Regulation (EU) 2017/589 requires investment firms within its scope to maintain pre-trade controls including: (European Commission, 2017)

- price collars;
- maximum order values;
- maximum order volumes;
- and other mechanisms intended to prevent inappropriate orders from entering the market.

In the United States, SEC Rule 15c3-5 similarly requires broker-dealers providing market access to maintain financial and regulatory risk-management controls. These include mechanisms designed to prevent orders exceeding specified credit, capital, price or size parameters from reaching exchanges or alternative trading systems (U.S. Securities and Exchange Commission, 2010).

These examples apply within their respective jurisdictions.

They nevertheless demonstrate a general architectural fact:

An order can be executable from the perspective of the investment strategy but inadmissible from the perspective of the market-access infrastructure.

This creates a distinction between investment intent and permitted order flow.

A model may request a particular position.

An execution algorithm may generate the corresponding order.

But an intermediate control layer may determine that the instruction should not proceed.

That outcome should not automatically be interpreted as a failure.

A properly functioning risk control may be doing exactly what it was designed to do.

From the perspective of realized strategy exposure, however, the consequence remains economically relevant.

The position intended by the model may not be obtained, or may be obtained only after the order has been modified, resubmitted or delayed.

Pre-trade controls therefore demonstrate how infrastructure and execution outcomes can interact.

The immediate cause of a blocked order may belong to the infrastructure or operational layer.

The resulting economic effect may appear as:

- delay;
- partial implementation;
- non-completion;
- or divergence between intended and realized exposure.

The causal category and the economic consequence should therefore remain analytically distinct.

5.5 Rejections, Cancellations and Execution Uncertainty

Not every order that leaves an execution algorithm becomes a trade.

Chapter 4 distinguished partial fills, non-fills and rejections.

The infrastructure layer adds another dimension: the rules governing whether an instruction remains valid and executable can differ across brokers, venues and market protocols.

An order can fail to become a transaction because:

- a pre-trade risk control blocks it;
- the venue rejects the order format or parameters;
- the available liquidity disappears before execution;
- the order expires;
- the execution system cancels or replaces it;
- or a market-specific protocol permits a liquidity provider to accept or reject a trade request.

The final example is particularly relevant in electronic FX.

Last Look

Certain FX trading arrangements use a process known as last look.

Under Principle 17 of the FX Global Code, last look is treated as a risk-control mechanism through which a liquidity provider may perform defined validity and/or price checks before accepting or rejecting a trade request (GFXC, 2024).

The Global Foreign Exchange Committee's guidance emphasizes that last look should be applied in a fair and predictable manner, used for those stated checks rather than unrelated purposes, and supported by appropriate ex-ante disclosure (GFXC, 2021b).

The FX Global Code, in its December 2024 version, is a global set of principles of good practice (GFXC, 2024).

It does not create legal or regulatory obligations.

Its significance for this paper is therefore market-structural rather than legal.

Last look illustrates a fundamental execution point:

A displayed or requested price does not always imply an irrevocable transaction.

Depending on the applicable protocol, execution may remain subject to an acceptance process.

This introduces an additional distinction between:

PRICE OBSERVATION

→ TRADE REQUEST

→ VALIDITY / PRICE CHECK

→ ACCEPTANCE OR REJECTION

→ EXECUTION

A systematic execution process operating in such an environment must therefore account not only for quoted price and latency, but also for the probability and timing of acceptance or rejection.

Again, rejection should not automatically be equated with poor execution.

A liquidity provider may reject a request according to a disclosed protocol.

A risk system may reject an order because it violates valid controls.

A venue may reject an instruction because it does not conform to permitted parameters.

The economic question is different:

What does the rejection do to the strategy's realized implementation?

It may require resubmission.

It may create delay.

The price available on the next attempt may differ.

Or the intended quantity may remain unexecuted.

Infrastructure rules can therefore influence execution outcomes even when every component operates according to its stated design.

5.6 Operational Resilience and Connectivity

Execution infrastructure must function not only in normal conditions but also when markets, systems or message flows become stressed.

This makes operational resilience part of the model-to-market problem.

Potential disruptions include:

- failure of market-data feeds;
- network or connectivity interruption;
- software malfunction;
- database or dependency failure;
- excessive message or order flow;
- failure of risk controls;
- incorrect processing of venue messages;
- and inability to cancel or disable an algorithm when required.

The purpose of listing these possibilities is not to suggest that such events are common.

It is to show that execution depends on a chain of technological processes, and the reliability of that chain influences whether an intended order can reach and interact with the market as designed.

European algorithmic-trading regulation provides a particularly explicit example.

Commission Delegated Regulation (EU) 2017/589 requires conformance testing designed to verify that trading systems communicate and interact properly with trading venues or direct-market-access providers and that market-data flows are processed correctly (European Commission, 2017).

It also requires testing designed to ensure that algorithmic systems continue to function effectively under stressed market conditions.

Business-continuity arrangements and the ability to shut down trading algorithms without creating disorderly market conditions form part of the same framework.

ESMA's supervisory briefing on algorithmic trading of 26 February 2026 reinforces the continuing importance of these issues (ESMA, 2026).

The briefing focuses particularly on:

- governance;
- testing;
- outsourcing;
- and pre-trade controls.

It is a non-binding supervisory-convergence tool rather than a new set of legal requirements.

Its significance lies in showing that these issues remain central to contemporary supervision of algorithmic trading infrastructure.

Operational resilience also extends beyond internally controlled systems.

Execution architectures may depend on:

- brokers;
- software vendors;
- cloud or hosting infrastructure;
- market-data providers;
- telecommunications providers;
- direct-market-access providers;
- and trading venues.

Outsourcing can relocate an operational dependency without eliminating it.

Within the EU regulatory framework, outsourcing or procuring relevant software or hardware also does not transfer the investment firm's regulatory responsibility for compliance (European Commission, 2017).

ESMA's 2026 supervisory briefing further emphasizes the need for clear allocation of roles and responsibilities and for firms to retain sufficient oversight and control over outsourced algorithmic-trading arrangements (ESMA, 2026).

Operational dependency and regulatory responsibility are therefore related but distinct concepts.

From a broader execution perspective, the implication is straightforward:

The reliability of the model-to-market process depends on the complete chain through which orders and information travel, not only on the code that generates the trading decision.

