

An EDHEC Risk and Asset Management Research Centre Publication

Transaction Cost Analysis A-Z

A Step towards Best Execution in the Post-MiFID Landscape

November 2008



Sponsored by

caceis
INVESTOR SERVICES

 **NYSE Euronext.**

SUNGARD®



EDHEC-RISK
Asset Management Research

Table of Contents

Foreword	3
Introduction.....	5
I. Transaction Cost Analysis as Part of the Investment Process	9
II. Transaction Cost Components and Drivers	13
1. Direct Explicit Transaction Costs	14
2. Indirect Explicit Transaction Costs	15
3. Implicit Transaction Costs.....	19
4. Transaction Cost Magnitude.....	23
5. Transaction Costs and Execution Methods.....	25
III. Measuring Transaction Costs with Post-Trade Analysis	27
1. Benchmark Comparison.....	28
2. Implementation Shortfall.....	29
3. Common Pitfalls and Shortcomings in Transaction Cost Measurement.....	32
4. Measuring Transaction Costs under MiFID	35
IV. Estimating Transaction Costs with Pre-Trade Analysis.....	39
1. Collection and Analysis of Pre-Trade Data.....	40
2. Cost and Risk Estimation.....	41
3. Optimisation and Efficient Trading Frontier	53
V. Trading Performance Measurement.....	57
1. Transaction Costs vs. Trading Performance.....	58
2. Current State of the Industry and Practices	58
3. Regulatory Pressure: MiFID and its Best Execution Obligation	64
VI. A New Framework: the EBEX Indicators.....	67
1. General Presentation	68
2. Detailed Presentation of the Indicators.....	70
3. Illustration of the Framework.....	77
Conclusion.....	85
References	87
About the EDHEC Risk and Asset Management Research Centre.....	90
About CACEIS	94
About NYSE Euronext.....	96
About SunGard.....	98

Foreword

After nearly ten years of debate, the final implementation of MiFID is radically transforming the European capital markets landscape.

New entrants such as Chi-X, Turquoise or Equiduct, new operating models developed by major brokerage firms or former central exchanges, along with the development of advanced execution technologies such as algorithmic trading, form what can be called the MiFID revolution.

Core to the change is the obligation of best execution, which is one of the pillars the regulator has imposed for a more competitive environment. But at the time the European directives were drafted, there was no consensus on what constituted best execution; indeed, there is still no consensus.

Transaction cost analysis (TCA) lies at the very heart of the best execution obligation and it is expected to become a tool that no intermediary and market participant can ignore. The literature on TCA is abundant but it remains difficult to find an overview of TCA techniques that allows investment firms to develop a view on which of the many approaches could or should be taken.

The objective of this report is to provide a comprehensive view of what TCA is, shed light on the main underlying concepts and document the tools and techniques that have been developed in the academic and professional worlds.

This A-Z is the first step of a number of research initiatives that will make possible a better understanding of execution risk and performance and ultimately provide tools and technologies that lead to more efficient trading systems.

Far from being restricted to equity markets alone, MiFID has so far prompted reaction mainly on infrastructure related to trading in listed securities; more changes can be expected in other markets. As such, our effort will continue both in the fixed income and the listed derivatives space in the very near future.

Finally, I would like to take the opportunity to thank the partners that have made possible the creation of the 'MiFID and Best Execution' research chair hosted by the EDHEC Risk and Asset Management research Centre. CACEIS, NYSE Euronext and SunGard have committed a significant amount of time as well as financial and technical support to the team, allowing us to offer material that, we hope, will be useful to investment firms involved in the execution process.



Noël Amenc
Professor of Finance
Director of the EDHEC Risk and Asset Management
Research Centre

About the authors



Catherine D'Hondt is an Associate Professor of Finance at the Louvain School of Management and the Catholic University of Mons in Belgium. Dr D'Hondt teaches various graduate courses on financial markets and her primary research area is market microstructure, with a specific focus on traders' behaviour and order submission strategies. Within the EDHEC Risk and Asset Management Research Centre, Catherine D'Hondt contributes to the *Best Execution and Operational Performance* programme. She conducts work on the long-term implications of MiFID for European capital markets and develops expertise in transaction cost analysis and trading performance measurement. Her doctoral dissertation received the *Euronext – French Finance Association* award in 2003 and her research in the area of market microstructure and transaction cost analysis has been presented at major academic and practitioner conferences and published in refereed journals such as the *Review of Finance*, the *Journal of Asset Management* and *Finance Letters*. With Jean-René Giraud, she has co-authored the reference text on MiFID: *MiFID: Convergence towards a Unified European Capital Markets Industry* (Risk Books, 2006) and an influential position paper: *MiFID: the (in)famous European Directive?* (EDHEC, 2007). Catherine D'Hondt holds a Doctorate in Management Sciences from the Catholic University of Mons and the University of Perpignan.



Jean-René Giraud is the Director of Development of the EDHEC Risk and Asset Management Research Centre and the head of the "MiFID and Best Execution" Research Chair supported by Caceis, NYSE Euronext, and Sungard. Jean-René Giraud is in charge of developing the relationship between EDHEC and key industry players. Before joining EDHEC, he was Managing Principal at Capco (London) with responsibility for buy-side advisory services. Previously, Jean-René was a Senior Project Manager with Barclays Capital designing and implementing transaction systems, and prior to that Head of Consulting at SIP Software. As a Research Associate with EDHEC, Jean-René Giraud heads the *Best Execution and Operational Performance* programme and specialises on hedge fund operational risks and best execution. His work has appeared in refereed journals such as the *Journal of Alternative Investments* and the *Journal of Asset Management*. He has authored a large number of articles in industry publications, contributed to *Operational Risk, Practical Approaches to Implementation* and co-authored the noted *MiFID: Convergence towards a Unified European Capital Markets Industry* (Risk Books, 2006). Jean-René Giraud is a frequent speaker at major industry conferences on various topics related to operations and risk issues. He has a Master's in Information Technology from the Polytechnic School of Engineering at the University of Nice-Sophia-Antipolis.

Introduction



Introduction

From retail to more professional investors and practitioners, all are concerned with transaction costs, as it is an established fact that lower transaction costs automatically enable higher returns. To provide their clients with competitive portfolio returns, investment firms need to be proactive and put in place the appropriate means of effective transaction cost management. However, when implementing their client decisions, investment managers often cope with issues regarding how transaction costs can be properly identified, measured, forecasted as well as how the quality of execution should be evaluated. Most of the time, these fundamental questions remain open because relatively little information is directly available. In some cases, facing the multiple commercial tools that are offered, investment managers choose and apply the most popular indicators in the industry, without really knowing if they provide useful and meaningful answers to their initial concerns. Some even admit that they are sometimes confused, especially when they are offered tools based on sophisticated models resembling “black boxes”.

Transaction costs are becoming both a topical and relevant issue in Europe with MiFID—the Markets in Financial Instruments Directive¹—now in place. In the near future, the role of transaction cost analysis (TCA) is expected to grow substantially as a result of the best execution obligation. According to this new obligation, investment firms must “take all reasonable steps to obtain the best possible result when executing orders for their clients”. Although so far a mix of both principle- and rule-based regulation, this new obligation is a key element for

investor protection in a marketplace that is completely open to competition.

Best execution has consequently become a very fashionable concept. Nevertheless, because the regulator has brought neither a clear definition nor a measurable objective to make up for the current absence of consensus, this concept is not always well understood and does not mean the same thing to everyone. We can easily find evidence of this phenomenon by simply reviewing the press and the multiple occasions where liquidity providers, MTFs,² platforms, and technology vendors claim to provide “best execution”, even though there is as yet no consensus on what exactly “best execution” entails or, more importantly, while some of those platforms are not yet operational. This ongoing and widespread confusion around best execution—and, to a larger extent, around TCA as a whole—is likely to have serious side effects such as increased moral hazard, greater adverse selection and a false sense of security given to end clients. These consequences tend to emerge when a piece of regulation attempts to legislate on elements that may not be factually demonstrated.

To clear up the above-mentioned confusion, we cover in the present work a broad range of material related to TCA and best execution. As understanding transaction costs is crucial to properly assessing the quality of implementation decisions and complying with the best execution obligation in the post-MiFID environment, our objectives are the following:

- provide a state of the art of TCA fundamentals;
- do a critical review of existing post-trade TCA techniques;

1 - MiFID is the second step in the harmonisation of the European capital markets industry and aims to adapt the first Investment Services Directive (ISD 1, issued in 1993) to the realities of the current market structure. Part of the European Financial Services Plan (FSAP), the “MiFID” (Directive 2004/39/EC, formerly known as Investment Services Directive II) was ratified by the European Parliament on April 21st, 2004. The implementing Directive and Regulation were approved by the Parliament over the summer of 2006 and provide detailed implementation guidelines applicable to all member states. MiFID came into force in November 2007. In a nutshell, MiFID sweeps away the very concept of central exchange and obligation of order concentration currently existing in several European countries, and recognises the need to include all participants in the execution cycle and all financial instruments under a consistent regulatory framework. At the same time, MiFID offsets the opening of the execution landscape to full competition with a set of obligations whose goal is to increase transparency and investor protection in order to maintain the efficient and fair price discovery mechanisms as well as the overall integrity of European markets in the face of inevitable fragmentation.

2 - MTF: multilateral trading facility

Introduction

- define a new and complete approach wherein checking whether best execution is achieved lies in the entire transaction cost management process.

To cover all the aspects related to these objectives, we have structured the present work in six sections.

Section I provides insights into how transaction costs arise during the investment cycle as well as an introduction to TCA's main goals. After a brief description of the investment phases, we first highlight the importance of being proactive in managing transaction costs to obtain competitive returns. We then present TCA as a very valuable decision-making tool since it allows monitoring both the transaction costs of different implementation strategies and the trading performance of different intermediaries. At this stage, we briefly introduce the three TCA missions—**transaction cost measurement, transaction cost estimation and trading performance assessment**. We conclude by showing how TCA can effectively contribute to enhancing total portfolio performance.

Section II provides a thorough investigation into transaction cost components and drivers. Transaction costs include all costs associated with trading, which are usually divided into explicit and implicit costs. Explicit costs include brokerage commissions, market fees, clearing and settlement costs, and taxes/stamp duties. These costs do not rely on the trading strategy and can be quite easily determined before the execution of the trade. By contrast, implicit costs represent the invisible part of transaction costs and consist of

spread, market impact and opportunity costs. As variable components, they offer the opportunity to improve the quality of execution. This section reviews all the costs of both categories in detail and gives insights into why they arise when investment decisions are made. In addition, we provide empirical evidence of the differences in total transaction costs as well as their composition across regions and trading venues, in Europe in particular.

Section III is devoted to the two fundamental approaches to measuring transaction costs. Benchmark comparison measures the cost of trading by the signed difference between the trade price and a specified benchmark price. This method provides cost indicators that depend on whether they are built on pre-trade, intraday or post-trade benchmarks. Implementation shortfall defines the cost of trading as the difference between the actual portfolio return and its paper return benchmark. This approach has become the standard for transaction cost measurement and we document how the primary framework may be expanded to fine-tune cost identification. Before reviewing the most popular indicators used in the industry, we emphasise several issues to consider when measuring costs. We conclude with new practical questions and complications that emerge with MiFID.

Section IV deals with pre-trade analysis, whose primary purpose is to forecast both the transaction costs and the risks associated with various strategies for a given trade or a specific trade list not yet executed. In this section we refer to the approach developed by Kissell and Glantz (2003) and show that it makes it possible to determine

Introduction

the cost profile of any trading strategy through simultaneous estimation of cost and risk, develop optimisation techniques to derive an efficient trading frontier and devise the most appropriate strategy to meet the investor's goal or comply with his preferences.

We would like to take the opportunity to express our sincere thanks to Robert Kissell and Morton Glantz, who have made, amongst other things, a very significant contribution to a better understanding and approach to modeling pre-trade transaction cost and risk estimates. Our aim here is only to provide a synthetic overview of the question; we have therefore summarised their material, but a good understanding of the approach cannot be obtained without reading in detail the full methodology available in Kissel and Glantz's 2003 *Optimal Trading Strategies*.

Section V is dedicated to trading performance evaluation. After having exposed how trading performance measurement differs from transaction cost measurement, we come back to the benchmark comparison approach, the most common practice in the industry, and we review its major shortcomings when it comes to measuring quality of execution. We then focus on the concept of best execution in general and see that although it is very fashionable, it is often misunderstood and does not mean the same thing to everyone. We conclude with the MiFID best execution obligation and show that the regulator has provided neither a clear definition nor a measurable objective to make up for the current absence of consensus in the industry.

Section VI gives us the opportunity to introduce a framework developed by the EDHEC Risk and Asset Management Research Centre team that makes it possible to address the crucial question of best execution evaluation for traders and investment managers. The model is explained in full and can be easily deployed and customised to financial institution's specific needs.

I. Transaction Cost Analysis as Part of the Investment Process



I. Transaction Cost Analysis as Part of the Investment Process

Four distinctive phases are usually identified in the investment decision cycle. They may be summarised as follows. The *asset allocation* phase refers to the distribution of funds across investment classes (equities, bonds, cash, commodities, derivatives, hedge funds, real estate, etc.) with the objective of diversifying risk and targeting a specified return. The *portfolio construction phase* is the phase wherein decisions about the exact instruments to buy or sell in each investment class are made. The *execution services phase* is the phase during which investment decisions are acted on. It involves decisions regarding the trading strategy: where, when and how to buy/sell. Finally, the *portfolio attribution phase* involves assessing portfolio performance and analysing reasons for missing the targeted return.

Most financial research is devoted to asset allocation, portfolio construction and performance attribution. We do not find the same abundance of literature on the implementation of investment decisions. For the total performance of a portfolio, however, the quality of the implementation is as important as the decision itself. The reason is that the implementation of investment decisions is not free and that the associated costs usually reduce portfolio returns with limited potential to generate upside potential. These costs can turn high-quality investments into moderately profitable investments or low-quality investments into unprofitable investments. These costs are usually referred to as transaction costs.

To provide investors with competitive portfolio returns, investment managers

must manage transaction costs proactively, because lower transaction costs mean higher portfolio returns. Ineffective transaction cost management may be very damaging in today's competitive environment, in which the difference between success and failure may be no more than a few basis points, or in which tiny spreads offer arbitrage strategies that must be protected from transaction costs. The aim of *transaction cost analysis* (TCA) is to provide a scorecard that helps investment managers assess understand how well their decisions have been acted on and how they can improve their overall performance. On the one hand, as different trading strategies correspond to different risk-cost trade-offs, investment managers need to know the real cost of implementing a given trading strategy. On the other hand, since bad execution can impact the total performance of even the best investment decision, investment managers must be able to assess the trading performance of their intermediaries (brokers, traders or even algorithms). Roughly, then, TCA is a tool for monitoring both the transaction costs of trading strategies and the trading performance of intermediaries.

Going into greater detail, we can attribute to TCA three distinct and specific tasks:

- Transaction cost measurement
- Transaction cost estimation
- Trading performance assessment

It is first very important to understand the difference between cost measurement and cost estimation, because, as we will see, each requires a specific methodology. In a nutshell, transaction cost measurement involves assessing the cost of completed

I. Transaction Cost Analysis as Part of the Investment Process

trades and hence occurs *ex post*. By contrast, transaction cost estimation attempts to forecast the cost of proposed trades. Basically, transaction cost measures and estimates differ in two main ways. First, transaction cost measures result in single point values expressed in either monetary units or basis points per share, while transaction cost estimates are expressed as a probabilistic distribution defining both an expected cost and a risk parameter. Next, estimation is essentially based on cost components and drivers, while identifying and measuring each cost component is not so obvious *ex post*.

When dealing with TCA, it is also of great importance to understand how trading performance evaluation differs from transaction cost measurement, even if both are done *ex post*. Performance evaluation attempts to assess how well intermediaries perform when executing trades. The ultimate aim is to determine the most effective intermediaries by market, instrument and trading strategy. With this information at hand, investment managers can considerably reduce the time required to select the best intermediary for the execution of a specific trade. While transaction cost measurement focuses on determining how large transaction costs are and where they arise, the analysis of trading performance involves determining whether the costs are justified or result from poor implementation decisions and could have been avoided.

By nature, all the tasks assigned to TCA must be performed within the execution services phase of the investment decision process. It is there that quantitative procedures and approaches can be developed to

measure the transaction costs of past trades, estimate the transaction costs of future trades and allow comparisons across trading strategies for a specific trade list, as well as measure the trading performance of the intermediaries. All the information collected in this phase must then be used in the other phases to avoid misallocation of funds, inefficient portfolio mixes and ineffective intermediaries. TCA thus contributes to enhanced total performance over the entire investment decision cycle.

To allow the reader to view this report in light of hard figures, Karceski, Livingston and O'Neal (2004) established that actively managed equity portfolios bear a total of 0.98% transaction costs per year. To a significant degree, those costs are implicit and are therefore not reported in total expense ratios.

I. Transaction Cost Analysis as Part of the Investment Process

II. Transaction Cost Components and Drivers



II. Transaction Cost Components and Drivers

Understanding exactly what transaction costs are and why they arise is the first and fundamental step when dealing with TCA. Transaction costs have nine components, usually categorised as implicit or explicit, as shown in figure 1. In this section, we describe the components of both categories in detail and give insights into why these costs arise when implementing investment decisions.

Figure 1: Typology of transaction costs

Brokerage Commissions	Explicit Costs
Market Fees	Explicit Costs
Clearing and Settlement Costs	Explicit Costs
Taxes/Stamp Duties	Explicit Costs
Bid-Ask Spread	Implicit Costs
Market Impact	Implicit Costs
Operational Opportunity Costs	Implicit Costs
Market Timing Opportunity Costs	Implicit Costs
Missed Trade Opportunity Costs	Implicit Costs

■ Explicit Costs ■ Implicit Costs

Those costs can be considered direct when related to individual orders/transactions, or indirect when accounted for globally as part of the provision of the transaction services. To provide a full picture, we also introduce the concept of direct and indirect explicit transaction costs.