5.7 Infrastructure and Market Functioning

Infrastructure should not be treated only as a source of failure.

Automation and algorithmic trading can, in particular market structures, improve measured liquidity and the efficiency with which trading interest is matched.

This is important because a one-sided discussion of execution infrastructure would produce a misleading conclusion.

Hendershott, Jones and Menkveld examine the effect of increased algorithmic trading on liquidity following the automation of quote dissemination on the New York Stock Exchange (Hendershott et al., 2011).

Using that market-structure change as an instrument for algorithmic trading, they find that for large stocks in their sample increased algorithmic trading narrowed spreads, reduced adverse selection and improved quote informativeness.

The finding is important precisely because it is specific.

It does not demonstrate that every form of automation improves liquidity in every market.

It demonstrates that algorithmic and electronic infrastructure can, under particular market structures and conditions, improve market-quality measures rather than merely introduce new technological risks.

The BIS reaches a similarly balanced conclusion in FX (BIS Markets Committee, 2020).

It concludes that FX execution algorithms can improve overall market functioning by increasing the efficiency of matching between liquidity providers and consumers in a highly fragmented market.

At the same time, the BIS identifies new risks, including shifts in the location of execution risk and reduced transaction visibility associated with internalisation (BIS Markets Committee, 2020).

The relevant conclusion is therefore not:

More automation is better.

Nor is it:

More automation creates more risk.

A more defensible conclusion is:

Automation changes the mechanism through which liquidity is accessed, execution risk is allocated and market participants interact. Its economic effects depend on the architecture and market environment in which it is deployed.

5.8 Why the Same Intended Trade Can Produce Different Executions

The infrastructure layer allows the model-to-market argument to be stated more precisely.

Consider a conceptual comparison in which two execution processes receive the same intended trade at the same reference time.

Both seek the same instrument.

Both seek the same quantity.

Both apply the same economic urgency.

At the investment-decision level, the intended trade is the same.

Their realized executions nevertheless need not be.

One execution process may have access to several liquidity sources while another reaches only one.

One broker may support a particular order type or execution protocol that another does not.

One system may reach the market earlier.

One execution path may encounter a pre-trade control.

One market-access arrangement may route to a different venue.

In FX, one liquidity relationship may operate under a different acceptance protocol from another.

Different infrastructure can therefore expose the same intended trade to different opportunity sets and execution paths, which may result in different:

- fill probabilities;
- execution timings;
- realized quantities;
- and realized prices.

This is a conceptual implication of the evidence reviewed in this and previous chapters.

RP-001 does not claim a universal empirical magnitude for these differences.

The important point is structural:

The same intended trade does not uniquely determine a realized execution independently of the execution environment through which it is implemented.

Between investment intention and fill lies an execution environment.

That environment defines what is observable, what is accessible, what is permitted and what is operationally possible.

This strengthens the distinction introduced in Chapter 1 between model performance and realized strategy performance.

The model generates an intention.

Execution logic determines how the strategy attempts to implement it.

Infrastructure defines part of the environment through which that attempt must pass.

The realized transaction emerges from the interaction of all three.

Chapter Synthesis — Infrastructure Defines the Feasible Execution Set

Previous chapters have examined execution primarily as an economic and decision problem.

This chapter adds the institutional and technological layer through which those decisions must operate.

The infrastructure layer determines:

- which markets and liquidity sources can be reached;

- how orders are transmitted and processed;
- which controls they must pass;
- which execution protocols apply;
- how systems respond under stress;
- and whether intended instructions ultimately become executable transactions.

Some infrastructure constraints are protective.

Pre-trade controls can prevent erroneous or excessive orders from reaching markets.

Testing and resilience requirements can reduce the risk that automated systems behave unexpectedly.

Some infrastructure developments broaden execution opportunity.

Electronic venues, aggregation and algorithmic execution can improve access to fragmented liquidity and, in specific empirical settings, improve market-quality measures.

Other characteristics create new forms of uncertainty.

Fragmentation can reduce visibility.

Acceptance protocols can create execution uncertainty.

Connectivity and external dependencies can introduce operational risk.

The relevant analytical conclusion is therefore not that infrastructure is good or bad.

It is that:

Infrastructure defines part of the feasible execution set through which a systematic strategy can become economically real.

This also clarifies the relationship among the risk categories introduced in Chapter 1.

An infrastructure failure or control action may originate as Infrastructure / Operational Risk.

Its consequences may then appear through execution as delay, rejection, partial completion or altered implementation conditions.

Those consequences can create a gap between intended and realized exposure.

The resulting position remains exposed to Market Risk while the implementation process continues.

The categories are analytically distinct even when their economic consequences interact.

Infrastructure does not determine the investment intention.

It defines part of the feasible environment through which that intention can become a realized transaction.

The next chapter turns from execution mechanics to measurement.

It examines Measuring What Actually Happened — why execution requires a benchmark, how implementation shortfall can be decomposed, why alternative benchmarks answer different questions and how transaction-cost analysis can be used to monitor the model-to-market transition.

6. Measuring What Actually Happened

Execution creates an observable outcome. Orders are submitted. Some are filled. Others are partially executed, rejected, cancelled or left incomplete. Prices change while implementation is underway. Explicit costs are incurred. The realized position may differ from the position originally intended.

Yet none of these observations, taken alone, answers a fundamental question: Was the execution economically consistent with the objective of the investment decision?

Answering that question requires measurement. Measurement, in turn, requires a reference point. An execution price cannot be described as favourable or unfavourable without specifying the price against which it is being compared. A broker cannot be judged solely by an average fill price without considering the orders, market conditions and objectives associated with those fills. A venue cannot be evaluated solely on quoted prices if execution probability, available quantity and timing are economically relevant.

Execution measurement is therefore not merely a process of collecting transaction data. It is a process of defining what comparison is economically meaningful.

Meaningful evaluation of execution quality requires a defined objective and an appropriately specified benchmark.

Different benchmarks answer different questions. Implementation shortfall asks what happened between investment decision and realized implementation. Arrival price asks how execution compared with the market when the order entered the execution process under the chosen methodology. VWAP and TWAP ask how execution compared with market prices observed over a defined interval. A closing-price benchmark may be appropriate when portfolio valuation or index tracking creates a specific economic reason to transact near the close. None provides a universal definition of good execution.

6.1 Why Execution Requires a Benchmark

Suppose a strategy buys an instrument at 101. Is that a good execution? The number alone provides no answer.

If the relevant market price was 100 when the investment decision was made, the execution may appear unfavourable relative to that decision point. If the market had already risen to 102 by the time the order was released for execution, an execution at 101 may appear favourable relative to arrival. If the volume-weighted average market price during the execution interval was 100.50, the same fill may again be evaluated differently.

The transaction has not changed. Only the benchmark has.

Institutional execution frameworks distinguish among several benchmark categories. Arrival price can be relevant when an investment manager wishes to evaluate execution relative to the market available when implementation begins; VWAP and TWAP can be relevant when execution is distributed across a defined interval; and closing prices may have direct relevance for portfolios linked to the market close (CFA Institute, 2026a).

From the analytical perspective adopted in this paper, benchmark selection should follow the investment objective rather than be chosen retrospectively to fit the observed execution outcome. Multiple benchmarks can still provide useful diagnostic information, but they answer different questions.

A benchmark does not simply measure execution. It defines the economic question against which execution is being evaluated.

6.2 Implementation Shortfall Revisited

Chapter 2 introduced implementation shortfall as a framework for comparing a paper portfolio with actual implementation. Its importance can now be stated more precisely: implementation shortfall retains the investment decision as the economic reference point.

Perold's original framework was developed to capture the difference between the performance of a paper portfolio and the actual portfolio created through real-market implementation (Perold, 1988). Contemporary institutional treatment similarly compares actual implementation with a paper portfolio based on the decision price and attributes the resulting difference to components such as delay, execution and opportunity cost (CFA Institute, 2026a).

$$\text{Implementation Shortfall} = \text{Paper Portfolio Outcome} - \text{Actual Portfolio Outcome}$$

The precise algebra and sign convention can vary. Implementation shortfall preserves movement between investment decision and order submission, movement during execution, explicit charges and intended quantity that remains unexecuted.

PAPER PORTFOLIO OUTCOME

↓

Delay Component

↓

Execution Component

↓

Explicit Trading Costs / Opportunity Cost

↓

ACTUAL PORTFOLIO OUTCOME

Figure 4. Conceptual Implementation Shortfall Attribution. RP-001 expository decomposition. Formal implementation-shortfall attribution conventions may differ, and the execution component should not be interpreted as a pure causal estimate of market impact.

The final qualification is essential. An observed difference between the decision benchmark and realized execution does not automatically reveal why that difference occurred. The execution component may reflect both buying or selling pressure associated with the order and market movements occurring during the execution interval.

Implementation shortfall therefore measures an economic difference. It does not perfectly identify the causal source of every part of that difference.

Measurement can identify the size of the gap more readily than it can establish the precise cause of the gap.

A fully hypothetical numerical illustration of this decomposition is provided in Appendix B.

6.3 Alternative Execution Benchmarks

Arrival Price

An arrival-price benchmark refers to a market reference observed when the order is released for execution, with the exact timestamp and price convention defined by the applicable methodology.

Conceptually, it asks how execution performed relative to the market available when implementation began. This can be useful for urgent trades, but it excludes at least part of the decision-to-arrival interval from the immediate execution comparison.

Volume-Weighted Average Price

$$VWAP = \Sigma(P_t V_t) \,/\, \Sigma V_t$$

VWAP compares execution with the volume-weighted average transaction price over a defined interval. The market universe, transaction set and interval used in the calculation must be specified.

Its interpretation becomes particularly sensitive in fragmented or OTC markets where comprehensive market-wide transaction volume may not be observable. A VWAP from one venue or data set should not automatically be interpreted as a market-wide benchmark.

Time-Weighted Average Price

$$TWAP = (1/N)\ \Sigma P_t$$

TWAP averages selected market-price observations through time rather than weighting them by volume. The precise value depends on observation frequency, sampling convention and defined interval. It therefore answers a time-distribution question rather than a liquidity-participation question.

Closing Price

A closing-price benchmark compares execution with a specified end-of-session or official closing price. It can have direct economic relevance for portfolios whose valuation or tracking objective is linked to the close, while being economically arbitrary for other strategies.

A benchmark is appropriate when it reflects the economic objective of the trade, not merely because it is easy to observe.

Benchmark	Primary Question
Decision Price / Implementation Shortfall	What happened to the investment decision from decision to realized implementation?
Arrival Price	How did execution perform relative to the market reference when implementation began?
VWAP	How did execution compare with volume-weighted market trading during the defined interval and market universe?
TWAP	How did execution compare with average sampled prices through the selected interval?
Close	How did execution compare with the market's specified closing reference price?

6.4 Transaction Cost Analysis

A benchmark comparison becomes more useful when it is applied systematically across a body of execution data. This is the role of Transaction Cost Analysis, or TCA.

TCA can be understood as structured analysis of execution outcomes relative to relevant benchmarks, order characteristics and market conditions. Institutional frameworks use it to evaluate execution quality, brokers, algorithms, venues and trading arrangements (CFA Institute, 2026a).

In FX, the GFXC TCA Data Template provides a standardized framework for parent- and child-order data that algorithmic execution providers can make available to clients for transaction-cost analysis (GFXC, 2021c). The template is voluntary and is not formally part of the FX Global Code. Standardisation can facilitate comparison but does not by itself make executions economically comparable when objectives and market conditions differ.

Depending on the analytical objective, useful data may include decision or order timestamps, parent and child orders, quantities, execution prices, benchmarks, cancellations, residual quantities, venues or liquidity sources, rejection or acceptance information and market conditions.

Execution measurement is constrained by the data architecture of the execution process itself.

6.5 Execution Monitoring

Execution monitoring can include individual-order analysis and recurring evaluation across many executions. The objective is not simply to label one trade favourable or unfavourable, but to identify persistent patterns that warrant further investigation.