1. Direct Explicit Transaction Costs

Brokerage commissions, market fees, clearing and settlement costs, and taxes/stamp duties are explicit costs. They are said to be explicit because they do not depend on the trade price and are usually documented separately from it. As these costs do not rely on the trading strategy, they can be known in advance, before the execution of the trade.

Brokerage commissions are paid to intermediaries for executing trades. They can be expressed on a per share basis or based on a total transaction value, most of the time in basis points and subject to volume discount. Although they differ from one intermediary to another, they are a fixed and visible transaction cost component.

Market fees are paid to trading venues for executing trades on their platforms. They are usually bundled into brokerage commissions for investors. These fees vary. On average, higher volume markets have the lowest costs. In recent years, competitive pressure has led to a significant reduction in these explicit costs. Like brokerage commissions, market fees are considered a fixed and visible transaction cost component.

Clearing and settlement costs are related to the process whereby the ownership of securities is transferred finally and irrevocably from one investor to another. When the trading venue owns the clearing and settlement system, these costs, which are a fixed and visible transaction cost component, are usually included in market fees. Like the latter, clearing and settlement costs differ from one trading venue to another. This is illustrated in table 1, which exhibits statistics about the cost per execution (in €) in the main European stock exchanges.

Table 1: Cost per execution in Europe (€)

Cost per Execution	Mean	Maximum	Minimum
Exchange	2.85	4.94	1.57
Clearing	1.24	2.57	0.38
Settlement	1.22	2.18	0.52
Total cost	5.31	7.00	4.01

Source: Various public sources (2007)

II. Transaction Cost Components and Drivers

Taxes/stamp duties are visible and variable transaction cost components. They are visible because tax rates or specific stamp duties are known in advance but variable because they often vary by type of return or trade.

In theory, explicit costs could be determined before the execution of the trade. In practice, their measurement is not so obvious because brokerage commissions are often paid for bundled services, not only for order execution. Research, analytics and trading technology are often bundled services provided by intermediaries. A trend towards unbundling is being observed in Europe and the US. New unbundling of commission regulations, which separates the payment for deal execution and the payment for broker research, will surely facilitate the measurement of explicit costs.

2. Indirect Explicit Transaction Costs

In addition to direct costs, explicit costs also include a number of indirect costs having to do with the processing supporting the execution and the counterparties involved. Even though the determination of these costs in today's environment is very difficult (and perhaps impossible on a trade-by-trade basis), their importance should not be underestimated when venues or particular types of transaction are opted for.

These indirect explicit transaction costs encompass:

- capital costs related to operational risks
- capital costs related to counterparty (and credit) risks

Capital requirements related to these costs are currently assessed on a global basis; allocating these costs to specific trading activities remains very ambitious even though these costs are often estimated to support strategic organisational decisions.

2.a Indirect costs related to operational risks

Basel II defines operational risk as the risk of loss resulting from inadequate or failed internal processes, people and systems, or from external events. From the perspective of managing transactions in financial instruments, operational risks relate to the risk of loss resulting from inadequate or failed internal processes, people and systems in the handling of the transaction cycle.

The definition provided in the Basel II capital requirement framework encompasses situations such as:

- Internal fraud—misappropriation of assets, tax evasion, intentional mismarking of positions
- External fraud—theft of information, hacking damage, third-party theft and forgery
- Business disruption and systems failures—utility disruptions, software failures, hardware failures
- Execution, delivery, and process management—data-entry errors, accounting errors, failed mandatory reporting, negligent loss of client assets.

The complex nature of the entire trade processing cycle, from execution to settlement, therefore has a direct impact on the operational risks borne by the financial institution.

II. Transaction Cost Components and Drivers

Basel II and various supervisory bodies have prescribed a number of soundness standards for operational risk management for banks and similar financial institutions. To complement these standards, Basel II has provided guidance on three broad methods of capital calculation for operational risk:

- Basic indicator approach
- Standardised approach
- Advanced measurement approach (AMA)

Banks using the **basic indicator approach** must hold capital for operational risk equal to the average over the previous three years of a fixed percentage (denoted alpha) of positive annual gross income.

In the **standardised approach**, bank activities are divided into eight business lines: corporate finance, trading and sales, retail banking, commercial banking, payment and settlement, agency services, asset management, and retail brokerage. Within each business line, gross income is a broad indicator that serves as a proxy for the scale of business operations and thus the likely scale of operational risk exposure within each of these business lines. The capital charge for each business line is calculated by multiplying gross income by a factor (denoted beta) assigned to that business line (18% for sales and trading, 12% for retail banking, 18% for payments and settlement, 12% for asset management and retail brokerage).

In the **advanced measurement approach**, the regulatory capital requirement will equal the risk measure generated by the bank's internal operational risk measurement system using quantitative

and qualitative criteria. Use of the AMA is subject to supervisory approval and makes it possible to reduce the minimum capital requirements significantly.

Under this last regime, widely opted for by major institutions, the actual operational risks measured by means of quantitative and qualitative analysis have a direct impact on capital requirements.

The following factors are likely to make a direct impact on the internal measure of operational risks:

- Nature of the execution process
- Extent of trade processing automation/manual interaction
- Nature of interfaces with third parties (paper/electronic/fax)
- Extent of operations outsourcing

The nature of the execution process (direct trading, algorithmic trading, phone/electronic interface) provides an important input to the extent of operational risks. As a result, the choice of execution venue and the maturity of the interface between the financial institution and the liquidity pool are key factors behind operational risks and hence capital requirements.

The extent to which the processing of trades is automated and the nature of interfaces with third parties are determined by the maturity of the markets on which transactions occur. For example, there are significant differences between emerging/developing and historic markets and exchanges. Similarly, transactions executed on exchanges and those that take place over-the-counter (OTC) are not equally automated. In the latter situation,

II. Transaction Cost Components and Drivers

financial institutions may not benefit from the advanced and proven market infrastructure that makes it possible to streamline and automate transactions.

As a result, in the post-MiFID environment, financial institutions must address the consequences of choosing alternative trading venues not only in relation to the execution process but, more importantly, in light of the impact on the entire trade processing cycle as well. The choice of alternative platforms implying OTC transactions (e.g., dark liquidity pools) might present significant advantages in terms of execution efficiency and quality but is likely to impose operational constraints that need to be addressed to avoid possibly significant capital implications.

Likewise, the extent to which operations are outsourced is likely to modify the firm's overall operational risk profile. Basel II does not allow firms to eliminate capital requirements by transferring operations to third-party providers, but those third parties may boast more modern and scalable operations tailored to processing transactions in the most efficient manner, implicitly reducing operational risks. When a firm outsources its transaction processing cycle to a third party, it is usually to a partner that specialises in providing this service; the firm can thus benefit from economies of scale and gain access to state-of-the-art infrastructure.

The quest for operations efficiency and the associated reduction in operational-risk costs account for much of the recent success of outsourcing offerings developed by pure back-office providers

as well as by back-office intermediaries such as custodian banks.

So it would be a mistake to believe that the choice of execution venue or changes in execution methods have no implications on the back-end of the processing cycle. Custodian banks have understood the importance of post-trade processing in the search for best execution perfectly well and, through better integration of alternative trading cycles in their operational processes with the aim of both reducing operational constraints and risks and enhancing overall efficiency, are well positioned to support the new requirements created by the fragmentation of liquidity venues.

2.b Indirect costs related to counterparty (and credit) risks

Counterparty risk exists when a financial institution transacts with another firm that may default on its obligation to settle a transaction or deliver securities related to a transaction. This situation can arise when the counterparty is in financial distress (we are then confronted with credit risk) or when the counterparty faces a pure operational or cash management/stock management glitch.

As a consequence, counterparty risk can occur when a financial institution lends cash or securities, but it could also be influenced by occasional problems in the settlement cycle. Failure to collect payment or deliver securities can have a damaging impact on a financial institution, as it may cause the firm to default on other transactions and create a cascade of events that can, if not dealt with in due time, lead to default.

II. Transaction Cost Components and Drivers

We will not deal here with the question of credit risk, *i.e.*, the assessment of the ability of the firm's counterparties to meet their obligation to repay/deliver, as this element is part of normal banking activity. The risk incurred by firms as a result of the settlement process of transactions should nonetheless be taken extremely seriously in light of the changes resulting from the implementation of MiFID.

The direct consequence of the development of MTFs and systematic internalisation is the likely significant increase in OTC transactions that involve a direct settlement to the counterparty of the trader rather than the clearing of transactions through a clearing house and a central counterparty. In most established markets, central counterparties and clearing houses have made it possible to reduce the complexity of the settlement process significantly (thanks to the netting of transactions) as well as the counterparty risk (thanks to the use of a single central counterparty managing the individual counterparty risks through appropriate processes such as delivery versus payment).

Once again, the choice of venue and the selection of a mode of execution will have a direct impact on the final level of risks and costs incurred by the financial institution.

2.d Increased importance of back-office providers

Allocating the exact indirect explicit costs of operational and counterparty risks to single transactions is obviously too ambitious, but these costs do represent significant expenditures or capital requirements that

have huge consequences on the economic profitability of trading activity.

So, to benefit from the economies of scale and state-of-the-art infrastructure offered by firms specialising in post-trade processing, financial institutions have begun outsourcing part of their back-office responsibilities to third parties.

Three broad benefits are to be expected:

- A decrease in the investments required to cover a growing number of liquidity pools with regards to post-trade processing. As specialist firms build infrastructure centrally for a large number of clients, they indirectly let their clients benefit from economies of scale that allow them to cover a larger number of pools than individual financial institutions would otherwise be able to.
- A significant increase in processing efficiency, leading to better control of explicit direct costs (settlement costs, costs of handling the clearing process, costs of reporting, streamlined processes for managing mistakes)
- A significant decrease in operational risk and application of the most advanced management processes, making it possible to reduce counterparty risk (or at the very least improve the management of breaks and intraday defaults) and the capital set aside as part of the Basel II framework.

But third-party back-office providers and custodians have also realised that their status as recipients of transaction data puts them in an ideal position to develop and propose a set of value-add services directly related to transaction cost analysis. Much as with the development of risk and performance allocation services

II. Transaction Cost Components and Drivers

delivered by back-office providers on the basis of portfolio and transaction data collected to manage positions on behalf of clients, these intermediaries can now offer transaction cost analysis services based on transaction data transmitted for settlement purposes.

The firms that attempt to provide services for the monitoring of best execution will face two main challenges:

- Access to data and consolidation of external market tick and transaction data
- Enrichment of existing front-to-back data flows to include front-office trading information (timestamps)

The second challenge is under the control of the third-party provider (though we do not underestimate the operational and technical complexity of handling trade data in addition to data usually transmitted for settlement and reporting purpose only), but creating an adequate repository of market and transaction data that will make it possible to perform such advanced forms of transaction cost analysis as peer group analysis will require heavy investment and proper coordination throughout the industry, as MiFID failed to deliver an industry model that would simplify the collection and re-distribution of public information on markets and transactions.

All the same, it is clear that back-office intermediaries are the only firms currently involved in the transaction cycle in a position to offer an independent assessment of transaction costs on behalf of the final investor. Like independent asset valuation, likely to be in demand following the recent banking crisis, independent

analysis of transaction costs is—because it is a prerequisite to the smooth running of a fragmented execution market—likely to become one of the essential value-added services provided by custodian banks to financial institutions managing third-party assets.

3. Implicit Transaction Costs

Transaction costs are more than just brokerage commissions, market fees and taxes. Implicit costs are the transaction costs that, invisible, cannot be known in advance because they are included in the trade price. They depend mainly on the trade characteristics relative to the prevailing market conditions. In other words, they are strongly related to the trading strategy and, as variable costs, provide opportunities to improve the quality of execution. These implicit costs can be broken down into their components: *spread, market impact and opportunity costs*.

(1) Spread

This component is compensation for the costs incurred by the liquidity provider. When taking liquidity (by buying at the best ask or selling at the best bid), we pay the spread. In the microstructure literature, three kinds of cost are usually associated with the bid-ask spread. *The order processing* cost is compensation for supplying an immediacy service to the market (Demsetz 1968). The ability to trade immediately rather than to have to wait for the opposite trade provides certainty for market participants. Liquidity is thus provided at that cost.

II. Transaction Cost Components and Drivers

The *inventory control* cost is compensation for the risk of bearing unwanted inventories (Ho and Stoll 1981). Accommodating other market participants' trades makes liquidity providers deviate from their optimal inventory based on their own risk-return preference. To restore their optimal position, they adjust their bid-and-ask prices to attract and/or avoid some trades. The *adverse selection* cost is compensation for the risk of trading with informed traders (Copeland and Galai 1983). Informed traders have a certain amount of private information that allows them to know or better estimate the true value of a security. As liquidity providers lose when they trade with informed traders, they widen their bid-ask spread for all market participants to cover their potential losses.

The spread represents the implicit cost of a round-trip for a small trade and, as such, may be measured from market data by the simple difference between the best ask and best bid quotes. The ease of obtaining measures for spread makes some people consider it a visible implicit cost. In any case, this cost component is variable since spreads vary over time according to trading conditions. In limit order book systems, spreads mechanically widen after a large trade consuming more than the


quantity available at the best opposite quote. The spread cost is therefore particularly sensitive to the timing of execution.

Figure 2 illustrates the variation of the volume-weighted bid-ask spread across several major European stock exchanges. For a given trading venue, spreads fluctuate across stocks. In general, spreads are negatively associated with market capitalisation and liquidity and positively associated with volatility and information asymmetry.

(2) Market impact

Market impact is the price to pay for consuming the liquidity available on the market beyond the best quote: to complete their "large" orders, buyers must pay premium prices and sellers must offer discounts. In other words, market impact is the price shift that is due to the trade size. Its main determinants are the trade size and the market liquidity available at the time of the trade. Accordingly, market impact can be viewed as a positive function of the trade size and a negative function of the liquidity available. For a given level of liquidity, market impact increases with the trade size. For a given trade size, market impact increases with the lack of liquidity.

Figure 2: Weighted average bid-ask spreads in %



Source: Various public sources, EDHEC-Risk Advisory

II. Transaction Cost Components and Drivers

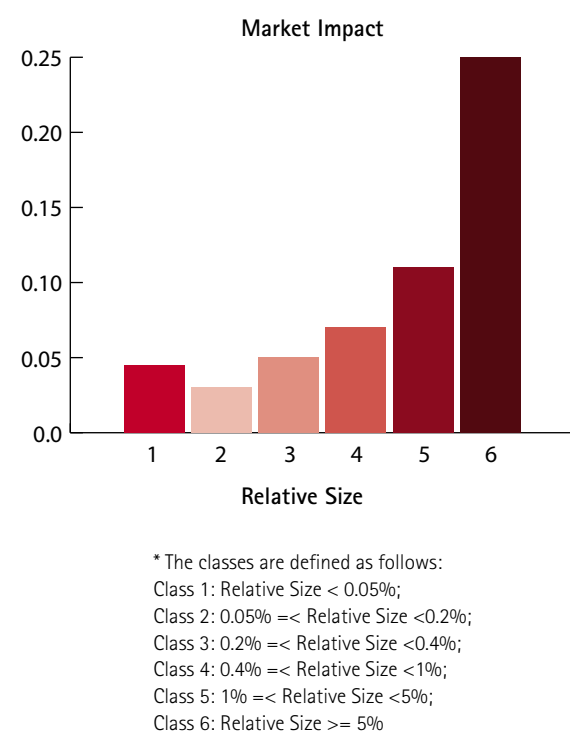
Trades affect market prices for two main reasons. First, if the trade is large relative to the available liquidity, the trade causes a market supply-demand imbalance and mechanically shifts the market price towards less favourable prices because it needs to attract additional liquidity (consumption of several quotes in the order book). This mechanical market impact lasts until liquidity is replenished. Second, trades can affect prices when they are perceived as motivated by new information. When the trade brings new information to the marketplace, a market correction occurs and the related market impact is permanent.

Market impact is one of the most costly implicit components because it generates mainly adverse price movements and therefore reduction in portfolio returns. It is an invisible and variable transaction cost: invisible because it cannot be easily determined before the execution of the trade and variable because it depends greatly on the trade size, the available liquidity and the specified trading strategy.

Figures 3 and 4 illustrate the relationship between market impact and trade size or liquidity. They are built on statistics based on the Sinopia Asset Management databases spanning sixteen developed countries and refer to trades executed during the period from 1999 to 2000. Figure 3 shows the relationship between market impact and the trade relative size. The latter is defined as the size of trade (expressed in number of shares) divided by the daily average volume for the stock. Market impact increases greatly from class four—trade size approaching or in excess of 1% of average daily volume—onwards.

Figure 4 focuses on the relationship between market impact and liquidity. Here, market liquidity is measured by turnover, defined as the ratio of the average daily volume to the total number of shares outstanding. As we can see, the lack of liquidity becomes expensive when turnover falls below 0.026% in the sample.