Metric	Main Analytical Question	Important Qualification
Implementation Shortfall	What was the total economic gap between decision and implementation?	Does not by itself establish causal attribution
Arrival-Price Slippage	How did execution compare with the market reference when implementation began?	Excludes earlier decision-to-arrival effects
Effective Spread	How did the realized transaction compare with the contemporaneous quoted midpoint?	Less complete for large multi-part executions

Metric	Main Analytical Question	Important Qualification
Fill Rate	How much of the intended executable quantity was completed?	High fill rate does not imply favourable price
Unexecuted Quantity	What target exposure was not obtained?	Economic consequence depends on subsequent market path
Execution Duration	How long did implementation take?	Faster is not inherently better
Participation Measure	How large was the strategy relative to the defined market-activity reference?	Reference activity may be observed or estimated
Rejection Rate / Frequency	How often did eligible submitted orders or trade requests fail an acceptance or control stage, and for what stated reason?	Requires a defined denominator and classification of rejection causes
Broker / Venue / Algo Analysis	Do outcomes differ across execution channels or methods?	Requires conditioning on relevant order and market characteristics

Table 4. Execution Monitoring Framework. RP-001 analytical framework for contextual evaluation of execution outcomes. No universal performance thresholds are implied.

Execution metrics require conditioning on economically relevant trade and market characteristics before comparative conclusions are drawn. The same metric should be interpreted in relation to trade objective, order size, instrument, liquidity conditions, volatility, execution horizon, participation, market structure and benchmark selection.

6.6 Evaluating Brokers, Venues and Execution Algorithms

Execution analysis can be aggregated across brokers, venues, liquidity providers, algorithms and routing arrangements, but raw rankings can be misleading because execution opportunities are not allocated randomly. One broker may receive larger or more difficult orders, one venue may be used disproportionately during volatile periods, and one algorithm may be selected specifically for urgent trades.

Comparative analysis should therefore, where data permit, condition or stratify results by economically relevant order and market characteristics.

The GFXC Algo Due Diligence Template provides a standardized framework through which users can obtain information about algorithmic execution services, while the TCA Data Template supports more standardized execution data (GFXC, 2021a; GFXC, 2021c). Neither template prescribes which provider is superior.

MiFID II Article 27 provides a jurisdiction-specific regulatory example of multidimensional execution assessment: for firms within scope, factors such as price, costs, speed, likelihood of execution and settlement, size and nature can be relevant to achieving the best possible result (European Parliament & Council of the European Union, 2014).

Professional execution governance treats execution arrangements as something that must be monitored and reassessed rather than selected once and assumed to remain optimal indefinitely.

6.7 From Measurement to Feedback

Measurement can remain descriptive, but it can also create information for subsequent implementation decisions. Consider an illustrative monitoring result in which a particular class of orders repeatedly experiences larger shortfall when participation exceeds a defined range. That observation does not prove the cause; it creates a research question.

The next analysis may examine order size, liquidity, urgency, benchmark choice, volatility, broker or venue effects, or the rate at which the underlying signal decays. Potential responses may include reassessing execution horizon, participation logic, order size, broker selection, venue access, algorithm choice, benchmark selection or monitoring thresholds.

At a broader level, persistent execution evidence may become relevant to portfolio construction or strategy capacity. This does not imply that TCA automatically improves performance or that historical relationships will persist.

Measurement can convert realized execution outcomes into structured information that informs subsequent implementation decisions.

The formal integration of that information into strategy design is developed in Chapter 7.

Chapter Synthesis — A Benchmark Defines the Question

Execution cannot be meaningfully measured without specifying the reference against which it is judged. Implementation shortfall, arrival price, VWAP, TWAP and closing-price benchmarks answer different economic questions.

TCA extends those comparisons across orders and execution channels, while monitoring asks whether outcomes reveal patterns that justify further investigation. But a benchmark can identify a difference without proving its cause, and standardized data can improve comparability without eliminating differences in trade objectives or market conditions.

Execution measurement is not simply a scorekeeping exercise. It is a structured comparison between intended objectives and realized implementation.

7. Execution as Part of Systematic Strategy Design

The preceding chapters have treated execution as an economic process, a liquidity problem, a set of implementation choices, an infrastructure dependency and a measurement problem. Taken together, these observations raise a broader question: should execution begin only after a systematic strategy has been designed, or should implementation constraints influence the design of the strategy itself?

If execution affects the price, timing and quantity at which target exposures become realized positions, assumptions about implementation can influence which signals are economically actionable, how urgently positions should be established, how frequently a portfolio can rebalance, how large positions can become, how much capital a strategy can deploy and how theoretical performance should be interpreted.

Institutional execution frameworks already link trading strategy to motivation, urgency, order characteristics and market conditions, and connect expected alpha decay with execution urgency (CFA Institute, 2026a). RP-001 takes this logic one step further.

Execution should be treated not merely as the final operational step of a systematic strategy, but as one of the constraints that define its economically realizable form.

This does not imply that execution replaces model design or that every model requires a complex microstructure simulation. The narrower proposition is that a strategy intended for real-market deployment should be evaluated under assumptions consistent with the environment in which its decisions must actually be implemented.

7.1 Execution-Aware Model Design

Quantitative opportunities differ in how difficult they are to implement. One signal may remain economically relevant for several days; another may decay over minutes or hours. Institutional execution research supports treating expected alpha decay and trading urgency as related considerations (CFA Institute, 2026a).

This evidence supports treating signal horizon and feasible execution horizon as related economic considerations. It does not imply a universal mapping between the two.

From the analytical perspective of RP-001, if the time required for realistic implementation becomes material relative to the expected economic life of a signal, theoretical and implementable opportunity should no longer be assumed to be equivalent.

> *Can the opportunity identified by the model remain economically relevant long enough to be implemented under realistic market conditions?*

7.2 Turnover Versus Expected Alpha

Higher turnover is not inherently undesirable. If repeated trading captures sufficiently valuable opportunities, the benefit may justify the implementation cost. But every rebalance creates another transition across the model-to-market boundary.

The relevant comparison is therefore the expected economic value of a portfolio adjustment versus the cost and risk of implementing it. RP-001 refers to this comparison as an economic implementation test; the term is an analytical concept rather than a standardized industry measure.

Patton and Weller provide an empirical reason for making the distinction between theoretical profitability and practical implementability, while not establishing a universal relationship between turnover and realized performance (Patton & Weller, 2020).

A signal can appear economically attractive in a theoretical portfolio while offering substantially less implementable value if the expected benefit of trading is small relative to plausible implementation frictions. The objective is therefore not to minimize turnover mechanically, but to determine whether portfolio changes survive an economic implementation test.

7.3 Liquidity-Aware Position Sizing

Position sizing is commonly associated with expected return, volatility, diversification, concentration and other portfolio objectives. Once a position must be implemented, size also becomes an execution variable.