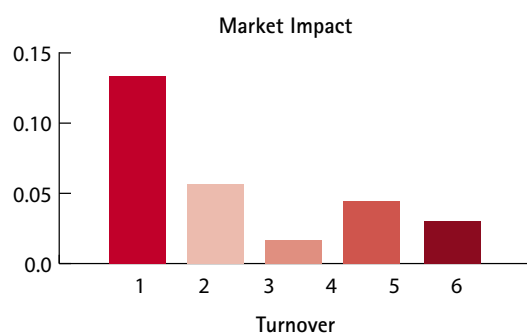
Figure 3: Average costs by trade relative size



Source: Boussema et al (2002)

II. Transaction Cost Components and Drivers

Figure 4: Average costs by turnover



* Share turnover = average volume traded during the previous month/Number of shares outstanding.
The classes of liquidity are defined as follows:

Class 1: Turnover < 0.019%;
Class 2: 0.019% ≤ Turnover < 0.026%;
Class 3: 0.026% ≤ Turnover < 0.034%;
Class 4: 0.034% ≤ Turnover < 0.046%;
Class 5: Turnover ≥ 0.046%

Source: Boussema et al. (2002)

(3) Opportunity costs

The decision to trade and the actual trade do not usually take place at the same time. Market prices can move for or against the proposed trade. The costs related to price fluctuation during the time required to act on an investment decision are opportunity costs. They arise when prices move between the time the trading decision is made and the time the order is executed.

In general, two elements are responsible for changes in opportunity costs during the time required to complete the trade: price appreciation and timing risk. Price appreciation represents the natural price trend of the security. Timing risk has to do with the uncertainty associated with the price movement. In other words, the first element has to do with expected price movements and the second has to do with unexpected price movements.

At each stage between the investment decision and the execution of the trade, there are potential delays and thus price movements that could positively or negatively affect portfolio returns. Opportunity costs are consequently invisible and variable implicit costs that depend greatly on the speed of execution.

As the source of the time between trade decision and execution can vary, opportunity costs can be broken down into three components:

- operational costs
- market timing costs
- missed trade costs

Operational opportunity costs arise when the time required to trade is operational and unintended (transmission delays between buy and sell sides, for example). When the delay results from market timing under the control of the broker (for example, the broker splits the order into small lots over a period to minimise market impact), we refer to *market timing opportunity costs*. Operational and market timing costs are sometimes both considered delay costs. We prefer the distinction as it enables identification of the party responsible for the costs incurred: the investor or his intermediary. Finally, *missed trade opportunity costs* are incurred when investors fail to fill their orders. Some trades may not be fully completed either because price movements have led to the cancellation of the initial trading decision or because there is no more security available (lack of liquidity). If a predetermined trading strategy is not completed, the resulting opportunity cost can be high. Failing to trade can be costly for the investor, who will have missed the

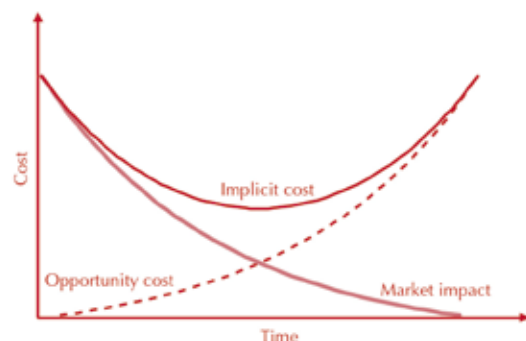
II. Transaction Cost Components and Drivers

opportunity to make an investment in the security requested.

The contribution of opportunity costs to total implicit costs is not independent of market impact. Attempting to reduce one can lead to the increase of the other. For example, splitting large orders over time to reduce market impact can lead to larger opportunity costs and vice versa. This issue is summarised in figure 5. Minimising market impact and reducing opportunity costs form a set of conflicting objectives usually referred to as "the trader's dilemma". Balancing these conflicting costs is a real challenge.

3 - For a discussion about the long-term impact of MiFID on the European trading landscape, see D'Hondt and Giraud (2007).

Figure 5: Relationship between market impact and opportunity costs



Source: Giraud (2004)

4. Transaction cost magnitude

Several empirical studies have examined the relative importance of explicit and implicit transaction costs. Most have focused on the costs of equity trading for institutional investors. We can draw several conclusions from these empirical findings.

First, transaction costs have fallen dramatically in recent years, essentially due to technological innovations and increased competition among trading venues and among intermediaries.


This fall is reported by Domowitz et al. (2001) in an analysis of equity trading costs in a sample of forty-two countries. Boussema et al. (2001) also document this fall in transaction costs. For a sample of trades included in the Sinopia Asset Management database, the authors find that total transaction costs amounted to about 0.38% from 1996 to 1999 while they fell to 0.18% from 1999 to 2000.

We observe a similar finding in Munck (2005), a more recent study devoted to transaction costs at the larger stock exchanges in Europe and in the north. Munck shows that total transaction costs have dropped in recent years. He attributes this drop to the significant drop in explicit costs. Figure 6 shows that the explicit costs of trading on the OMX exchanges, Euronext and Deutsche Börse have fallen over the past eight years and are fairly clustered. The higher explicit costs on the London Stock Exchange are the result of a special stamp duty of fifty basis points on all buy trades. Figure 7 shows the pattern of implicit costs over the same period. Although there are fluctuations, it is clear that costs at these exchanges are clustered. In late 2004, implicit costs ranged from ten to fifteen basis points.

Today, with MiFID just in place, it is expected that execution fees will keep falling in Europe. The impact of the Directive on total transaction costs, however, is not so obvious, as implicit costs depend on the efficiency of market structures, on which the long-term implications of MiFID are as yet unknown.³


II. Transaction Cost Components and Drivers

Figure 6: Explicit transaction costs



Source: Munck (2005)

Figure 7: Implicit transaction costs




Source: Munck (2005)

Second, transaction costs vary from region to region and even within regions. Although they are falling, they remain economically significant, especially in emerging markets. According to Domowitz *et al.* (2001), over the third quarter of 2000, total transaction costs ranged from twenty-two basis points in the Netherlands to 184 basis points in Venezuela, with a cross-country mean of about sixty basis points. Boussema *et al.* (2001) also compare transaction costs in developed and emerging markets. In their sample of trades, explicit costs (including only brokerage commissions) average 0.15% in developed markets and 0.61% in emerging markets. This phenomenon is similar for implicit costs (market impact and delay costs): they are about 0.23% in developed markets and approximately 0.58% in emerging markets. Checking for a possible correlation between costs and the use of trading systems, Munck (2005) identifies market system turnover as a statistically significant explanatory

variable for both explicit and implicit costs. He finds that both explicit and implicit costs fall as system activity increases. Figure 8 illustrates this relationship for total transaction costs.

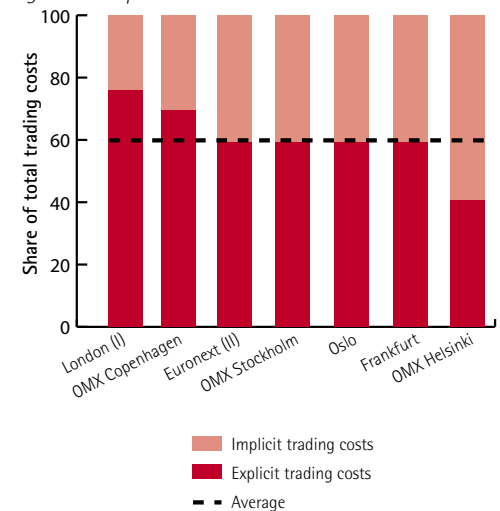
Figure 8: Total transaction costs vs. market system turnover



Source: Munck (2005)

Third, the composition of transaction costs can vary across trading venues. Figure 9, excerpted from Munck (2005), breaks down total transaction costs in 2004 for the European exchanges in his sample. Except on the OMX Helsinki market, explicit costs account for the lion's share of transaction costs, with an average of approximately sixty basis points. This cost composition is unlike that in US trading venues, where implicit costs tend to be higher.

Figure 9: Composition of transaction costs



Source: Munck (2005)

II. Transaction Cost Components and Drivers

5. Transaction Costs and Execution Methods

Transaction costs also depend on execution methods. There are two kinds of techniques for executing an order: agency trading or principal trading. The two methods differ in terms of the risk sharing between the investor and the executing broker. In fact, the choice between agency trades and principal trades involves a trade-off between a low certain explicit cost and an uncertain implicit cost.

Agency trading is the most common execution method. When sending orders to the market through an agency, the investor bears all the risks associated with the trade. The agency assumes no market risk, so all the implicit costs fall on the investor. In practice, the investor sends an order to the broker, specifying the name of the security and the number of shares to buy or sell. The agency then executes the order and the speed of execution depends mainly on the order size and the available liquidity. In some cases, the investor can give the agency price (closing price, daily VWAP, etc.) or volume (percentage of the daily volume, for example) constraints. Average brokerage fees are relatively low.

In principal trading, the broker provides the investor with guaranteed execution of the trade list at the market prices at a specific point in time (usually the close). Here, the entire market risk is transferred to the broker because he acquires the investor's position and all associated risk. As a consequence, brokerage fees are higher than for agency trades and depend on the risk associated with the execution of the trade list. These specific higher commissions are known as the principal bid premium and serve as an insurance policy for the broker.

The two methods and their implications are summarised in table 2.

Table 2: Execution methods and transaction costs

Execution method	Characteristics	Advantage	Disadvantage
Agency trade	Release of an order specifying the security and the quantity Execution on the market with or without a target (closing price, VWAP, etc.)	Low fees	Market risk (implicit costs)
Principal trade	Determination of a basket: size, securities, etc. Commitment of the broker to trade the basket at a predetermined price	No market risk	Higher fees

II. Transaction Cost Components and Drivers

III. Measuring Transaction Costs with Post-Trade Analysis



III. Measuring Transaction Costs with Post-Trade Analysis

As Harris (2003) notes, it is hard to manage effectively what you cannot accurately measure. Proper measurement of transaction costs is thus crucial for quality decision-making, since it permits assessments of portfolio profitability and final performance. More broadly, accurate transaction cost measures are also useful for building and testing transaction cost estimation models.

In this section, we discuss the two currently used fundamental approaches to measuring transaction costs. These approaches are the benchmark comparison and the implementation shortfall. We then identify common pitfalls to look out for when measuring costs and offer a critical review of the most popular indicators. The section concludes with a discussion of complications brought about by MiFID.

1. Benchmark Comparison

This method involves comparing the monetary value difference between the executed position and the position evaluated at a benchmark price. For a given transaction, the monetary value per share is given by the signed difference between the average price obtained for the trade and the benchmark price. A positive value indicates an execution that is less favourable than the benchmark price, while a negative value indicates execution more favourable than the benchmark. Negative costs, in other words, are savings. The choice of benchmark is primordial here. For one, the benchmark price should be easy to obtain or compute. For another, it should be ideal, in the sense that it should enable assessment of the price that would

have been observed if the trade had not taken place. In this case, the difference between this price and the trade price could be pinned entirely on the trade.

Different transaction cost indicators are used for different benchmarks. The indicators can be based on pre-trade, intraday or post-trade benchmarks. Furthermore, in each of these three categories, an additional distinction can be made between absolute indicators and time-related indicators. Absolute indicators are easier to compute because they do not require consideration of time stamping. In other words, these indicators do not take into account the exact time at which the trade is executed. By contrast, time-related indicators are dependent on the exact time at which the trade is completed. They rely on a benchmark price that is computed around or at the execution time.

(1) Indicators built upon pre-trade benchmarks

These indicators use benchmark prices prevailing on the market before the execution of the trade. The most frequent are based on the following benchmark prices, depending on whether they consider the accurate time of execution or not.

Absolute indicators:

- T-1 Close: previous night's closing price
- T Open: opening price of the day

Time-related indicators:

- Ask: last ask price before execution
- Bid: last bid price before execution
- Last: last trade price before execution
- Midpoint bid-ask: last bid-ask average before execution

III. Measuring Transaction Costs with Post-Trade Analysis

(2) Indicators built upon intraday benchmarks

These indicators are based on benchmark prices that represent a kind of average market price for the day of the trade execution. The most frequent are based on the following benchmark prices, depending on whether they focus on the exact time of execution or not:

Absolute indicators:

- Daily VWAP: the volume-weighted average price of the day
- Daily LHOC: the average of the lowest, highest, opening and closing prices of the day
- Midpoint H/L: the median of the highest and lowest prices of the day

Time-related indicators:

- Available VWAP: the VWAP computed from the market opening to the trade time
- Interval VWAP: the VWAP calculated over a fixed time interval around the trade execution time⁴

(3) Indicators built upon post-trade benchmarks

This final category contains indicators referring to benchmark prices that prevail on the market after the execution of the trade. These indicators are less heterogeneous than in the two previous categories and the most frequently used in the industry are listed below.

Absolute indicators:

- T Close: closing price of the day

Time-related indicators:

- Next mid bid-ask: next bid-ask average

2. Implementation shortfall

This method, widely recognised as the most effective means of measuring transaction costs, was initially proposed by Perold (1988) and its principle is easily understood. It involves assessing the impact of trading on portfolio returns by computing the difference between the net returns on a paper portfolio and those on a real portfolio.

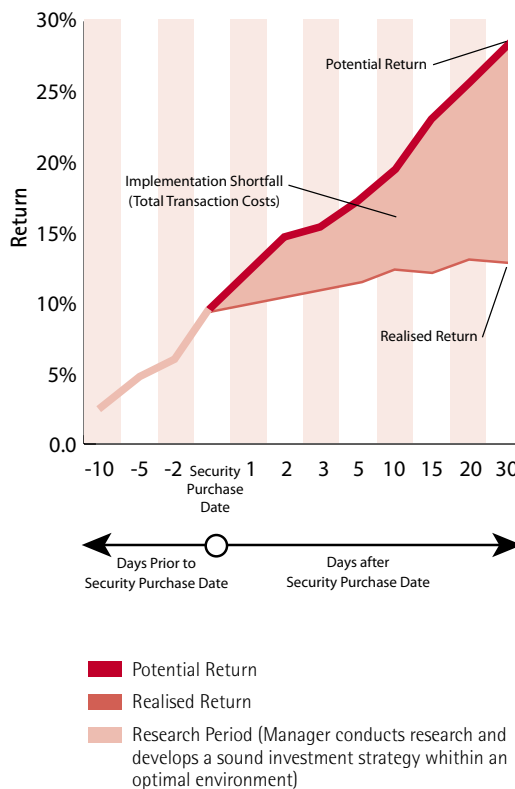
(1) Original framework

Determining implementation shortfall is a two-step process. First, at the time of the investment decision, the investor needs to build and price a paper portfolio. This portfolio will be an imaginary holding consisting of all the security positions the investor decides to take. It is assumed that these positions are acquired at the price on the market at the time it was decided to acquire them. Consequently, the resulting paper portfolio, unlike the corresponding actual portfolio, does not incur any transaction costs. The difference between the paper portfolio return and the actual portfolio return is the "implementation shortfall". Figure 10 gives a summary overview of the method.

4 - The VWAP computed over a certain long period of time (n days) around the trade execution time also exists. This multi-day VWAP can be relevant for trades on very illiquid securities.

III. Measuring Transaction Costs with Post-Trade Analysis

Figure 10: Principle of the implementation shortfall



5 - If an investor simultaneously buys and sells shares in a security and recorded buys at the best ask and sells at the best bid, the resulting paper portfolio will incur a loss due to the spread.

Source: Plexus Group

Much as with the benchmark comparison, the key here is the choice of price for valuing the paper portfolio. Perold's definition requires that all executions occur at the spread midpoint prevailing on the market at the time of the investment decision. This specific valuation is necessary to avoid charging the paper portfolio with one-half the spread cost on average.⁵

Calculated with the decision spread midpoint, transaction costs are then divided into what the author terms *execution costs and opportunity costs*. Execution costs cover all the costs that can be put down to the actual trade, such as explicit costs and market impact. However, as market impact is here measured by the difference between the market price at the decision

time and the actual trade price, it actually delivers a nearly complete measure of implicit costs. Indeed, in this formulation, Perold's market impact includes costs resulting from spread, market impact, operational delay and market timing delay. Therefore, the opportunity cost component refers only to the missed trade opportunity cost, *i.e.*, to the cost related to unfilled or partially filled orders.

The mathematical notations for measuring the implementation shortfall for a given trade in a particular security can be presented as follows.

X = total quantity to execute (total trade size)

x_i = number of shares executed at price i
 $X, x_i > 0$ for a buy; $X, x_i < 0$ for a sell

P_0 = quotation midpoint at the time of the investment decision

P_T = quotation midpoint at the end of trading

P_i = execution price of i^{th} trade

EC = all the explicit costs associated with the trade

When the trade is completed, the implementation shortfall is:

$$IS = [XP_T - XP_0] - \left[\left(XP_T - \sum_{i=1}^N x_i P_i \right) - EC \right]$$

$$IS = \sum_{i=1}^N x_i P_i - XP_0 + EC$$

When the trade is not completed, the implementation shortfall becomes:

$$IS = \left[\sum_{i=1}^N x_i P_i - \sum_{i=1}^N x_i P_0 \right] + \left(X - \sum_{i=1}^N x_i \right) (P_T - P_0) + EC$$

III. Measuring Transaction Costs with Post-Trade Analysis

We can illustrate this primary version of the implementation shortfall with a very simple example. Suppose an investment decision to buy 500 shares on a given security traded in € whose implementation details are the following:

$$X = 500$$

$$x_1=200 \text{ and } x_2=200$$

$$P_0 = 20$$

$$P_T = 22$$

$$P_1 = 21 \text{ and } P_2 = 21.5$$

Based on the above data, the implementation shortfall is calculated as follows:

$$IS = [(200 \times 21 + 200 \times 21.5) - 400 \times 20] + (500 - 400)(22 - 20) + EC$$

For this investment decision, the missed trade opportunity cost measure is €200 or €0.4 per share, while the other implicit costs (spread, market impact, operational and market timing opportunity costs) amount together to €500 or €1 per share. These cost measures can also be expressed in basis points (bp). In this case, we obtain 200bp for the missed trade opportunity cost and 500bp for the other implicit costs.

As we have seen, the implementation shortfall is easily understood and implemented. Although it seems to be an effective means of measuring transaction costs, Perold's framework delivers interesting information about how large transaction costs are but not really about where they are incurred. To better identify where they are incurred, the original version must be expanded.