For purposes of RP-001, liquidity-aware position sizing means considering the market-implementation requirements created by a target position alongside its portfolio-level risk and return characteristics.

A larger desired position may represent a greater share of available market activity, require a longer execution horizon and/or higher participation depending on execution policy, interact with additional layers of liquidity or increase the period during which realized exposure remains away from target.

Execution conditions need not scale proportionally with desired position size. Kyle and Obizhaeva link transaction-cost characteristics to the scale of risk transfers, volume and volatility, while Capponi, Menkveld and Zhang demonstrate theoretically that large-order execution can depend jointly on size and duration when liquidity supply responds endogenously (Kyle & Obizhaeva, 2016; Capponi et al., 2025). Neither provides a universal position-sizing rule.

Position size and implementation conditions cannot always be treated as independent variables.

7.4 Capacity-Aware Portfolio Construction

Chapter 3 defined strategy capacity for RP-001 as the scale of capital or exposure that can be implemented while maintaining execution conditions consistent with the strategy's economic objective. This creates an analytical distinction between Model Scalability and Implementation Scalability.

A theoretical portfolio can often be scaled mechanically without changing its logic. The market does not necessarily scale in the same way. A larger portfolio can produce larger orders, greater order size relative to market activity, more turnover in currency terms, longer implementation periods and greater participation requirements.

Within the RP-001 framework, capacity-aware portfolio construction asks not only how much capital a model can allocate, but whether the required portfolio transitions remain implementable under conditions consistent with the strategy's economic objective.

Diversification across instruments does not necessarily diversify execution demand if several portfolio changes must be implemented during the same market window or are exposed to common liquidity conditions.

> *Portfolio construction intended for real-market deployment should consider not only the desirability of the target portfolio, but the feasibility of transitioning from the current portfolio to that target.*

7.5 Execution-Aware Backtesting

A backtest necessarily represents real-market implementation through assumptions: when information becomes available, when the strategy can act, the price at which it is assumed to execute, spreads, explicit charges, executable quantity, market volume, delay, partial fills, market impact and the treatment of orders that could not have been completed.

Execution-aware backtesting is used in RP-001 as an analytical principle rather than a standardized methodology. Its purpose is to make implementation assumptions explicit and to examine how sensitive a strategy's theoretical economics are to plausible execution constraints.

A fixed transaction-cost deduction can be useful for preliminary analysis but does not by itself make a backtest execution-aware when costs and constraints vary with order characteristics or market state.

Detailed historical reconstruction is also limited by the counterfactual problem. Historical market or order-book information may be incomplete, and even with detailed data the liquidity that would have been available had the hypothetical strategy actually traded cannot generally be observed directly because the strategy's own orders could have altered the market state.

Execution-aware backtesting should therefore not be interpreted as a reconstruction of the live outcome that would certainly have occurred. Within RP-001, it is a sensitivity framework for testing whether a strategy remains economically plausible under explicitly stated implementation assumptions.

- alternative spread assumptions
- alternative execution delays
- limits on participation
- limits on executable quantity
- partial-fill assumptions
- more conservative execution prices
- varying transaction costs across volatility or liquidity states
- reduced strategy scale

These are examples of analytical tests, not universal requirements. If a strategy's economic conclusion changes materially under modest changes in plausible implementation assumptions, that sensitivity is itself relevant information.

> *This is how the model behaves under the specified assumptions about how its decisions could have been implemented.*

7.6 The Execution Feedback Loop

Chapter 6 established that realized execution produces measurable evidence. Within the analytical framework of RP-001, that evidence can create a feedback loop.

MODEL / SIGNAL DESIGN
PORTFOLIO CONSTRUCTION
EXECUTION DESIGN
MARKET IMPLEMENTATION
TCA & EXECUTION MONITORING
LIQUIDITY / CAPACITY ASSESSMENT

↺ FEEDBACK TO MODEL / PORTFOLIO / EXECUTION DESIGN

Figure 5. Execution Feedback Loop. RP-001 analytical framework illustrating how realized implementation evidence can inform subsequent research, portfolio construction and execution decisions. The framework does not imply that such feedback automatically improves future investment performance.

Execution Design. Monitoring may indicate that certain order types, execution horizons or participation levels are associated with different outcomes under particular market conditions.

Broker and Venue Selection. Repeated analysis may reveal conditional differences associated with brokers, venues or liquidity sources after relevant order and market characteristics are considered.

Position Sizing. If larger positions are repeatedly associated with implementation characteristics inconsistent with the strategy's objective, sizing assumptions may warrant reassessment.

Capacity. If increasing capital is associated with materially different execution duration, participation or implementation shortfall, capacity assumptions may require reconsideration.

Model Design. Persistent implementation difficulties may prompt reassessment of the signal horizon, the assumed implementation process or the economic usefulness of pursuing the signal under the relevant constraints. Execution evidence alone does not identify which explanation is correct.

Feedback therefore requires interpretation rather than automatic optimization. Realized execution generates information about the feasibility of translating model decisions into actual positions, and that information can become an input into subsequent strategy research.

7.7 The Integrated Model-to-Market Framework

The argument developed across Chapters 1–7 can now be consolidated into several analytically distinct but interacting layers. The risk labels identify the primary analytical lens associated with each layer and are not mutually exclusive: an event originating in one layer can generate economic consequences in another.

Layer	Core Question	Analytical Role
Layer 1 — Model	What exposure does the strategy intend to obtain?	Primary analytical risk lens: Model Risk
Layer 2 — Portfolio Construction	How should signals be converted into target positions and portfolio exposures?	Expected opportunity, risk, concentration, turnover, liquidity and implementation capacity
Layer 3 — Execution Design	How should the intended change in exposure be implemented?	Primary analytical risk lens: Execution Risk
Layer 4 — Infrastructure	Through which brokers, venues, liquidity sources, systems and controls can the execution decision operate?	Primary analytical risk lens: Infrastructure / Operational Risk

Layer	Core Question	Analytical Role
Layer 5 — Market Interaction	What prices, quantities and liquidity conditions are encountered while implementation occurs?	Execution-related effects and Market Risk interact; causal separation is important
Layer 6 — Realized Exposure	What position was actually obtained, at what price, at what time and in what quantity?	Cumulative result of the preceding model-to-market process
Layer 7 — Measurement and Feedback	How did actual implementation compare with the economic objective?	Analytical layer, not an additional source of economic exposure

MODEL → TARGET PORTFOLIO → EXECUTION DESIGN → INFRASTRUCTURE → MARKET INTERACTION → REALIZED EXPOSURE → MEASUREMENT ↺ FEEDBACK

The framework is an RP-001 analytical synthesis rather than a standardized industry model. Its purpose is to make explicit that a systematic strategy exists economically not only as a rule for generating target positions, but as a process through which those positions can be created in real markets.

Chapter Synthesis — Execution Is Part of the Realized Strategy

A model produces a target exposure. Realizing that exposure requires interaction with liquidity, time, order size, execution logic, infrastructure and changing market conditions. These interactions can alter execution price, timing, realized quantity, exposure path, implementation cost and strategy capacity.

Execution is not merely the final step of a systematic strategy. It is part of the process that determines what the strategy becomes economically.

On the basis of the evidence and framework developed in RP-001, the analytical implications are to consider signal horizon relative to feasible execution horizon; assess turnover relative to the economic value of the portfolio changes it generates; recognize liquidity and implementation constraints in position sizing; evaluate capacity as an interaction between scale, portfolio transitions and market structure; make material implementation assumptions explicit in backtests; and use realized execution evidence as input to subsequent portfolio and execution decisions without treating it as automatically causal or predictive.

8. Conclusion

Systematic trading is often described through the quality of its models. Models identify relationships, generate signals and determine target exposures. But a model does not create a market position by itself.

Between an intended trade and a realized position lies an implementation process shaped by transaction costs, liquidity, order size, execution timing, market structure, infrastructure and changing market conditions.

RP-001 has examined this transition as the model-to-market problem.

The central analytical conclusion is that the difference between theoretical and realized strategy outcomes cannot generally be understood as a single transaction-cost deduction applied after a model has generated its return.

Execution can alter not only price, but also timing, realized quantity, exposure path, implementation cost and strategy capacity.

Implementation shortfall provides one framework for measuring the resulting economic gap, but measurement does not by itself establish its precise causal origin. Adverse execution outcomes may reflect interaction with liquidity, broader market movements, infrastructure constraints or several mechanisms operating simultaneously.

Liquidity further matters because markets do not generally offer unlimited executable quantity at a fixed price or on a predetermined timetable. Order size, participation, volatility and execution duration can therefore change the conditions under which a strategy is implemented.

This explains why RP-001 distinguishes between model scalability and implementation scalability. A portfolio can often be scaled mathematically without implying that the required market transitions will scale proportionally.

Execution logic determines how a strategy operates within these constraints. Different order types, execution horizons, participation choices and routing approaches distribute risk differently between price, timing, liquidity demand and completion. No execution method is universally superior independently of the economic objective of the trade.

Infrastructure adds another layer. Brokers, venues, liquidity relationships, market-data systems, connectivity and risk controls determine part of what is observable, accessible, permitted and operationally possible.

The same intended trade can therefore encounter different opportunity sets and execution paths depending on the environment through which it is implemented.

Measurement makes these outcomes observable. Implementation shortfall, arrival price, VWAP, TWAP and closing-price benchmarks answer different economic questions. Within the analytical framework of this paper, benchmark selection should therefore be aligned with the economic purpose of the trade.

RP-001 accordingly treats Transaction Cost Analysis as an analytical process rather than a simple ranking mechanism. Execution evidence can identify patterns and inform subsequent research and implementation decisions without being treated as automatically causal or predictive.

This leads to the principal analytical implication of the paper.

A quantitative model produces a target exposure. An execution process attempts to transform the resulting intended trade into a realized exposure. The market determines the conditions under which that transformation can occur. Infrastructure defines part of the feasible execution environment. Measurement reveals how the realized result differed from the intended one. And that evidence can inform subsequent portfolio, execution and research decisions.

The resulting strategy is therefore more than its model.

> *Execution is not merely the final operational step of systematic trading. It is part of the process that determines what the strategy becomes economically.*

For investors evaluating systematic strategies, the framework therefore raises two related questions:

> *Does the model identify attractive opportunities?*

> *Can those opportunities be translated into the intended exposures, at the required scale and under market conditions consistent with the strategy's economic objective?*

Considering both questions provides a more complete view of the distance between model and market.

Appendix A — Glossary of Execution Terms

Arrival Price

A market reference observed when an order is released for execution. The exact timestamp and price convention depend on the applicable methodology.

Bid-Ask Spread

The difference between the quoted bid and ask prices at a specified market, venue or liquidity source. Relative to the contemporaneous midpoint, a simple execution at the ask for a purchase or at the bid for a sale implies approximately half the quoted spread before considering price improvement, order size or depth.

Capacity

Within RP-001, the scale of capital or exposure that can be implemented while maintaining execution conditions consistent with the strategy's economic objective. Capacity is treated as an analytical concept rather than a standardized industry metric.

Decision Price

The market price associated with the point at which an investment decision is made. In implementation-shortfall analysis, it can serve as the reference price for the paper portfolio.

Effective Spread

A transaction-cost measure conventionally expressed as twice the signed distance between the transaction price and the contemporaneous quoted midpoint.

Execution Horizon

The period over which an intended trade is implemented. A shorter or longer execution horizon changes the balance between immediacy, liquidity demand, timing risk and completion risk.

Execution Risk

The risk that an intended exposure is not obtained at the intended price, time or quantity. It is analytically distinct from Model Risk, Market Risk and Infrastructure / Operational Risk.

Fill

The execution of all or part of an order or accepted trade request, resulting in a realized transaction quantity.

Partial Fill

Execution of only part of the intended order quantity, leaving a residual quantity unexecuted.

Implementation Shortfall

The economic difference between a paper portfolio based on the investment decision and the actual portfolio produced through real-market implementation. Within the RP-001 expository framework, the economic gap may be discussed through delay, execution, explicit trading costs and opportunity cost, while formal attribution conventions may differ.

Last Look

In certain FX trading arrangements, a risk-control mechanism through which a liquidity provider may perform defined validity and/or price checks before accepting or rejecting a trade request.

Liquidity

The conditions under which risk can be transferred in a market. In RP-001, liquidity is treated as multidimensional and may involve spread, depth, market activity and price response rather than a single observable quantity.

Market Impact

The price effect causally attributable to an order's interaction with available liquidity. Because this is a causal concept, observed execution-price differences should not automatically be interpreted as market impact.