(2) Expanded framework

There are several ways to expand Perold's approach to fine-tune the measurement of transaction costs. Some are presented in Kissell and Glantz (2003). The expanded version we propose here is adjusted for the classification of implicit cost components provided in section II.

To measure transaction costs and properly identify where they occur, the investor can consider four points on the time line: the investment decision time, the order release time, the time at which the broker begins to implement the trade and the time at which the broker stops trading. With the method below, the investor can then identify four reference prices to calculate the implementation shortfall:

X = total quantity to execute (total trade size)

x_i = number of shares executed at price i

$X, x_i > 0$ for a buy; $X, x_i < 0$ for a sell

P_0 = quotation midpoint at the time of the investment decision

P_R = quotation midpoint at the time the order is released to the broker

P_S = quotation midpoint at the time the broker starts working the order

P_T = quotation midpoint at the end of trading

P_i = execution price of i^{th} trade

EC = all the explicit costs associated with the trade

III. Measuring Transaction Costs with Post-Trade Analysis

$$IS = \underbrace{\left[\sum_{i=1}^N x_i P_R - \sum_{i=1}^N x_i P_0 \right]}_{(1)} + \underbrace{\left[\sum_{i=1}^N x_i P_S - \sum_{i=1}^N x_i P_R \right]}_{(2)} + \underbrace{\left[\sum_{i=1}^N x_i P_i - \sum_{i=1}^N x_i P_S \right]}_{(3)} + \underbrace{\left(X - \sum_{i=1}^N x_i \right) (P_T - P_0)}_{(4)} + EC$$

In this equation, part (1) measures the operational opportunity cost, part (2) measures the market timing opportunity cost, part (3) measures the cost due to both spread and market impact and part (4) measures the missed trade opportunity cost.

To illustrate the degrees of interpretation made possible with this expanded approach, we may apply this formula to the previous example. For this purpose, the following information must be added: $P_R = 20.5$ and $P_S = 21$.

$$IS = [400 \times 20.5 - 400 \times 20] + [400 \times 21 - 400 \times 20.5] + [(200 \times 21 + 200 \times 21.5) - 400 \times 21] + (500 - 400)(22 - 20) + EC$$

As observed before, the missed trade opportunity cost measure is €200 or €0.4 per share, while the other implicit costs amount together to €500 or €1 per share. In addition, we are now able to split the global cost of 500 into its main components: €200 for operational opportunity cost, €200 for market timing opportunity cost and €100 for both spread and market impact. This breakdown is relevant as it allows an understanding of exactly where costs are incurred. In our simple example, it seems obvious that minimising investor-to-intermediary transmission delays could reduce transaction costs substantially.

3. Common Pitfalls and Shortcomings in Transaction Cost Measurement

The two approaches that have been described in depth both require an appropriate benchmark price. Identifying this price is not always as easy as it seems. Transaction cost indicators are numerous and their quality often depends on the choice of benchmark. The "ideal" transaction cost indicator should provide an accurate, complete and relevant measure of implicit costs. This definition emphasises several issues that should be considered when building or choosing transaction cost indicators.

(1) Which benchmark for which purpose?

In a nutshell, the benchmark on which the indicator of transaction costs is built determines what is really measured. Accordingly, different benchmarks serve different purposes. As explained above, benchmark prices can be categorised as pre-trade, intraday and post-trade. Both the completeness and the relevance of the measures they deliver vary greatly.

The *pre-trade benchmarks* provide the best measures of the implicit costs related to the implementation of an investment decision and should be favoured. They include the costs due to spread, market impact, operational and market timing delays; only missed trade opportunity costs are ignored. Although indicators

III. Measuring Transaction Costs with Post-Trade Analysis

built on pre-trade benchmarks are quite complete measures, their accuracy varies with the gap between the investment decision time and the time at which the benchmark price is determined. The most accurate cost indicator is built on the price prevailing at the decision time. When this price is unknown, the prior night's closing price or the opening price is typically used as a proxy for the decision price, but the indicators it delivers are then less accurate.

The *intraday benchmarks* do not provide good measures of the implicit costs. In fact, indicators built on intraday benchmarks do not really give a measure of implicit costs but rather an indication of the execution price relative to the average market price of the day or of a given time period. Indicators built on intraday benchmarks are thus neither proper nor useful cost measures. For example, knowing that a trade has underperformed or outperformed the daily VWAP by a given number of basis points does not provide the investor with any insight into how large the implicit costs are and how they can be better managed. Instead, this information allows the investor to gauge the quality of the price obtained, if we consider that the daily VWAP is a good indicator of the fair market price. As we will see in section V, intraday benchmark prices are best used to assess the trading performance of intermediaries.

Like the previous category, the *post-trade benchmarks* do not really deliver a proper and complete indicator of implicit costs, although they are sometimes used for this purpose. Instead, they can provide an indicator of market impact only

through a measure of price reversion. However, whatever the specific post-trade benchmark price, this measure refers only to the temporary impact, because the potential permanent impact is incorporated into future prices and is impossible to isolate. Although their usefulness as measures of cost is limited when taken alone, post-trade benchmarks are more relevant when combined with pre-trade benchmarks to give a measure of missed trade opportunity costs. Such a combination is present in the implementation shortfall approach.

(2) Critical review of the most popular indicators

Although we assert that an effective measure of implicit transaction costs is the difference between the decision price (its proxy when not available) and the execution price, we know that this assertion has not always won unanimous backing. In this subsection, we offer an in-depth presentation of the transaction cost indicators most often used by the industry. We look at both the advantages and disadvantages of these indicators.

(a) Spread midpoint benchmark

An indicator based on the spread midpoint relies on the signed difference between the trade price and the average of the bid-and-ask at a given time. Spread midpoint indicators are very popular, essentially because they are very easy to implement and interpret. The cost of a buy at the "ask" (or a sell at the "bid") is one-half of the spread. Varying the time at which the spread midpoint is determined delivers various indicators.

III. Measuring Transaction Costs with Post-Trade Analysis

The quotation midpoint prevailing at the investment decision time is the reference price for computing the *implementation shortfall*. The implementation shortfall is the most effective of the indicators of transaction costs founded on the spread midpoint. The only disadvantage of this indicator is that it requires a large amount of data, *i.e.*, intraday quotation data as well as decision time and order data. Collecting all this data and dealing with it can be difficult or quite costly for investment managers. In practice, time-stamped information on individual trades may be incomplete or not gathered in one place but held in different systems and perhaps not in the same firm. Therefore, the implementation shortfall indicator indirectly requires that data acquisition and treatment costs not be a consideration.

The *effective spread* relies on the quotation midpoint prevailing at the time of the trade execution. The effective spread is twice the liquidity premium, which is the signed difference between the trade price and the time-of-trade midpoint. This indicator ignores all the opportunity costs but gives a measure of market impact when compared with the quoted spread. When the trade is completed at the quoted price (a buy at the ask/a sell at the bid), the effective spread equals the quoted spread and market impact is zero. When the trade is completed outside the spread (the large order is filled at prices outside the best quote), the effective spread is larger than the quoted spread and market impact is positive. When price improvements mean that the trade can be completed inside the spread, the effective spread is smaller than the quoted spread and market impact is negative. This last case may be a little

counterintuitive because market impact is mainly associated with adverse price movement. The main shortcomings of the effective spread are twofold. First, it can provide a measure of market impact but without a distinction between temporary and permanent parts. Second, the effective spread may be a poor indicator of market impact for large orders completed through multiple trades. When orders are split into smaller lots, the indicator should measure the total impact of executing the entire order and not simply add up the cost of each lot. As the first lots impact the market, they make the remaining lots more expensive. An indicator based on the effective spread will underestimate the total market impact costs if a different midpoint is used for each trade.

The *realised spread* is based on a post-trade quotation midpoint. It equals twice the signed difference between the trade price and the midpoint observed at a specified time after the trade. Like the effective spread, this indicator attempts to measure market impact. However, this measure misses the permanent impact cost, as previously explained. Furthermore, the greater the time gap between trade execution and midpoint determination, the noisier the cost indicator will be.

(b) VWAP

The volume-weighted average price (VWAP) indicator is the signed difference between the trade price and the market volume-weighted average price computed over a given interval. The most frequently used indicator relies on the daily VWAP, *i.e.*, the volume-weighted average of all prices in the trading day. Nowadays, almost all data vendors as well as some

III. Measuring Transaction Costs with Post-Trade Analysis

trading venues compute and publish daily VWAP in real time. However, the window can be shorter (interval/available VWAP) or longer than one day (multi-day VWAP). The VWAP-related indicator is widely used in the industry, mainly because it is easily interpreted: it indicates whether the trader received a higher or lower price than did the "average trader" of the measurement interval. Nevertheless, the VWAP benchmark is inappropriate for correct measurement of transaction costs.

(c) Closing price benchmark

The indicator based on the closing price is the signed difference between the trade price and the closing price of the day. This benchmark is very attractive in the industry because it is easy to get and requires little data processing. All other popular indicators require intraday market trade data (VWAP, LHOC) or quotation data (spread midpoint), while the indicator relying on the closing price needs only summary daily market data. Furthermore, many investment firms prefer to use closing prices as benchmarks because they already value their portfolios at these prices.

Although attractive, the closing price may not provide an effective measure of implicit costs. It can only deliver indicators of temporary market impact, because permanent impact is incorporated into future prices. Moreover, such indicators are noisy when the time gap between the execution of the trade and the market close is long.

(d) Average of LHOC

The indicator relying on the LHOC is the

signed difference between the trade price and the average of the lowest, highest, opening and closing prices of the day. This indicator is widely recognised in the industry, even though it suffers from the same shortcomings as VWAP indicators. Furthermore, as a simple average of prices irrespective of traded quantities, the LHOC is less representative of the fair market price since it misses the dimension of market depth included in the VWAP.

4. Measuring Transaction Costs under MiFID

Under pressure from MiFID, TCA is set to grow and transaction cost measurement in particular is likely to become both a widespread and a more than essential practice. However, two main issues are emerging with the new regulation and are expected to make transaction cost measurement more complex in the near future. The first has to do with market fragmentation and the second with data availability and processing.

(1) Market fragmentation

The European trading landscape is likely to become more fragmented⁶ post-MiFID since the new regulation clearly encourages the proliferation of liquidity pools and promotes competition by putting an end to the well-established order concentration rule. This rule had hitherto made it possible for most of the traditional exchanges to maintain monopolistic positions and be viewed as official providers of reference prices in a great majority of European countries. By allowing multiple trading venues to compete, MiFID is, by design, allowing liquidity pools to fragment and put

6 - Fragmentation occurs when not all orders relative to a given security interact with each other in a single execution system. Markets are fragmented when people can trade the same securities in various trading venues. Markets are consolidated when all market participants trade on the same execution venue. In practice, we obviously refer to fragmentation as soon as a given security is cross-listed or may be traded on multiple execution venues.

III. Measuring transaction costs with post-trade analysis

the very existence of an acknowledged reference price at risk. At first sight, both benchmark comparison and implementation shortfall are made conceptually difficult in the absence of an acknowledged unique price for the securities traded. When a security is traded at the same time on various execution venues, which of the coexisting prices is the best reference for the security? Several elements can be put forward to address this point.

Under MiFID, investment firms have potential access to a wide range of execution venues in the form of regulated markets,⁷ MTFs⁸ and systematic internalisers.⁹ However, as part of the best execution obligation, MiFID requires that investment firms determine their own unique set of execution venues and define an execution policy documenting the selected trading venues as well as the factors that led to this selection. This selection process reduces the number of liquidity pools to consider when executing orders and creates a “frame of reference” within which each firm must meet its best execution obligation. Hence, when it comes to measuring transaction costs, only the list of selected execution venues will be relevant to determine reference prices.

When only one trading venue is selected for a given security, we are back to a pre-MiFID situation with a single recognised reference price; therefore, transaction costs can be measured as simply as previously described. However, this case is likely to be an exception and, most of the time it is likely that more than one trading venue has been selected for a

given security. It is then necessary to pay attention to the venue(s) the order is routed to.

When an order is routed to and entirely executed on a single trading venue (for whatever reason¹⁰), it seems quite consistent to use as benchmarks prices available on that venue. But when an order is split into smaller lots that are routed to several execution venues to sweep the market by taking liquidity from all sources at once—as may be frequent for large orders—complications arise and measuring transaction costs becomes a true brain-teaser. When there are multiple price sources for an order, the ability to keep track of and consolidate all these prices is eventually the major determinant of the ability to measure costs properly. Theoretically, each lot should be considered a single order for which transaction costs could be measured based on benchmarks prevailing on its execution venue. Next, the total cost of trading for the initial order should be obtained by adding all costs related to each lot. This approach requires great data storage and processing capabilities.

(2) Data availability and treatment

Accurate transaction cost measurement requires a large amount of time-stamped data about orders and trades. Besides, as we have just mentioned, post-MiFID market fragmentation is likely to tighten these requirements substantially. In practice, investment firms can have trouble collecting information about their own trades as well as collecting market data. The reason is twofold:

- Incompleteness of internal databases
- Time-stamped information on individual

7 - This regime refers to the traditional exchanges in the MiFID language.

8 - Multilateral trading facilities are similar to regulated exchanges in that they allow clients to enter into negotiations without taking part in a transaction as counterparty. They include all forms of multilateral negotiations such as order books, block trades, periodic auctions and any other mechanism resulting in negotiations between two counterparties.

9 - This new regime allows investment firms on an organised, frequent and systematic basis to deal on their own account by executing client orders outside a regulated market or MTF.

10 - This can be a constraint set by the client, an automated choice when the venue quotes the best of the prices available, etc.

III. Measuring transaction costs with post-trade analysis

trades may be incomplete. For example, time-stamped data can become available only from the time of the release to the broker. In this case, no data including decision time is available, with the consequence that some transaction cost indicators cannot be calculated with accuracy. Another possibility is that time-stamped information is held not in one place but in different systems and perhaps not in the same firm.

- Access to intraday market data

Both benchmark comparison and implementation shortfall require intraday market data, in addition to information about the trades being analysed. When several executive venues have been selected, more data and greater processing capability are required. Consequently, the firm's ability to manage all trading data is key. Indeed, one of the most important aspects is the ability to consolidate prices from multiple and unconnected systems. Clearly, the ability to source prices and integrate them through a single system is a great advantage.

- Data quality, completeness and integrity

For the implementation of MiFID, the regulator has left to the industry the choice of organising the infrastructure needed to consolidate market and reported data. As no central solution has emerged, the risk of having data reported in numerous places and incorrectly consolidated is great, resulting in an absence of quality, completeness or integrity and making the implementation of post-trade analysis impossible.

MiFID indirectly addresses the above issue with harmonised post-trade public transparency requirements that facilitate aggregation for all transactions on equities,

whether executed on a regulated market or not. Such consolidated pan-European market data is being offered by some data vendors and new offers should appear in coming months. All these solutions in the routing and transformation of real-time market and trade data that aggregate feeds from multiple contributors and publish this information to a firm's market data system in real-time are likely to help build relevant databases and tools for proper transaction cost measurement. Nevertheless, aggregated market data is still often restricted to equities or liquid securities. Investment firms can collect consolidated intraday market data for other securities, but the acquisition cost can be higher.

It seems clear then that MiFID has huge operational and IT implications for investment firms. There is no doubt that these implications are going to strengthen the presence of technology providers in the European execution landscape.

III. Measuring transaction costs with post-trade analysis

IV. Estimating Transaction Costs with Pre-Trade Analysis



IV. Estimating Transaction Costs with Pre-Trade Analysis

This section is devoted to the second mission of TCA—transaction cost estimation. Its primary purpose is to forecast both transaction costs and risk associated with strategies for a given trade or trade list not yet executed. This assumes that the implicit costs of future orders can be estimated from the implicit costs of past trades, using trade-specific variables (order size, price limit, etc.), market-related variables (spread, depth, recent volume, recent price change, etc) and security-specific variables (average volume, capitalisation, volatility, etc.).

Since transaction cost estimation is done *ex ante* and is more than the mere provision of cost estimates, it is often referred to as pre-trade TCA or simply pre-trade analysis. It is essential for investment managers who want to offer competitive portfolio returns to do pre-trade analysis before acting on investment decisions. It helps them assess the difficulty of trades, compare trading strategies, and develop appropriate strategies for particular market circumstances. In line with these three objectives, we can present pre-trade TCA as a three-step process:

- Collection and analysis of pre-trade data
- Cost and risk estimation
- Optimisation of trading strategies

Each step needs an appropriate approach that requires a large amount of data and particular tools as well as the development of relatively complex models. We provide a detailed description of each step below.

1. Collection and Analysis of Pre-Trade Data

The first step of pre-trade TCA involves collecting and analysing pre-trade data

to understand the characteristics of the proposed trade and assess the degree of difficulty. The concept of “difficulty” is primordial, as it predicts the amount of work that will be required to execute the order and the overall need to adjust the trading strategy to changing market conditions.¹¹

All the data gathered from pre-trade analytics is usually presented in a *pre-trade report* wherein information may be categorised as security-specific, market-related and trade-specific. We review each category below and provide insight into how it can help gauge the difficulty of the order to execute.

(1) Security-specific information

All the data delivering relevant information about the security to trade fall into this category. It includes:

- the country in which the security was issued
- the economic sector
- the market capitalisation
- the venues where the security is traded
- some fundamentals such as price-to-earnings and price-to-book ratios
- the index the security belongs to, if any
- the recent historical price movement
- the price momentum
- the historical average daily trading volume
- the historical price volatility
- the average trade size
- the intraday volume pattern
- etc.

This list provides a quick reading of the risk associated with the security to be traded. It is not exhaustive and any variable investment managers consider relevant could be added.

11 - Kissell and Glantz (2003) define trading difficulty as a function of the amount of work required to execute an order and the overall need for traders to adjust to changing market conditions and make modifications to the specified trading strategy.