Marketable Order

An order that, under the relevant market or venue rules and prevailing quotes, is capable of interacting immediately with available liquidity. Marketability generally prioritizes immediacy relative to a non-marketable order, but does not guarantee execution price or complete execution.

Opportunity Cost

The economic difference associated with intended quantity that remains unexecuted relative to the relevant paper or benchmark position. Its magnitude depends on the subsequent market path.

Participation Rate / Participation Measure

The executed quantity of the strategy relative to a defined measure of relevant market activity over a specified interval. The reference activity may be observed or estimated and may be market-wide or venue-specific. Participation is not itself a direct measure of liquidity.

Realized Exposure

The position actually obtained after execution, including the realized price, timing and quantity of implementation.

Slippage

Within RP-001, the difference between the realized execution price and an explicitly stated reference price. Slippage is a measurement convention and does not by itself identify the causal source of the price difference.

Target Exposure

The position or change in position intended by the model or portfolio-construction process before real-market implementation.

Transaction Cost Analysis (TCA)

The structured analysis of execution outcomes relative to relevant benchmarks, order characteristics and market conditions. Within RP-001, TCA is treated as an analytical process rather than a simple ranking mechanism.

TWAP — Time-Weighted Average Price

A benchmark based on average market prices observed across a defined time interval. Its value depends on the observation frequency, sampling convention and selected interval.

VWAP — Volume-Weighted Average Price

A benchmark representing the volume-weighted average transaction price across a defined market universe and time interval. Its interpretation depends on the completeness and relevance of the underlying volume data.

Appendix B — Illustrative Implementation Shortfall Example

Purpose of the Example

The following example is entirely hypothetical and is included solely to illustrate the economic logic of implementation shortfall.

The prices, quantities, costs and resulting shortfall are not based on FINOVIS trading activity, do not represent empirical estimates and should not be interpreted as typical execution costs for any market, instrument or strategy.

All monetary values are expressed in hypothetical currency units.

For exposition, RP-001 separately displays delay, execution, explicit trading costs and opportunity cost associated with unexecuted quantity. Formal implementation-shortfall attribution conventions may differ.

Hypothetical Trade Setup

Assume a systematic strategy decides to purchase 10,000 units of a hypothetical instrument. At the time of the investment decision:

$$\text{Decision Price} = 50.00$$

The paper portfolio therefore assumes acquisition of the complete intended quantity at the decision price:

$$10{,}000 \times 50.00 = 500{,}000$$

For purposes of this example, the arrival-price methodology is assumed to use the market reference observed when the order is released for execution. The order is subsequently released for execution when:

$$\text{Arrival Price} = 50.10$$

8,000 units are executed;

the average execution price is 50.18;

2,000 units remain unexecuted through the selected evaluation point;

explicit trading charges total 400.

At the selected end-of-period evaluation point, the market price is:

$$\text{Evaluation Price} = 50.40$$

1. Delay Component

The market moves from the decision price of 50.00 to the arrival price of 50.10 before the order is released for execution. The delay component is calculated on the 8,000 units that are ultimately executed:

$$(50.10 - 50.00) \times 8{,}000 = 800$$

Illustrative Delay Component = 800

This measures the benchmark difference between the investment-decision point and the selected arrival point for the executed quantity. It does not, by itself, identify why the price changed.

2. Execution Component

The 8,000 executed units are purchased at an average price of 50.18 compared with an arrival-price benchmark of 50.10:

$$(50.18 - 50.10) \times 8{,}000 = 640$$

Illustrative Execution Component = 640

This represents an observed benchmark-to-execution price difference. It should not automatically be interpreted as causal market impact. The difference may reflect the strategy's interaction with liquidity, broader market movements during execution, or a combination of factors.

3. Explicit Trading Costs

Assume commissions and other directly identifiable transaction charges total:

Explicit Trading Costs = 400

For purposes of this RP-001 example, these costs are displayed separately to make their economic contribution visible.

4. Unexecuted Quantity and Opportunity Cost

The strategy intended to purchase 10,000 units but executed only 8,000. The residual intended quantity is therefore:

$$10{,}000 - 8{,}000 = 2{,}000$$

Unexecuted Quantity = 2,000 units

The market price subsequently reaches 50.40. Relative to the original decision price, the hypothetical opportunity cost associated with the unexecuted quantity is:

$$(50.40 - 50.00) \times 2{,}000 = 800$$

Illustrative Opportunity Cost = 800

This illustrates why execution risk includes a quantity dimension. Even if no adverse execution price exists for the unfilled quantity, failure to obtain the intended exposure can still have an economic consequence. Depending on the subsequent market path and the applicable sign convention, the economic effect associated with unexecuted quantity need not always be adverse.

5. Total Illustrative Implementation Shortfall

Component	Illustrative Amount
Delay Component	800
Execution Component	640
Explicit Trading Costs	400
Opportunity Cost	800
Total Implementation Shortfall	2,640

$$800 + 640 + 400 + 800 = 2{,}640$$

Illustrative Implementation Shortfall = 2,640

6. Paper Portfolio versus Actual Portfolio

For simplicity, the comparison below ignores financing, carry, taxes and other non-execution effects and considers only the implementation economics described above.

Paper Portfolio

The hypothetical paper portfolio purchases all 10,000 units at the decision price of 50.00. At the evaluation price of 50.40:

$$(50.40 - 50.00) \times 10{,}000 = 4{,}000$$

Paper Portfolio Outcome = 4,000

Actual Portfolio

Only 8,000 units are executed, at an average price of 50.18. At the evaluation price of 50.40:

$$(50.40 - 50.18) \times 8{,}000 = 1{,}760$$

After explicit trading costs of 400:

$$1{,}760 - 400 = 1{,}360$$

Actual Portfolio Outcome = 1,360

The difference is:

$$4{,}000 - 1{,}360 = 2{,}640$$

which equals the illustrative implementation shortfall calculated above.

7. Measurement Is Not Causal Attribution

The numerical consistency of the decomposition should not be confused with causal certainty. The example measures economic differences between defined reference points:

DECISION PRICE → ARRIVAL PRICE → REALIZED EXECUTION → REALIZED QUANTITY → EVALUATION POINT

It does not establish why each price movement occurred. In particular, the 640 execution component cannot be interpreted automatically as market impact caused by the hypothetical order. Likewise, the 800 delay component measures movement between two timestamps but does not establish that the execution process caused the delay or the market movement.