IV. Estimating Transaction Costs with Pre-Trade Analysis

(2) Market-related information

This category contains all the data delivering relevant information regarding the current market conditions for the security to trade. The purpose is here to assess the market risk of the proposed order. If several trading venues are available, comparisons must be made on the conditions prevailing at each venue. The variables most relevant to a description of market conditions are commonly:

- the last trade prices
- the best quotes and the quotation midpoint
- the quoted spread
- the quoted depth
- the recent traded volume
- etc.

This list should provide the best possible representation of the current market prices and liquidity available for the security to be traded. The list is exhaustive if the security is traded in transparent liquidity pools; otherwise, it is not. Indeed, in transparent markets, the prevailing liquidity conditions can be easily assessed through the displayed limit order book or market makers' quotations. In dark liquidity pools, by contrast, no information about trading conditions is disclosed.

(3) Trade-specific information

This category covers order size, trade direction (buy/sell) and other characteristics of the order to be executed. When time or price constraints are attached to the order, they must be added since they can make implementation harder. For example, if the client wants the trade to be completed before a certain time, this deadline should appear in the list.

With a pre-trade report summarising security-specific, market-related and trade-specific data, investment managers are able to gain a thorough understanding of the trade to be completed. They can then assess its difficulty and know at a glance whether the trade will be subject to substantial market impact (for example, the order is large in comparison with the quoted depth and the average trade) and/or opportunity costs (for example, the liquidity to execute the order immediately is lacking). Difficult trades will require permanent adjustments to the trading strategy.

2. Cost and Risk Estimation

The second step of pre-trade TCA is cost decomposition and estimation. Only implicit transaction costs are considered here because they depend on the trading strategy and may thus be controlled during implementation. The same is not true for explicit costs, which are known in advance and do not really vary with the trading strategy, other than with the choice of venue and intermediary.

The proper framework for estimating transaction costs attempts to provide a probabilistic distribution of both an expected cost and a risk parameter for a trade to be completed. Two essential principles emerge here and it is essential to understand them clearly. First, cost estimates, unlike transaction cost measures, are not single values. Computing cost estimates involves providing a complete picture of potential costs with their respective risk, where both parameters vary with the implementation strategy. Second, estimating a distribution of cost-risk couples makes it possible to assess alternative strategies and does

IV. Estimating Transaction Costs with Pre-Trade Analysis

the groundwork for implementation that meets the investor's goals and respects his preferences.

Various sophisticated tools and models are on offer today to this end. Our objective is not to review all of them; it is impossible because there are too many and there is little public information about how they work, essentially for commercial reasons. We intend instead to document the construction of a complete framework for proper forecasting of transaction costs. Accordingly, we have chosen to refer hereafter to the analytical approach proposed by Kissell and Glantz (2003). These authors put in the public domain a considerable amount of technical and very detailed material dealing with the means of modeling and forecasting implicit transaction costs. Their work allows us to illustrate the approach without promoting a particular commercial tool. We will see also that this approach makes it possible to determine the cost profile of any trading strategy through simultaneous estimation of cost and risk, which is the basis for developing optimisation techniques to derive an efficient trading frontier and determine the optimal strategy most closely aligned with the investor's final objective.

(1) Cost structure

Kissell and Glantz (2003) build their analytical framework for cost estimation on the following price trajectory formulation: $P_t = P_{t-1} + U_t + K_t + E_t$, where P_t is the price of trade t , U_t is the natural price appreciation from time $t-1$ to t , K_t is the market impact of trade t and E_t is the price volatility with $E_t \sim N(0, \sigma^2)$. In this price trajectory, we can identify price appreciation (U_t) and timing risk (E_t) as drivers for opportunity costs as well as

market impact (K_t). For the latter, we simplify reality since there is no distinction between permanent impact and temporary impact that dissipates over time.

From the above cost structure, estimating transaction costs results in making the sum of the forecasted value of each driver: price appreciation, market impact and timing risk. To get such values, we need estimation models that are able to predict cost values when applied to a specific trading strategy. We review below each cost driver and use Kissell and Glantz's material to show what such models look like and how they deliver estimates.¹²

(2) Price appreciation

As explained above, price appreciation is the natural price movement of a security. It is sometimes referred to as price trend. It is a truism to say that transaction costs depend on price appreciation: momentum traders tend to incur large transaction costs because they buy when prices have risen and sell when they have fallen. By contrast, contrarian traders buy when prices have fallen and sell when they have risen, so their transaction costs tend to be lower. Taking price appreciation into account thus adds substantial value to the implementation process.

To predict the cost of price appreciation, we need first to develop price forecasts over the trading horizon and then determine cost estimates for the specified execution strategy.

(a) Price forecast models

Both fundamental analysis and technical analysis provide price forecasts in the form of expected short-term or long-term price

¹² - We present here a summarised version of the models developed by the authors. For more detail, see Kissell and Glantz (2003).

IV. Estimating Transaction Costs with Pre-Trade Analysis

movement. We can derive from it target prices for all the periods in the trading horizon. Three models can be used for this purpose. To present them and facilitate understanding, we refer to the following numerical example.

Suppose a security that is currently traded at €100 (P_0) and is expected to increase to €125 (P_T) after one year. If we consider fifteen-minute periods, we have 8,500 trading intervals (n) over the forecasted horizon.¹³ We can first determine the expected price change over each trading interval and then calculate the expected price at the end of two days (sixty-eight trading intervals (m)).

Linear model

This model assumes constant price changes regardless of current prices. Applying this model to our example, we get the constant price change (ΔP) and the expected price (P_T) as follows:

$$\Delta P = \frac{1}{n}(P_T - P_0) = \frac{1}{8500}(125 - 100) = 0.0029$$

$$P_t = P_0 + m\Delta P = 100 + 68 \times 0.0029 = 100.20$$

Return model

This model assumes that the price changes by a constant percentage of the current price. Applying this model to our example, we get the percentage price change (r) and the expected price (P_T) as follows:

$$r = \left(\frac{P_T}{P_0}\right)^{\frac{1}{n}} - 1 = \left(\frac{125}{100}\right)^{\frac{1}{8500}} - 1 = 0.00026\%$$

$$P_t = P_0(1+r)^m = 100 \times (1 + 0.000026)^{68} = 100.18$$

Growth model

This model also assumes that the price

changes by a constant percentage of the current price. Applying this model to our example, we get the exponential price change (g) and the expected price (P_T) as follows:

$$g = \frac{1}{n} \ln\left(\frac{P_T}{P_0}\right) = \frac{1}{8500} \ln\left(\frac{125}{100}\right) = 0.0026\%$$

$$P_t = P_0 e^{(m \times g)} = 100 \times e^{(68 \times 0.000026)} = 100.18$$

Return and growth models are interchangeable since they deliver similar results for any short-term trading horizon.

(b) Cost estimates

Consistent with the implementation shortfall approach and if we apply the linear price change model, the transaction cost component due to price appreciation is estimated as follows.¹⁴

P_0 = security price at the beginning of trading

ΔP = linear price change forecast

P_j = forecasted price in period j ;

X = order size; $X > 0$ for buys and $X < 0$ for sells

x_j = number of shares traded in period j ;

$$\sum_{j=1}^n x_j = X$$

$U(X)$ = cost of price appreciation for order X

$$U(X) = \sum_{j=1}^n x_j P_j - XP_0$$

$$U(X) = \sum_{j=1}^n x_j P_0 + \sum_{j=1}^n x_j j \Delta P - XP_0$$

$$U(X) = XP_0 + \sum_{j=1}^n x_j j \Delta P - XP_0$$

$$U(X) = \sum_{j=1}^n x_j j \Delta P$$

13 - We consider that there are thirty-four intervals of fifteen minutes in a usual trading day (9:00AM to 5:30PM) and 250 trading days in a year.

14 - The method is similar when we consider return and growth models.

IV. Estimating Transaction Costs with Pre-Trade Analysis

$U(X)$ estimates the cost expressed in monetary units for the total order size. To get the cost per share or in basis points, we simply divide $U(X)$ by the order size (X) or the traded value at the beginning of trading (XP_0). Because $U(X)$ is a cumulative function, predicting the cost of price appreciation for a specific trade list that contains several orders and securities results in making the sum of cost estimates for each order and listed security.

We can illustrate the method above with a small example. Consider an order of 170,000 shares to be executed over three trading periods. The current security price is €50 and this price is expected to increase by €0.10 per period. If the execution strategy involves trading 60,000, 60,000 and 50,000 shares in periods 1, 2 and 3; we estimate the cost of price appreciation as follows:

$$U(X) = 60000 \times 0.10 + 60000 \times 2 \times 0.10 + 50000 \times 3 \times 0.10 = 33000$$

This cost of €33,000 is equivalent to €0.19 or 39bp per share.

(3) Market impact

Market impact is the price shift caused by a particular trade. It can be broken down into short-lived impact that dissipates over time and long-term disturbances in the price trajectory.

As with price appreciation, we need first a model that relates the market impact cost of past trades to a number of trade-specific, market-related, and security-specific variables. We then use the calibrated model to obtain cost estimates for the specified execution strategy. However, because market impact is a more complex cost component

than price appreciation and depends on elements such as imbalances, volatility, trading style and liquidity, the approach is more difficult and technical. We provide hereafter a summarised version of Kissell and Glantz's market impact model and illustrate how it can be used to derive cost estimates for given trading strategies.

(a) Market impact model

As explained just above, market impact depends on several factors:

- Order size and imbalances: market impact is positively related to order size and the liquidity supply-demand imbalances it causes;
- Volatility: higher volatility means greater price elasticity and higher transaction costs;
- Trading style: aggressive trading strategies lead to higher market impact than do passive strategies;
- Market conditions: market impact is negatively related to the available liquidity over the trading horizon.

Based on the previous factors, Kissell and Glantz (2003) model the total market impact cost as: $MI = \alpha I + (1 - \alpha)I$, where α is the percentage of temporary impact; $(1 - \alpha)$ is the percentage of permanent impact and I is the total market impact cost. The latter is expressed in monetary units and depends on both imbalance and volatility.

Assuming that only liquidity demanders incur temporary impact costs¹⁵ and that permanent impact is a function of the net imbalance,¹⁶ the market impact cost expressed in monetary units per share is rewritten as:

$$MI = \frac{\alpha I}{V_{side}} + \frac{(1 - \alpha)I}{Q} \quad \text{where } V_{side} \text{ is the}$$

15 - Liquidity demanders consume the available liquidity by trading at market prices and contribute to the supply-demand imbalance. By contrast, liquidity providers supply liquidity to obtain better prices for trades that they would like to complete. Hence, only liquidity demanders pay temporary impact costs.

16 - Permanent impact is due to the net effect of all market participants and all market participants incur permanent impact costs. This assumption is valid if (i) all trades convey some information to the market and the security current fundamental value is continuously assessed; (ii) the information content is proportional to the trade size and similar for buy orders and sell orders.

IV. Estimating Transaction Costs with Pre-Trade Analysis

traded volume from liquidity demanders and Q is the net imbalance.

In this top-down approach, the market impact cost for an order of X shares is then given by:

$$MI(X) = X \left[\frac{\alpha I}{V_{side}} + \frac{(1-\alpha)I}{Q} \right]$$

As we may observe, if the order size is equal to the imbalance and accounts for all liquidity demand ($X=Q=V_{side}$), the total market impact cost incurred by the investor is I . Hence, the market impact cost for a specific trading strategy $MI(x_k)$ can be derived according to the following principles. First, the percentage of temporary impact in any trading period k is equal to the percentage of the imbalance in that period. The total temporary impact cost can then be allocated to each trading period based on the percentage of imbalance in each period ($\frac{q_k}{Q}$ where q_k is the net imbalance in period k). Next, by applying an average permanent impact cost across all trades, allocation of total permanent impact across periods is not necessary. The cost of market impact for the trading strategy x_k is therefore calculated as follows:

$$MI(x_k) = \sum_{k=1}^n x_k \left[\frac{q_k}{Q} \frac{\alpha I}{V_{side}} + \frac{(1-\alpha)I}{Q} \right]$$

(b) Parameter estimation

To calibrate their model, Kissell and Glantz (2003) start with the total cost of market impact over an entire trading day when $X=Q$:

$$MI(X=Q) = Q \left[\frac{\alpha I}{V_{side}} + \frac{(1-\alpha)I}{Q} \right] = \frac{\alpha IQ}{V_{side}} + (1-\alpha)I$$

As previously explained, I is a function of both imbalance and volatility. By defining $\eta = \frac{V_{side}}{Q}$, the authors obtain the following market impact formulations:

$$MI(Q) = I(Q, \sigma) (\alpha \eta^{-1} + (1-\alpha)) = I(Q, \sigma) d(\eta)$$

$$MI_{bp} = I_{bp}(Z, \sigma) d(\eta)$$

The first equation specifies the total market impact cost in monetary units, the second in basis points, with imbalance stated as a percentage of average daily volume (Z). Both equations now present market impact as a product of two functions where $d(\eta)$ is the dissipation function and $I(Z, \sigma)$ is the instantaneous market impact function. The variables in each function have to be determined and the parameters have to be jointly estimated with advanced regression techniques. We report below what the authors propose given their own analyses.¹⁷

Variables

The instantaneous market impact function is built on two variables, Z and σ . The first represents the imbalance expressed as a percentage of average daily trading volume and is computed as follows:

$$Z = \frac{|Q|}{ADV} \times 100$$

where

$$Q = \sum_{i=1}^n \text{sign}(v_i)$$

with $\text{sign}(v_i)$ equal to the signed trade size and

$$ADV = \frac{1}{T} \sum_{t=1}^T v_t$$

with T usually equal to thirty trading days. The second variable σ refers to the price volatility factor and corresponds to the close-to-close measure of the volatility of logarithmic price change

17 - The results reported by Kissell and Glantz (2003) are based on numerous analyses on various sample sizes, volatility, participation rates and mixed data samples.

IV. Estimating Transaction Costs with Pre-Trade Analysis

for the previous T trading days (most of the time $T=30$). Mathematically, we have

$$\sigma = \sqrt{\frac{1}{T-1} \sum_{t=1}^T (g_t - \bar{g})^2}, \text{ where}$$

$$g_t = \ln \frac{P_t}{P_{t-1}}.$$

The dissipation function relies on the variable η that represents the participation number. It is used as a proxy of trading style and is calculated as the number of equal size orders executed in the market over the same trading period as the investor's order.

$$\text{Specifically, we have } \eta = \frac{V_{\text{side}}}{|Q|}.$$

Taking the absolute value of Q is required to ensure that the participation rate is positive. The lower the participation number is, the shorter the trading time horizon is, and the more aggressive the strategy is.

Finally, estimation of the various parameters of the market impact model requires an additional variable that represents the trading cost associated with the imbalance Q . This cost expressed in monetary units is computed according to the following equation:

$$\varphi_a = \sum x_j p_j - X P_0 = X P_a - X P_0 = X (P_a - P_0)$$

where P_0 is the market price at the beginning of trading and P_a is the volume-weighted average price for shares on the same side as the imbalance, that is:

$$P_a = \frac{\sum v_i p_i}{\sum v_i} \text{ for } i \in Q.$$

The cost of imbalance expressed in value per share and in basis points is given by $\varphi_a = \text{sign}(Q)(P_a - P_0)$ and

$$\varphi_a = \text{sign}(Q) \left(\frac{P_a}{P_0} - 1 \right) \times 10^4.$$

Parameter estimates

For the parameter α in the dissipation function, the authors found in almost all their regressions that it is about 95%; this estimate is not dependent on time and size. We can then rewrite the dissipation function as: $d(\eta) = 0.95\eta^{-1} + 0.05$. Since α refers to the temporary impact, Kissell and Glantz's findings suggest that temporary impact cost is by far the largest part of total market impact cost.

Kissell and Glantz (2003) investigate three possible structures for the instantaneous market impact function: the linear function, the non-linear function and the power function. Each requires i parameters termed a_i . We present below each possible function with both the general formulation and the formulation including the estimated parameters.¹⁸

- Linear function:

$$I_{bp}(Z, \sigma) = a_1 Z + a_2 \sigma + a_3$$

$$I_{bp}(Z, \sigma) = 8Z + 0.3\sigma + 90$$

- Non-linear function:

$$I_{bp}(Z, \sigma) = a_1 Z^{a_2} + a_3 \sigma + a_4$$

$$I_{bp}(Z, \sigma) = 35Z^{0.65} + 0.3\sigma + 15$$

- Power function:

$$I_{bp}(Z, \sigma) = a_1 Z^{a_2} \sigma^{a_3}$$

$$I_{bp}(Z, \sigma) = 25Z^{0.38} \sigma^{0.28}$$

(c) Forecasting market impact cost

The forecast of market impact cost for an order $K(X)$ and a particular trading strategy $K(x_k)$ is based on the following equations:

18 - According to Kissell and Glantz (2003), the three functions yield fairly consistent results across size and volatility. However, the non-linear and power functions exhibit smaller regression errors, suggesting that the true relationship between cost and size is non-linear.

IV. Estimating Transaction Costs with Pre-Trade Analysis

$$K(X) = I(0.95\eta^{-1} + 0.05)$$

$$K(x_k) = \sum_{k=1}^n x_k \frac{0.95Iq_k}{Qv_{k,side}} + \frac{0.05I}{Q}$$

To apply these equations, estimated values are needed for imbalance Q , liquidity demander volume by period $v_{k,side}$ and instantaneous cost I . Kissell and Glantz (2003) discuss two cases that require specific forecasting techniques.

First case: the imbalance is equal to the order size

This situation corresponds to the common assumption that $Q=X$ and each $x_k=q_k$ since intentions of other market participants are hard to assess before trading begins. Under this assumption, instantaneous cost is determined as follows:

$$I = I_{bp}(Z, \sigma)10^{-4}P_0Q \text{ where}$$

$$Z = \frac{|X|}{ADV}100$$

As for liquidity demander volume, the expected volume on the day $V(t)$ is first computed as: where $DOW(t)$ is a day-of-week adjustment factor accounting for the weekly volume effect.¹⁹ The percentage of volume in each trading period v_k is then obtained by applying the security volume profile $u_{k,t}$ which represents the percentage of the day's total volume traded in each period:

$$v_k = V(t)u_{k,t}$$

Since the total imbalance is assumed equal to the order size, v_k consists of an equal amount of buy and sell volume. Consequently, the liquidity demander volume $v_{k,side}$ will be equal to the imbalance in each period (x_k) plus half of the market volume. We have thus $v_{k,side} = x_k + 0.5v_k$.

We finally obtain that the market impact cost (in monetary units) for an order of X shares implemented over n -trading periods following strategy x_k is:

$$K(x_k) = \sum_{k=1}^n x_k \left[\frac{0.95Ix_k}{X(x_k + 0.5v_k)} + \frac{0.05I}{X} \right]$$

Since market impact cost is a cumulative function, the market impact cost estimated for a list of m -securities traded over n -periods is simply the sum of costs for all orders. That is:

$$K(x_k) = \sum_{i=1}^m \sum_{k=1}^n |x_{i,k}| \left[\frac{0.95I_i |x_{i,k}|}{|X_i|(|x_{i,k}| + 0.5v_{i,k})} + \frac{0.05I_i}{|X_i|} \right]$$

Second case: there is an incremental imbalance due to other market participants

This situation corresponds to the assumption that $Q=X+Y$, where Y is the expected net imbalance of other market participants.²⁰ Appropriate adjustments in the previous developments are necessary to incorporate this piece of information. Accordingly, the instantaneous market impact cost becomes

$$I = I_{bp}(Z, \sigma)10^{-4}P_0|X + Y| \text{ , where}$$

$$Z = \frac{|X + Y|}{ADV}100 \text{ with } X \text{ or } Y > 0 \text{ for buys}$$

and X or $Y < 0$ for sells. Next, the liquidity demander volume in each trading period is now:

$$v_{k,side} = |x_k + y_k| + 0.5v_k$$

19 - Patterns of daily trading volumes are used to exhibit a day-of-week effect. A factor specifying the day's historical percentage of the average daily volume can improve the daily volume forecast.
20 - This case is possible when investors can formulate realistic expectations about buying and selling pressure from other market participants. It often happens at times of public announcements or index reconstitution.

IV. Estimating Transaction Costs with Pre-Trade Analysis

It is worth noting that the absolute value function is necessary to account for the cases when the investor's order and the net imbalance of other market participants are on opposite sides. Furthermore, whenever the net imbalance of other market participants is larger and on the opposite side, the investor incurs a savings rather than a cost. To address this potential savings, the calculation must provide the investor with the correct sign for market impact cost. This is made through the following cost function:

$$\text{sign}(k) = \text{sign}(X) \cdot \text{sign}(X + Y)$$

From all the previous adjustments, the market impact cost (in monetary units) for an order of X shares implemented over n -trading periods following strategy x_k is:

$$K(x_k) = \text{sign}(k) \cdot \sum_{k=1}^n x_k \left[\frac{0.95I |x_k + y_k|}{|X + Y| (|x_k + y_k| + 0.5v_k)} + \frac{0.05I}{|X + Y|} \right]$$

For a list of m -securities traded over n -periods, the market impact cost forecasting equation becomes:

$$K(x_k) = \sum_{i=1}^m \sum_{k=1}^n \text{sign}(k_i) \cdot |x_{i,k}| \left[\frac{0.95I_i |x_{i,k} + y_{i,k}|}{|X_i + Y_i| (|x_{i,k} + y_{i,k}| + 0.5v_{i,k})} + \frac{0.05I_i}{|X_i + Y_i|} \right]$$

(d) Illustration

We now present a set of examples to illustrate how Kissell and Glantz's (2003) market impact model can, once it is calibrated, be used for easy cost estimates. For convenience, we change nothing and

assume that the parameter values obtained by the authors are suitable for our examples. However, pre-trade analysis assumes that the market impact of future orders can be predicted from the market impact of past orders. This assumption is valid when liquidity conditions do not change too much. So, we recommend to those interested in using this market impact model to determine their own calibration with their own recent data sample, since the data used by the authors may be not relevant for estimation on other markets, securities and periods.

Example 1: Forecasting market impact for an order when the imbalance is equal to order size

An investor has a buy order for 100,000 shares (X) on a particular security whose current market price is €30 (P_0), daily volatility is 200bp/day (σ) and average daily volume is 1,000,000 shares (ADV). The investor wants to execute the order over an entire day whose day-of-week adjustment factor is about 0.95 (DOW). Given the data at hand, how can he estimate the market impact cost for his order (in monetary units)?

The following five steps lead to an answer:

1. express the imbalance as a percentage of ADV:

$$Z = \frac{|100000|}{1000000} \times 100 = 10$$

2. determine the market impact instantaneous cost using the power function (for example):

$$I = (25 \times 10^{0.38} \times 200^{0.28}) \times 10^{-4} \times 30 \times 100000 = 79314$$

IV. Estimating Transaction Costs with Pre-Trade Analysis

3. compute the expected volume on the day:

$$V = 0.95 \times 1000000 = 950000$$

4. calculate the participation number:

$$\eta = \frac{100000 + 0.5 \times 950000}{100000} = 5.75$$

5. determine the expected market impact cost:

$$K(X) = 79314 \times (0.95 \times 5.75^{-1} + 0.05) = 17070$$

Example 2: Forecasting market impact for an order when there is an incremental imbalance

Suppose a situation similar to the previous example, except that the investor believes now that there will be an additional imbalance of 100,000 shares (Y) on the same side as his order. Given the data at hand, how can he estimate the market impact cost for his order?

Now the following six steps lead to an answer:

1. express the imbalance as a percentage of ADV:

$$Z = \frac{|100000 + 100000|}{1000000} \times 100 = 20$$

2. determine the market impact instantaneous cost using the power function (for example):

$$I = (25 \times 20^{0.38} \times 200^{0.28}) \times 10^{-4} \times 30 \times 200000 = 206429$$

3. compute the expected volume on the day:

$$V = 0.95 \times 1000000 = 950000$$

4. calculate the participation number:

$$\eta = \frac{|100000 + 100000| + 0.5 \times 950000}{|100000 + 100000|} = 3.38$$

5. determine the sign:

$$\text{sign}(100000) \times \text{sign}(100000 + 100000) = + \times + = +$$

6. compute the expected market impact cost:

$$K(X) = +206429 \times \frac{100000}{|100000 + 100000|} \times (0.95 \times 3.38^{-1} + 0.05) = 34214$$

Example 3: Forecasting market impact for a given strategy when the imbalance is equal to order size

An investor has a buy order for 200,000 shares (X) and wants to execute 40,000 shares (x_k) in each of the coming five periods. Forecasts for expected market volume in each period are 250,000 (v_1), 200,000 (v_2), 100,000 (v_3), 200,000 (v_4) and 250,000 (v_5). If the instantaneous market impact cost is estimated at €150,000 (I), how can he compute the cost related to his strategy?

$$\begin{aligned} K(x_k) &= \frac{I}{X} \sum_{k=1}^n x_k \left[\frac{0.95 x_k}{x_k + 0.5 v_k} + 0.05 \right] \\ &= \frac{150000}{200000} \sum_{k=1}^5 \left[\frac{0.95 x_k^2}{x_k + 0.5 v_k} + 0.05 x_k \right] \\ &= 50271 \end{aligned}$$

Example 4: Forecasting market impact for a given strategy when there is an incremental imbalance

An investor has a buy order for 200,000 shares (X) and wants to execute 40 000 shares (x_k) in each of the coming five periods. Forecasts for expected market volume in

IV. Estimating Transaction Costs with Pre-Trade Analysis

each period are 250,000 (v_1), 200,000 (v_2), 100,000 (v_3), 200,000 (v_4) and 250,000 (v_5). Furthermore, the investor believes that there will be an incremental market imbalance of 100,000 shares (Y). Incremental imbalances allocated to each period are 35,000 (y_1), 30,000 (y_2), 20,000 (y_3), 10,000 (y_4) and 5,000 (y_5) respectively. If the instantaneous market impact cost is estimated at €300,000 (l), how can he determine the cost related to his strategy?

$$K(x_k) = \text{sign}(k) \frac{l}{|X+Y|} \sum_{k=1}^n x_k \left[\frac{0.95|x_k+y_k|}{|x_k+y_k|+0.5v_k} + 0.05 \right]$$

$$K(x_k) = + \frac{300000}{|200000+100000|} \sum_{k=1}^5 \left[\frac{0.95x_k|x_k+y_k|}{|x_k+y_k|+0.5v_k} + 0.05x_k \right] = 83350$$

(4) Timing risk

Timing risk is related to the uncertainty associated with price movements and cost estimates. It can be broken down into price risk, which consists of price volatility, and liquidity risk, the result of fluctuations in market conditions. Price risk affects price appreciation estimates ($U(x_k)$) while liquidity risk affects market impact estimates ($K(x_k)$). Here again, we will refer to Kissell and Glantz's material to document how timing risk components can be modelled and forecasted.

Consistent with what we have seen earlier, we can write the following transaction cost estimate equation:

$$\phi(x_k) = U(x_k) + K(x_k) = \left[\sum_k x_k k \Delta p \right] + \left[\sum_k \frac{0.95x_k^2 l}{X(x_k + 0.5v_k)} + 0.05l \right]$$

where l is the instantaneous market impact

cost for the security, Δp is the expected price trend per period for the security, X is the order size, x_k is the number of shares to trade in period k and v_k the expected volume for the security in period k . Hence, the uncertainty surrounding the transaction cost estimate can be computed as the standard deviation of the previous equation, that is:

$$\mathfrak{R}(\phi) = \sqrt{E(\phi^2) - [E(\phi)]^2}$$

Assuming that volume and price movement are independent²¹ simplifies the calculation and allows the following equation:

$$\mathfrak{R}(\phi) = \sqrt{\sigma^2[U(x_k)] + \sigma^2[K(x_k)]}$$

where timing risk is obtained through a simple combination of price risk and liquidity risk.

(a) Price risk

The most common measure of price risk is the standard deviation of price returns. For trading purposes, price return volatility must be converted into monetary units per share as follows:

1. assuming that \bar{p} is not much different from each p_t , we can write:²²

$$r_t = \frac{p_t - p_{t-1}}{p_{t-1}} \cong \frac{p_t - p_{t-1}}{\bar{p}} = \frac{1}{\bar{p}} \Delta p_t$$

2. since \bar{p} is a constant:

$$\sigma^2(r) = \frac{1}{\bar{p}^2} \sigma^2(\Delta p_t)$$

3. for short time periods (≤ 5 days) and with $E(r) \leq 50\%$ and $\sigma(r) \leq 75\%$ on a yearly basis, the current security price p_0 may be substituted for \bar{p} with little loss of accuracy:

$$\sigma^2(r) \cong \frac{1}{p_0^2} \sigma^2(\Delta p_t)$$

21 - There is a negligible degree of correlation between price changes and market volume, but the absolute value of price changes and that of volume are correlated. This means that the presence of high volume does not indicate the direction of price movement.
22 - Natural log of price returns can also be used. Both calculations deliver very similar results for normal market conditions.

IV. Estimating Transaction Costs with Pre-Trade Analysis

4. hence, variance and volatility expressed in monetary units per share are easily obtained as follows: $\sigma^2(\Delta p_t) \cong p_0^2 \sigma^2(r)$ and $\sigma(\Delta p_t) \cong p_0 \sigma(r)$.

5. therefore, the total variance and volatility in monetary units for an order is: $\sigma^2(X) \cong X^2 p_0^2 \sigma^2(r)$ and $\sigma(X) \cong X p_0 \sigma(r)$.

Price risk of a specific trading strategy in a single security

Based on the variance of the security expressed in monetary units, computing price risk for a specific trading strategy results in computing the risk of a one-security position that fluctuates per period. If x_k is the number of shares traded in period k , r_k is the number of unexecuted shares at the beginning of period k and σ^2 is the per period variance for the security in monetary units/share, the total price risk of the trade schedule over n -periods is calculated as follows:

$$\sigma^2(x_k) = \sum_{k=1}^n r_k^2 \sigma^2$$

$$\sigma(x_k) = \sqrt{\sum_{k=1}^n r_k^2 \sigma^2}$$

We can illustrate this approach with a small example. Suppose a trader has an order for 12,000 shares (X) and wants to execute 4,000 shares (x_j) in each of the coming three periods. If the per period volatility of the security is estimated at €0.04/share, the price risk of the given strategy is computed as follows:

$$\sigma^2(x_k) = 12000^2 \times 0.04^2 + 8000^2 \times 0.04^2$$

$$+ 4000^2 \times 0.04^2 = 358400$$

$$\sigma(x_k) = \sqrt{358400} = 599$$

Price risk of a list of m -securities

If we consider a list of m -securities traded over n -periods, the approach described above must be extended to a changing portfolio whose risk depends on both individual security volatility and the covariance of price movement across all pairs of securities. To address this issue, we need first to understand how to convert the covariance matrix into monetary units per share and how to compute the total portfolio variance and risk.

Using the same approximation methodology as for individual security volatility, we can convert the covariance of returns of securities i and j as follows:

$$\text{Cov}(r_i, r_j) = E(r_i r_j) - E(r_i)E(r_j)$$

$$\cong \frac{1}{p_i} \frac{1}{p_j} [E(\Delta p_i \Delta p_j) - E(\Delta p_i)E(\Delta p_j)]$$

$$\text{Cov}(r_i, r_j) \cong \frac{1}{p_{0,i}} \frac{1}{p_{0,j}} \text{cov}(\Delta p_i, \Delta p_j)$$

Hence, the covariance of security prices expressed in monetary units/share is:

$$\text{Cov}(\Delta p_i, \Delta p_j) \cong p_{0,i} p_{0,j} \text{cov}(r_i, r_j)$$

$$\cong p_{0,i} p_{0,j} \rho_{i,j} \sigma(r_i) \sigma(r_j),$$

where $\rho_{i,j}$ is the correlation coefficient, and the covariance in monetary units for two orders X_i and X_j is then:

$$\text{Cov}(X_i \Delta p_i, X_j \Delta p_j) \cong X_i p_{0,i} X_j p_{0,j} \text{cov}(r_i, r_j)$$

Expanding to a list of m -securities, the covariance matrix C in monetary units/share is computed from the covariance matrix of return C_r as follows: $C = DC_r D$, where D is the diagonal matrix of the current price of the m -securities:

IV. Estimating Transaction Costs with Pre-Trade Analysis

$$D = \begin{pmatrix} p_1 & 0 & \dots & 0 \\ 0 & p_2 & \dots & 0 \\ \dots & \dots & \dots & \dots \\ 0 & 0 & \dots & p_m \end{pmatrix}$$

From the above approximations for security volatility and covariance matrix, we can easily compute the total variance and risk for a static portfolio of m -securities. If X represents a vector of security positions where x_i is the number of shares held in the i^{th} security and C is the covariance matrix expressed in monetary units per share, we obtain the total risk of the portfolio as follows: $\sigma^2(X) = X'CX$ and $\sigma(X) = \sqrt{X'CX}$.

If we consider now a list of m -securities traded over n -periods, its price risk has to be derived from the changing portfolio approach. If x_k is a vector where $x_{i,k}$ is the number of shares of security i to be traded in period k , r_k is a vector where $r_{i,k}$ is the number of unexecuted shares of security i in period k and C the per-period covariance matrix expressed in monetary units per share, we compute the total risk in monetary units for the specified trading strategy as follows:

$$\sigma^2(x_k) = \sum_{k=1}^n r_k' C r_k \quad \text{and} \quad \sigma(x_k) = \sqrt{\sum_{k=1}^n r_k' C r_k}$$

(b) Liquidity risk

Liquidity risk is the risk of unexpected movements in market conditions. It must be assessed to account for the total risk involving price impact estimates.

To formulate liquidity risk, Kissell and Glantz (2003) start from the following total market impact expression:

$$K(x) = \frac{l}{X} \sum_j \frac{x_j^2}{x_j + 0.5v_j}$$

and then make two assumptions. First, that successive volumes are independent from period to period. Second, that the estimation error of the instantaneous market impact function l may be ignored. Hence, they obtain that liquidity risk is estimated as follows:

$$\begin{aligned} \sigma^2(K(x)) &= \sigma^2 \left[\frac{l}{X} \sum_j \frac{x_j^2}{(x_j + 0.5v_j)} \right] \\ &= \left(\frac{l}{X} \right)^2 \sum_j \sigma^2 \left(\frac{x_j^2}{(x_j + 0.5v_j)} \right) \end{aligned}$$

Since the only random variable in this expression is v_j , the authors solve it with the theorem below, which relies on the assumption that the function is at least twice differentiable.²³

THEOREM

Let v be a random variable with $E(v)=\mu$ and $\sigma^2(v)=\sigma^2$. If $Y=H(v)$ then,

$$E(Y) \cong H(\mu) + \frac{H''(\mu)}{2} \sigma^2$$

$$\sigma^2(Y) \cong [H'(\mu)]^2 \sigma^2$$

If the expected value and variance of each period j is $E(v_j)=v_j$ and $\sigma^2(v_j)=\sigma^2_{v_j}$ respectively, liquidity risk is computed as follows:

$$H(v) = \frac{x^2}{(x + 0.5v)}$$

$$H'(v) = \frac{-x^2}{2(x + 0.5v)^2}$$

$$\sigma^2(Y) = \frac{x^4 \sigma^2(v)}{4(x + 0.5v)^4}$$

Hence, the liquidity risk of the execution schedule K_x is given by:

23 - For more detail, see Kissell and Glantz (2003), pp. 126-127.