> *Implementation shortfall can quantify the economic gap between intended and realized implementation more readily than it can establish the precise causal origin of that gap.*

References

Academic Research

Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. The Journal of Risk, 3(2), 5–39. https://doi.org/10.21314/JOR.2001.041

Capponi, A., Menkveld, A. J., & Zhang, H. (2025). Large orders in small markets: Execution with endogenous liquidity supply. Review of Finance, 29(1), 201–239. https://doi.org/10.1093/rof/rfae036

Hasbrouck, J., & Saar, G. (2013). Low-latency trading. Journal of Financial Markets, 16(4), 646–679. https://doi.org/10.1016/j.finmar.2013.05.003

Hendershott, T., Jones, C. M., & Menkveld, A. J. (2011). Does algorithmic trading improve liquidity? The Journal of Finance, 66(1), 1–33. https://doi.org/10.1111/j.1540-6261.2010.01624.x

Keim, D. B., & Madhavan, A. (1997). Transactions costs and investment style: An inter-exchange analysis of institutional equity trades. Journal of Financial Economics, 46(3), 265–292. https://doi.org/10.1016/S0304-405X(97)00031-7

Kyle, A. S., & Obizhaeva, A. A. (2016). Market microstructure invariance: Empirical hypotheses. Econometrica, 84(4), 1345–1404. https://doi.org/10.3982/ECTA10486

Patton, A. J., & Weller, B. M. (2020). What you see is not what you get: The costs of trading market anomalies. Journal of Financial Economics, 137(2), 515–549. https://doi.org/10.1016/j.jfineco.2020.02.012

Perold, A. F. (1988). The implementation shortfall: Paper versus reality. The Journal of Portfolio Management, 14(3), 4–9. https://doi.org/10.3905/jpm.1988.409150

Institutional and Professional Sources

Bank for International Settlements, Markets Committee. (2020). FX execution algorithms and market functioning. Markets Committee Papers No. 13, 30 October 2020.

CFA Institute. (2026a). Trade strategy and execution. 2026 CFA Program Curriculum, Level III, Portfolio Management.

CFA Institute. (2026b). Trading costs and electronic markets. 2026 CFA Program Curriculum, Level III, Portfolio Management and Wealth Planning.

Federal Reserve Board. (2019). Financial Stability Report, November 2019. Box: What Has Been Happening to the Liquidity of U.S. Treasury and Equity Futures Markets?

Global Foreign Exchange Committee. (2021a). Algo Due Diligence Template. Version as of 15 July 2021.

Global Foreign Exchange Committee. (2021b). Execution Principles Working Group Report on Last Look. August 2021.

Global Foreign Exchange Committee. (2021c). Transaction Cost Analysis (TCA) Data Template. Version as of 15 July 2021.

Global Foreign Exchange Committee. (2024). FX Global Code: A set of global principles of good practice in the foreign exchange market. December 2024 version.

Regulatory and Legal Sources

European Commission. (2017). Commission Delegated Regulation (EU) 2017/589 of 19 July 2016 supplementing Directive 2014/65/EU with regard to regulatory technical standards specifying the organisational requirements of investment firms engaged in algorithmic trading. Official Journal of the European Union, L 87, 31 March 2017, 417–448.

European Parliament & Council of the European Union. (2014). Directive 2014/65/EU of 15 May 2014 on markets in financial instruments and amending Directive 2002/92/EC and Directive 2011/61/EU (recast). Official Journal of the European Union, L 173, 12 June 2014, 349–496.

European Securities and Markets Authority. (2026). Supervisory briefing on algorithmic trading in the EU. ESMA74-1505669079-10311, 26 February 2026.

U.S. Securities and Exchange Commission. (2010). Risk management controls for brokers or dealers with market access. Securities Exchange Act Release No. 34-63241, 3 November 2010, 75 Federal Register 69791.

U.S. Securities and Exchange Commission, Office of Investor Education and Advocacy. (2011). Investor Bulletin: Trading Basics — Understanding the Different Ways to Buy and Sell Stock. SEC Publication No. 141 (3/11).

Research Disclaimer

It does not contain a recommendation concerning the present or future value or price of any specific financial instrument or issuer.

Nothing in this paper should be interpreted as a recommendation or representation that any particular investment, trading strategy, execution method or technological approach is appropriate or suitable for any person or institution.

The paper discusses systematic trading, execution, liquidity, market structure, infrastructure and related topics at a general analytical level. It does not disclose or describe the proprietary trading strategies, algorithms, models, execution logic, routing methodology, position-sizing rules or other confidential intellectual property of FINOVIS.

Any numerical examples included in this publication are illustrative and hypothetical only, unless expressly stated otherwise. They do not represent actual FINOVIS trading activity, historical performance, expected performance or empirical estimates of transaction costs or execution outcomes.

Historical market behaviour and observed execution relationships may not persist. Market conditions, liquidity, volatility, execution quality, transaction costs and infrastructure performance can change materially over time and across instruments, venues, brokers and jurisdictions.

References to academic research, regulatory frameworks, industry standards and institutional publications are provided to support the analysis presented in this paper. Regulatory or legal references are jurisdiction-specific and should not be interpreted as statements regarding the applicability of any particular legal or regulatory requirement to a specific person, firm, transaction or activity.

The FX Global Code is a set of principles of good practice and does not itself impose legal or regulatory obligations or substitute for applicable regulation. Individual jurisdictions may, however, recognise, adopt or otherwise use elements of the Code within their regulatory or supervisory frameworks.

FINOVIS has sought to use authoritative and relevant sources available at the time of preparation. Information, regulatory frameworks and external publications may subsequently be amended, superseded or reinterpreted. No representation or warranty, express or implied, is made as to the continued accuracy or completeness of external information after the publication date.

Views, interpretations and analytical frameworks identified as FINOVIS or RP-001 analysis represent the analytical conclusions developed for this paper and should be distinguished from the findings or positions of the external sources cited.

Readers should conduct their own analysis and, where appropriate, obtain independent investment, legal, regulatory, tax or other professional advice before making decisions based on matters discussed in this publication.

FINOVIS RESEARCH

EXECUTION & RISK · RP-001

From Model to Market

Execution Risk in Systematic Trading

March 2026

finovis.ai