IV. Estimating Transaction Costs with Pre-Trade Analysis

$$\begin{aligned}\sigma^2(K(x)) &= \left(\frac{1}{X}\right)^2 \sum_j \sigma^2\left(\frac{x_j^2}{x_j + 0.5v_j}\right) \\ &= \left(\frac{1}{X}\right)^2 \sum_j \frac{x_j^4 \sigma^2(v_j)}{4(x_j + 0.5v_j)^4} \\ \sigma(K(x)) &= \sqrt{\sigma^2(K(x))}\end{aligned}$$

The liquidity risk for a trade list of several securities can be determined as above but by making an additional assumption of independence based on zero correlation of excess volumes²⁴ across securities and across periods for the same securities. So, assuming independence, the liquidity risk for a trade list is:

$$\begin{aligned}\sigma^2(K(x)) &= \sum_i \left(\frac{1_i}{X_i}\right)^2 \sum_j \frac{x_{ij}^4 \sigma^2(v_{ij})}{4(x_{ij} + 0.5v_{ij})^4} \\ \sigma(K(x)) &= \sqrt{\sigma^2(K(x))}\end{aligned}$$

Now that we know how to compute both price and liquidity risk, and assuming that volume and price movement are independent,²⁵ we forecast the total timing risk for a given trade schedule through a simple combination of them, that is:

$$\mathfrak{R}(x_k) = \sqrt{\sum_j r_j^2 \sigma_j^2 + \sum_i \sum_j \left(\frac{1_i}{X_i}\right)^2 \frac{x_{ij}^4 \sigma^2(v_{ij})}{4(x_{ij} + 0.5v_{ij})^4}}$$

It must be said that, in practice, the estimation and incorporation of liquidity risk into the timing risk term is complicated. It is for that reason that the timing risk of a strategy is often referred to only as the price risk of the strategy.

The estimation models and techniques for price appreciation, market impact

and timing risk provide very interesting material that, when we put all the pieces together, makes it possible to determine the cost profile of any trading strategy through simultaneous estimation of cost and risk. This combination of material is the foundation for the development of techniques to determine optimal trading strategies and to derive an efficient trading frontier. We address these points in the next subsection.

3. Optimisation and Efficient Trading Frontier

The cost profile of a specific trading strategy x_k summarises the cost and risk estimates associated with the strategy. It is expressed as $\theta_k = (\phi(x_k), \mathfrak{R}(x_k))$, where θ_k is the cost profile, $\phi(x_k)$ is the expected implicit transaction cost and $\mathfrak{R}(x_k)$ is the forecasted timing risk for the strategy x_k . Based on cost profiles, comparisons of several strategies are easy to make.

Consider a numerical example for an order to be executed and two possible strategies. We show below how we determine the cost profile of each strategy and how we can compare the findings.

Let us first suppose a VWAP strategy for an order of 120,000 shares of XYZ stock. The market price is currently equal to €80 and is expected to reach €82 at the end of the trading day. Daily volatility is estimated at 120bp. The execution strategy and expected market conditions are summarised in table 3 and the instantaneous market impact cost for the stock is estimated at €150,000.

24 - Volume quantities above/below the mean in each trading interval.

25 - There is a negligible degree of correlation between price changes and market volume but the absolute value of price changes and that of volume are indeed correlated. This means that the presence of high volume does not indicate the direction of price movement.

IV. Estimating Transaction Costs with Pre-Trade Analysis

Table 3: Expected market conditions over the trading horizon

Trading period	Expected volume		VWAP strategy	
	Buy	Sell	Executed shares	Residual shares
1	200,000	200,000	24,000	120,000
2	150,000	150,000	18,000	96,000
3	100,000	100,000	12,000	78,000
4	50,000	50,000	6,000	66,000
5	50,000	50,000	6,000	60,000
6	100,000	100,000	12,000	54,000
7	150,000	150,000	18,000	42,000
8	200,000	200,000	24,000	24,000

To determine the cost profile for this VWAP strategy, we must compute the expected price appreciation cost, market impact cost and timing risk. For the latter, we focus on price risk alone. The final results are presented in table 4 and some helpful details for each calculation are noted below.

Table 4: Cost profile of the VWAP strategy

	Implicit costs				Risk
	PA*		MI**	Costs***	
	linear	growth			
Total euros (€)	135,000	134,607	22,768	157,768	70,310
Euros/share (€/share)	1.13	1.12	0.19	1.31	0.59
Basis points (bp/share)	141	140	24	164	73

* price appreciation,

** market impact,

*** total costs with the linear model

Note: As for price appreciation, the linear model delivers a constant price change of 0.25 per trading period while the growth model provides an exponential price change of 0.0031. Concerning timing risk, the stock variance per trading period (1,800bp) is estimated from daily volatility. To express it in monetary units, we multiply it by the square root of the current stock price.

With the same method, we are also able to determine the cost profile of an alternative strategy such as a 10% participation strategy (see tables 5 and 6).

Table 5: Expected market conditions over the trading horizon

Trading period	Expected volume		10% participation strategy	
	Buy	Sell	Executed shares	Residual shares
1	200,000	200,000	40,000	120,000
2	150,000	150,000	30,000	80,000
3	100,000	100,000	20,000	50,000
4	50,000	50,000	10,000	30,000
5	50,000	50,000	10,000	20,000
6	100,000	100,000	10,000	10,000
7	150,000	150,000	0	0
8	200,000	200,000	0	0

Table 6: Cost profile of the 10% participation strategy

	Implicit costs				Risk
	PA*		MI**	Costs***	
	linear	growth			
Total euros (€)	77,500	76,972	30,350	107,850	53,343
Euros/share (€/share)	0.65	0.64	0.25	0.90	0.44
Basis points (bp/share)	81	80	32	112	56

* price appreciation,

** market impact,

*** total costs with the linear model

Comparing tables 4 and 6, we observe now that $\theta_{bp}(VWAP) = (164.73)$ and $\theta_{bp}(10\%) = (112.56)$. So, the second strategy is clearly less costly and risky than the first. In fact, given expected market conditions, the 10% participation strategy offers the lowest risk and cost because it completes trading at the sixth period and does not expose the order to adverse price movements over the entire trading day.

(1) Optimisation formulation

In uncertain market conditions, the preservation of asset value is essential for most investors and we can define the ultimate goal of implementation as asset value preservation in the presence of risk.

IV. Estimating Transaction Costs with Pre-Trade Analysis

This definition means that we need to consider cost and risk jointly. Accordingly, optimisation must be viewed as a means of determining the trading strategy that resolves the trader's dilemma, which will find the suitable trade-off between cost and risk. Indeed, market impact and timing risk are conflicting terms; an optimal strategy results therefore in finding the proper trade-off.

Broadly, optimisation in the context of trading can be formulated in three ways:

- Minimise cost subject to a specified level of risk, that is:

$Min \phi(x_k)$ s.t. $\mathcal{R}(x_k) \leq \mathcal{R}^*$; \mathcal{R}^* is the maximum risk exposure and $\phi(x_k)$ the expected transaction cost.

- Balance trade-off between cost and risk, that is:

$Min \phi(x_k) + \lambda \mathcal{R}(x_k)$; λ is the level of risk aversion or the marginal rate of substitution between cost and risk.

- Maximise probability of price improvement, that is:

$Max Pr[\phi(x_k) \leq L^*]$; L^* is the highest acceptable cost.

Specifically, we can illustrate the second formulation with the following proposal, which incorporates some real-world constraints, from Kissell and Glantz (2003):

Objective function

$$\text{Min} \sum_{i=1}^m \sum_{k=1}^n k \cdot x_{i,k} \Delta p_i + \sum_{i=1}^m \sum_{k=1}^n |x_{i,k}| \left[\frac{0.95 l_i |x_{i,k}|}{|X_i| (|x_{i,k}| + 0.5 v_{i,k})} + \frac{0.05 l_i}{|X_i|} \right] + \lambda \cdot \sqrt{\sum_{k=1}^n r_k^2 C r_k}$$

Constraints

- Completion: $\sum_{k=1}^n x_{i,k} = X_i$

This constraint ensures that optimisation will provide a strategy that executes orders over the specified trading horizon.

- Shrinking portfolio: $|r_{i,j}| \leq |r_{i,j-1}|$

This constraint ensures that the portfolio is continuously decreasing in shares and prevents optimisation from making the position either longer or shorter than the initial position.

- Participation rate: $\frac{|x_{i,k}|}{|x_{i,k}| + v_{i,k}} \leq \alpha$

This constraint places an upper bound on participation volume in each period.

- Cash balancing

$$-D_{min} \leq \sum_i \sum_k x_{i,k} \left[p_i + k \cdot \Delta p_i + \frac{0.95 x_{i,k} l_i}{|X_i| (|x_{i,k}| + 0.5 v_{i,k})} + \frac{0.05 l_i}{|X_i|} \right] \leq D_{max}$$

This constraint is usually used in the implementation of a trade list including buy and sell orders to ensure that the net cash position in any period is within a specified range. In the above approach, the objective function calculates the total cost by adding price appreciation and market impact plus λ units of timing risk. $\lambda=1$ indicates that investors are equally concerned with cost and risk. $\lambda > 1$ (< 1) refers to investors who are more (less) concerned with risk than cost and who thus prefer aggressive (passive) strategies. So, for a given risk aversion, the optimisation output includes the number of shares of each security to be traded in each period (trading schedule) as well as cost estimates for each component (cost profile).

(2) *Efficient trading frontier*

The approach described above is in


IV. Estimating Transaction Costs with Pre-Trade Analysis

the same vein as Markowitz's portfolio optimisation framework, a very popular tool from classic finance theory. It shows the trade-off between risk and return in an investment portfolio: to achieve higher returns, investors must take on more risk. The same is true for the question at hand: to incur lower transaction costs, the trader must implement a riskier strategy. As in portfolio theory, this takes us to the notion of an efficient frontier.

The notion of a *trading frontier* is explored in several academic and professional works that deal with execution optimisation.²⁶ It can be defined as the set of strategies that contain the lowest cost for the specified level of risk, in which strategies are determined through optimisation for all possible values of λ . Figure 11 shows a trading frontier. The cost curve is a convex function in which each point corresponds to a unique degree of risk aversion (λ). This function decreases over the range $0 \leq \mathcal{R} \leq t$ and increases when $\mathcal{R} > t$. Strategy C ($\theta(z,t)$) corresponds to the minimum cost of the frontier. Based on the above optimisation equation, this minimum cost will occur at frontier. Based on the above optimisation equation, this minimum cost will occur at $\lambda=0$ since risk is excluded. For positive values of λ , we obtain points with higher costs and lower risk than at the minimum cost point. For example, strategy A ($\theta(y,s)$) is on the curve to the left of C since $s < t$ and $y > z$. By contrast, points on the curve on the right of C are associated with negative values of λ . It is the case for strategy B whose cost profile is $\theta(y,u)$. An important element here is that when $\lambda < 0$ the optimisation is minimised at the value where the quantity of risk is maximised. This corresponds to a strategy wherein trading is spread over as long as possible a

period. However, we can observe situations in which the incremental price appreciation cost begins dominating the market impact cost reduction, resulting in an increase of total transaction costs. It is for this reason that optimisation with negative risk aversion is not really relevant for investors.

Figure 11: Trading frontier



Like the efficient frontier in portfolio theory, the efficient trading frontier (ETF) is the set of all the dominating trading strategies, that is, the set of strategies that contain the lowest cost for the specified risk and the least risk for a specified cost. It is quite easy to see in figure 11 that only strategies on the curve to the left of C are optimal. Although A and B exhibit the same cost level (y), only A is a true optimal strategy because it contains less risk for the cost ($s < u$). Hence, only strategies lying on the ETF are rational trading decisions. Any other strategy is irrational because at least one alternative strategy can always be found with less risk for the same cost, lower cost for the same risk, or with both lower cost and risk. Consequently, the ETF is a valuable means of identifying the cost profile of optimal trading strategies and of choosing the most appropriate execution (in view of the investor's preferences or final objectives).

²⁶ - See, for example, Almgren and Chriss (2000) and Nevmyvaka et al. (2005).

V. Trading Performance Measurement



V. Trading Performance Measurement

The last task assigned to TCA is the assessment of trading performance. Like transaction costs, trading performance must be measured after the execution of the trade. Here, it is crucial to understand clearly how evaluation of trading performance differs from transaction cost measurement, although both are done *ex post*.

1. Transaction Costs vs. Trading Performance

Even today there is widespread confusion of transaction cost and trading performance measures; this confusion leads to the frequent industry adoption of inappropriate practices. Although the main reason for this ongoing confusion is undoubtedly the elusive and sometimes controversial notion of *best execution* (we will come back to this important notion), it must be said that some in the industry also contribute to its currency. Some intermediaries, for example, have incentives to argue that trading performance evaluation is extremely complex and that it should be founded on market impact that cannot be observed or measured. They can thus consistently come up with arguments, some more relevant than others, for declaring benchmarks or peer groups unfit for assessing the quality of their trades. In this way, they avoid being subjected to performance assessment. In addition, some consultants have a vested interest in this confusion, since it justifies their lucrative business. In any case, effective transaction cost management would require clarification.

Post-trade analysis consists of two parts: transaction cost measurement and trading

performance assessment. As we have seen, transaction cost measurement attempts to determine the source and magnitude of trading costs. Trading performance assessment, by contrast, looks into the justifications for these costs and seeks to identify any costs caused by poor implementation decisions. In other words, performance evaluation attempts to assess the performance of intermediaries when they execute trades. The ultimate aim is to determine the most effective intermediaries by market, instrument and trading strategy. With this information at hand, investment managers can considerably reduce the time required to select the best intermediary for the execution of a specific trade.

Given the above objective, a good trading performance measure should help investors determine if the actual transaction costs incurred were reasonable under the market circumstances that prevailed at the time of the trade. Proper performance evaluation must therefore assess the real capability of intermediaries and, to a larger extent, determine if they add value. Distinguishing between skill and luck can make it possible to determine if execution can indeed be termed best.

2. Current State of the Industry and Practices

The most common performance measurement practice in the industry is the benchmark comparison described above. So, here we address its shortcomings as a gauge of trading performance. Next, we will introduce the relative performance measure developed by Kissell and Glantz (2003). Although this measure is not

V. Trading Performance Measurement

widely used, it appears to be an improved metric that is more consistent and stable because it allows comparisons across securities and across days. Furthermore, this approach shares the principle of peer group review with the innovative and systematic framework for assessing the quality of execution that we propose and will develop further in the document.

(1) Benchmark comparison

This approach involves comparing the monetary value difference between the executed position and the position evaluated at some benchmark price. For a given transaction, the monetary value per share is given by the signed difference between the average price obtained for the trade and the benchmark price. A positive value indicates an execution that is less favourable than the benchmark price, while a negative value indicates execution more favourable than the benchmark. As we have mentioned above, the benchmarks usually used include pre-trade, intraday and post-trade reference prices.

(a) Critical review of benchmarks

The *pre-trade benchmarks* (T-1 close, T open, last trade price, last bid, last ask) provide transaction cost measures, not performance measures. They deliver good proxies for the costs of spread, market impact, operational and market timing delays, but they do not provide any insight into actual execution performance since they do not help gauge whether the costs incurred are acceptable given market conditions at the time of trading.

The *intraday benchmarks* (VWAP, LHOC) compare the execution price and a kind of average market price of the day or of a

given time period. The belief behind these benchmarks is that you did a good job when you did better than the "average trader" of the day or period. In this sense, intraday benchmarks are the most consistent with a performance measure, assuming that the benchmark is a good indicator of the fair market price. However, although intraday benchmarks account for market conditions and trading activity, they do not provide a meaningful metric that is consistent across days and across stocks. We will look into this key problem later.

The *post-trade benchmarks* (T close, next mid bid-ask) attempt to measure the market impact of trades. When used in trading performance analysis, they serve more as a measure of investor skill than as a measure of trader performance. For funds that track a benchmark index, post-trade benchmarks indicate the contribution of execution to total tracking error.

(b) Major shortcomings

As a whole, when used to assess trading performance, benchmark comparison has major shortcomings that make it inefficient and/or misleading. These shortcomings are listed and described below.

Benchmark comparison may be biased

The benchmark comparison may provide biased performance measures that essentially depend on the reason for the trade. According to Harris (2003), biases may arise when trading decisions depend on past price changes or when investors are well informed about future price changes. Some benchmark prices can deliver performance measures that will be systematically good or bad depending on whether the investor uses a momentum

V. Trading Performance Measurement

or contrarian strategy. For example, the opening price delivers performance measures that are easily biased. Momentum investors will push down the trading performance measure because they buy (sell) when prices have risen (fallen), so the opening price is low (high). By contrast, contrarian investors buy (sell) when prices have fallen (risen), so the opening price is high (low). This pushes up the trading performance measure.

Benchmark comparison may be gamed

Gaming problems arise when intermediaries know the benchmark that will be used to measure their execution performance and consequently time their trades based on their evolving benchmark score.

The consideration of trade-timing is important for investors who give their intermediaries discretion over the timing of their trades. Proper trading performance measures should help investors ensure that their brokers make appropriate use of this discretion. Trade-timing effects are best measured when the benchmark does not depend on the time of the trade. When the benchmark relates to the execution time, intermediaries can accelerate/decelerate trades based on their evolving benchmark score. This behaviour is questionable because it can result in higher total transaction costs for the end investor.

Harris (2003) describes several situations in which brokers can game their evaluations. For example, brokers who have discretion over how aggressively they fill orders can easily game a performance measure based on the spread midpoint prevailing at the time of the trade. They will simply supply

liquidity and never take it. Hence, they always buy at the "bid" or sell at the "ask". This behaviour makes their performance look great but comes at the expense of increased timing risk.

Benchmark comparison needs a unique and relevant reference price

Benchmark comparison assumes that the benchmark is an appropriate price for the value of the security. The only "true" value of a security is the price at which an actual trade was made. Sometimes, the benchmark is not such a price. For example, the spread midpoint, the VWAP or the LHOC may be prices at which no trade has taken place. Furthermore, market fragmentation and proliferation of liquidity pools make the determination of the right benchmark more difficult. Indeed, the benchmark comparison becomes conceptually difficult in the absence of a recognised unique price for the securities traded. When a security is traded at the same time on various execution venues, which of the coexisting prices is the benchmark price for the security? With MiFID around the corner, this issue is of great importance. By allowing multiple trading venues to compete, this new piece of European regulation is, by design, allowing liquidity pools to fragment and putting the very existence of recognised benchmarks in danger.

Benchmark comparison is not standardised

Finally, for at least two reasons, benchmark comparison does not offer a unified framework to enable easy assessment of execution performance across a series of trades at any aggregate level.

V. Trading Performance Measurement

First, benchmarks provide absolute measures of costs, which are often expressed in basis points. Determining whether the trading performance is bad, normal or good is not straightforward and benchmarks do not make comparison easy because they deliver inconsistent metrics. For example, a measure of 10bp could mean good performance in one security but bad performance for another. Similarly, it could also mean good performance today but bad tomorrow. Next, existing measures are built on benchmark prices that often depend on trade characteristics. Consequently, using only one benchmark price to analyse the execution quality of a universe of trades can be criticised, unless all the trades take place at the same time and are of a similar size and difficulty.

The absence of a standardised framework has led to the absence of consensus about a universal indicator and accounts for the availability in the industry of a great many indicators. As a result, investors are often confused.²⁷ Furthermore, from the absence of consensus emerges the difficulty of comparing trading performance and execution quality across different intermediaries. The heterogeneity of trades and diversity of measures complicate any attempt at comparative analysis of trading performance.

(2) Relative performance measure

The relative performance measure (RPM) is a metric proposed by Kissell and Glantz (2003). Founded on peer group review, it compares the average execution price of the trade to all market activity over the trading period. For a given trade, the RPM is the percentage of all activity in the market that traded at a price less favourable than

its average execution price. Modelled after the percentile ranking, the RPM attempts to provide a descriptive and meaningful measure of trading performance that is consistent across days and across stocks. Not unlike VWAP-based measures, in that it properly accounts for contemporaneous market conditions and trading activity, the RPM is more robust because it allows multiple comparisons.

(a) Calculation and interpretation

The RPM is built on a combination of both market volume and number of trades, which makes it possible to account for the potential skew of large block trades at extreme prices. Accordingly, if P^* is the average execution price of the order and P_i is the price of the i^{th} trade, the RPM is calculated as follows for a buy:²⁸

$$RPM = \frac{1}{2} [RPM_{trades} + RPM_{volume}] , \text{ with}$$

$$RPM_{trades} = \frac{\sum trades_{P_i > P^*}}{\sum trades} \text{ and}$$

$$RPM_{volume} = \frac{\sum volume_{P_i > P^*}}{\sum volume}$$

Graphically, we present the RPM by plotting the percentage of activity traded at a price less favourable than or equal to the average execution price. For this purpose, activity needs to be sorted from lowest to highest prices for sells and, by contrast, from highest to lowest prices for buys. Figure 12 illustrates the RPM calculation for a sell order with an average execution price of €31. The figure shows an RMP of 80% for the order, meaning that 80% of all market activity was traded at prices lower than €31. The interpretation is that


27 - The domination of the sell-side in the TCA services provision is also problematical. TCA tools and services have been at the heart of the development of the sell-side industry over the last five years. With the multiplication of electronic trading capabilities (direct market access or algorithmic trading), brokers have endeavoured to provide substantial tools for their clients to help them measure transaction costs and, to a larger extent, assess trading performance. Created as independent boutiques, most TCA providers were quickly acquired by the major brokerage firms so that they could become a significant and visible element of the sell-side value proposition. That the sell-side has totally taken over the market for TCA confirms the strategic importance of this function for the industry.

28 - The RPM calculation is symmetric for a sell order ($P_i < P^$).*

V. Trading Performance Measurement

the trader did a good job since 80% of market activity traded at less favourable prices than he did. If the order had been a buy, all market activity higher than €31 would deliver a RPM of 20%, indicating poor performance.

Figure 12: Relative performance measure



29 - We have a similar phenomenon on the sell side. An aggressive strategy in a falling (rising) market will deliver a high (low) RPM for sells, suggesting good (poor) performance. A passive strategy will result in the opposite.

The RPM is superior to any other benchmark-based performance measure because it provides a meaningful metric that is consistent across days and across stocks. Suppose, for example, that a trader's performance measure built on a benchmark is 15bp in stock A and 40bp in stock B. As we mentioned earlier, this information cannot help determine if this was good or bad performance or make any comparisons across stocks. By contrast, if the trader has an RPM of 20% in stock A and 75% in stock B, we can easily determine that the trader did badly in stock A but very well in stock B. This means that the trading performance was better in stock B than in stock A, even though the trader has a better benchmark score in stock A. These conclusions are meaningful and easy to reach with the RPM since it is consistent over time and across stocks.

Furthermore, using a sufficient number of observations and analysing the distribution of RPM facilitates a reading of the consistency of intermediaries. Those

who exhibit little deviation in performance measures are consistent, while those who often have extreme scores are gambling or taking risks.

(b) Weighted RPM

When intermediaries have discretion over how they execute orders within a specific horizon, the RPM based on all market activity traded in that horizon reflects their ability to work orders and seize the best available trading opportunities. However, when the trader is given a predefined execution strategy, using the RPM over the entire trading horizon may provide noisy performance measures when the instructions or constraints set by the investor are not neutral. This is the case when they force traders to execute the majority of shares at times of the least or most favourable prices. For example, an aggressive strategy in a falling (rising) market will refer to a low (high) RPM for buys, suggesting poor (good) performance.²⁹

To avoid noise and really distinguish between trading performance and investor constraints, the RPM must be adjusted upon the specified strategy as follows:

$$RPM^* = \sum_j \frac{x_j}{X} RPM_j ,$$

where RPM_j is the RPM computed over the trading period j , x_j is the quantity to execute in period j and X the total order size. This weighted RPM attempts to provide insight into the quality of prices obtained by the trader during the times he is requested to trade.

(c) Measure of value-added

When intermediaries deviate from

V. Trading Performance Measurement

instructions or constraints set by the investor (on purpose or not), the final result may be higher or lower total transaction costs. In such situations, what is important beyond the RPM is to determine whether the trader's decision has added value (lower costs) or hurt overall performance (higher costs). Let us take an example to illustrate this concern.

Suppose a trader who is requested to execute a given order before midday but decides (for whatever reason) to work it through the entire day. Although he minimises market impact cost, the investor incurs higher total transaction costs because of adverse price movement. It is possible for the trader to get a favourable RPM but it will not give any insight into the effect of his decision on final performance.

To address this issue and quantify the value-added by the trader's own market timing, Kissell and Glantz (2003) define a new metric (VA) to calculate the percentage of the total RPM attributable to the deviation decision, that is:

$$VA = \frac{RPM(x^*) - RPM(x)}{RPM(x^*)}$$

$$= \frac{\sum_j (x_j^* - x_j) RPM_j}{RPM(x^*)}$$

where $RPM(x)$ is the RPM that the trader would have got had he followed investor instructions/constraints and $RPM(x^*)$ is his actual RPM.³⁰ So, by comparing the deviation in strategy, we can assess the value that the market timing of the trader adds to the implementation.

It is easy to see that, when traders follow the prescribed strategy, the value-added is zero. When they deviate, a positive value-added indicates good market timing ability (the trader achieves better prices and lower costs) and, by contrast, a negative value-added means poor market timing ability (the trader achieves worse prices and higher costs). The VA metric thus assesses the real capability of intermediaries and contributes to distinguishing between skill and luck.

A step further is to assess the overall contribution to cost attributable to the trader's market timing. This could be done by plotting together on a chart the normalised difference between the actual and the expected transaction costs and the VA of the trader.

(d) RPM vs. VWAP

The RPM is somewhat similar to VWAP-based measures since both properly account for the prevailing market conditions by comparing the average execution price and all other contemporaneous market activity. In that sense, both measures suffer from the same shortcoming. When an order accounts for the main part of market activity in the actual trading period, its average execution price converges to the VWAP computed over that period and the RPM converges to 50%, suggesting average performance. In this case, interpreting these results becomes quite hard since the execution being analysed is its own reference.

The main advantage of the RPM over any VWAP-based indicators is, as already mentioned, that the RPM is more robust because it allows multiple comparisons, i.e., across securities and in the same

³⁰ - For periods where $x_j=0$, the RPM_j is computed assuming that the trader achieves the average price in the period.

V. Trading Performance Measurement

security across days or periods. This is due to the consistency of the RPM, which makes comparison easy. A performance of 75% means exactly the same thing whatever the day, the security or even the trader. Furthermore, unlike the RPM, the VWAP exhibits an asymmetric distribution since most market activity is likely to be on one side of the VWAP. This means that 50% of prices and volume are not higher or lower than the VWAP. We can thus observe situations in which intermediaries underperformed the VWAP but still outperformed the majority of market activity. In those cases, the VWAP scores may lend themselves to misleading interpretations.

3. Regulatory Pressure: MiFID and its Best Execution Obligation

The absence of a standardised framework for easy assessment of trading performance becomes problematic with the arrival of the MiFID. Indeed, with this new piece of European regulation, investment firms are going to enter an environment where they will have to demonstrate that they have executed at the best possible conditions for their clients while taking into consideration potential multiple liquidity pools. To address this issue, we will first focus on the concept of best execution in general. We will see that nowadays this concept is very fashionable in the industry, although it is not always well understood and does mean not the same thing to everyone. Then, we will move on to the MiFID best execution obligation and show that the regulator has provided neither a clear definition nor a measurable objective to make up for the current absence of consensus in the industry.

(1) Best execution, this elusive concept

Best execution means many things for many people, but everybody agrees that it includes at least trading performance assessment. In theory, best execution is often defined as a measure of how well investors' trades are executed. This definition is vague as it encompasses components such as the trade price, the execution speed, the opportunity for price improvement, the probability of full execution, the respect of anonymity and the level of explicit costs. Although the total charges paid by the investor are still the main dimension, the others are not to be ignored.

Everyone admits that the lowest transaction cost does not necessarily imply best execution, but the consideration of all components intensifies debate when it comes to checking whether best execution has been achieved, essentially because some of these components are not easy to measure and/or combine. Either the information is not directly available (execution speed, price improvement) or the component itself is not measurable (anonymity is respected).

According to Kissell and Glantz (2003), best execution can be categorised as price, time, and size factors whose importance varies with the goals and objectives of investors. For example, value and passive investors are concerned mainly with price factors, growth and momentum investors focus on time factors, while large investors and mutual funds pay special attention to size factors. In this respect, Kissell and Glantz (2003) define best execution as "the whole process of managing transaction costs

V. Trading Performance Measurement

throughout all phases of the investment cycle to ensure that the portfolio realizes its highest returns possible”.

Even if the above definition is not yet widely recognised in the industry, we strongly believe that it is the right way to clear up most of the confusion surrounding best execution. Consistent with this definition, Kissell and Glantz (2003) claim that checking whether best execution is achieved results in analysing the entire transaction cost management process. This thorough check thus requires a complex post-trade procedure that should consist of the following steps:

1. Evaluate the implementation decision: Is the trading strategy optimal? Does it lie on the ETF?
2. Measure transaction costs actually incurred: Where and why did they occur?
3. Estimate transaction costs *ex post*: What are cost estimates with known market conditions?
4. Compare actual transaction costs with *ex-post* cost estimates: Is there over/underperformance?
5. Measure execution performance: Has the intermediary executed at fair market prices? Did he add value?
6. Compare the cost difference with the value-added: Is over/underperformance related to superior/poor execution or to favourable/adverse price movement?
7. Record performance of intermediaries

Although this procedure looks attractive and quite complete from a theoretical standpoint, it is rarely applied, as debate on the definition of best execution and/or trading performance measurement is ongoing.

(2) MiFID best execution obligation

This new rule is a crucial element of investor protection and must be viewed as the natural counterpart of a full liberalisation of the marketplace. This obligation is found in the now famous Article 21, which has been at the heart of the arguments put forward by many opponents of MiFID. After having summarised the main details of the provision, we will show that although MiFID puts the spotlight on best execution, it has not provided a clear definition or a measurable objective that would make it possible to determine whether trades have been executed in the best possible fashion.

In the 2004 Directive, the best execution obligation has been defined as an obligation of means whereby investment firms are required to have taken all reasonable steps to obtain the best possible result for the client. This obligation is therefore structured around three major principles: (i) an obligation of means to achieving the best possible result for the client, involving factors that determine whether or not this best possible result has been achieved;³¹ (ii) documentation of an execution policy that includes the selected execution venues and documentation of the parameters that justify this selection; (iii) an obligation for investment firms to demonstrate, at the demand of the client, that execution has been carried out in accordance with the agreed execution policy and that the execution policy allows achievement of the best possible result on a consistent basis.

MiFID considers best execution in the context of client categorisation and in a tiered manner. Accordingly, the assessment

31 - Article 21 states, however, that whenever there is a specific instruction from the client the investment firm shall execute the order following the specific instruction.

V. Trading Performance Measurement

of best execution for professional clients must consider factors that include not only fees but a wide range of elements that can be used to qualify the quality of a trade, namely "price, costs, speed, likelihood of execution and settlement, size, nature or any other consideration relevant to the execution of the order". Next, retail clients are said to have achieved best execution when they have received the best price net of expenses. This means that the net price should usually be the most important element; the other factors play a limited role. Finally, trades executed between firms categorised as "eligible counterparties"³² fall outside the scope of the best execution obligation.

in what the best execution obligation is able to achieve with regard to the fair protection of all parties.³³

32 - This category refers to the most experienced firms operating in the marketplace, such as professional trading firms and asset managers executing client orders.

33 - For more detail on MiFID and its best execution obligation, see D'Hondt and Giraud (2007).

Begun as an obligation of result in a principle-based regulatory approach, the best execution obligation has been actively fought by industry representatives and has slowly turned into a more modest obligation of means that remains complex and ambiguous. Furthermore, as a direct result of the underlying complexity of defining best execution, the regulator has shifted the focus to the means rather than simply to the result. By providing only partial guidelines on how investment firms should assess the quality of execution through a double system of criteria and factors to be taken into consideration, but without determining how they will actually work, the regulator has frustrated all industry participants, with those seeking a principle-based approach left only with an overly prescriptive rule and those seeking the definition of precise criteria for complying with the obligation left with too much flexibility in the application of Article 21. This situation has led to significant resistance and mistrust

VI. A New Framework: the EBEX Indicators



VI. A New Framework: the EBEX Indicators

In the post-MiFID environment, the role of TCA is set to take on even greater importance. As we have seen, however, the current absence of a standardised and widely recognised framework for assessing the quality of the entire trading process is problematic. In this section, we try to fill this gap first by offering a unified framework for measuring *ex post* the quality of execution and then by showing how this framework can be incorporated into the analysis of the entire transaction cost management process.

Because benchmark comparison suffers from several shortcomings, we have developed a standardised framework in the same vein as the RPM³⁴ and have opted for an absolute measure of the quality of the price obtained: a score between 0 (bad performance) and 1 (good performance).

Our approach, the EBEX framework (EDHEC Best Execution), is founded on indicators that facilitate comparisons of a large universe of trades and provide insightful information not only about final performance (the absolute EBEX indicator) but also about the possible justification of the performance (the directional EBEX indicator); a measure of the quality of the market timing is thereby provided.

EBEX has several advantages over current industry practices. It is very simple, provides a standardised framework for assessing the quality of execution across a series of trades aggregated at any level, and complies with the MiFID requirement to demonstrate that the target has been reached. It also delivers absolute and meaningful measures of execution quality, allowing straightforward and objective interpretation, and includes

trade-timing consideration. EBEX allows the investor to determine easily whether or not the trading performance of his intermediary, trader or even algorithm is consistently at the top of the class.

EBEX is being actively discussed by professionals and academics, and we are confident that positive developments will be proposed in the very near future to allow it to cope with specific situations that the first version did not make allowances for, such as investor constraints or the cost for the investor associated with the broker performance.

After having exposed the objectives and principles behind this new methodology, we will describe in detail how our indicators are built, as well as how they can be interpreted. Then, to illustrate both the framework and the level of interpretation made possible, we will present the results of an empirical study conducted on a sample of orders for Euronext blue chips. Next, we will compare EBEX and other approaches to highlight the advantages of our framework.

1. General presentation

Our approach aims to provide a simple answer to the following question:

Given a transaction handed over to a broker, trader or algorithm and executed for a given price at times that are recorded under given time constraints, to what extent have other brokers, traders or algorithms executed comparable volumes to this transaction, either before or after this transaction, at a better price?

34 - We became aware of the existence of the RPM after having defined our own tools.

VI. A New Framework: the EBEX Indicators

The answer to this question can be split into four important elements:

- The time at which the order is handed over (release time) to an intermediary (a broker, a trader or an algorithm) is the first point of reference, while the time at which the order is entirely filled (execution time of the last lot referring to the order in the event of splitting) is the second point of reference. The third and final point of reference is the next market close if the intermediary has discretion over the day to execute the trade or, in the event of specific instructions from the investor, the time at which the intermediary must stop trading.
- The size of "competing trades" is not important as such; the relevant measure is how many times a volume comparable to the order has been executed at a better price, which is a first measure of the quality of the price obtained. The price must be compared to small trades executed at better prices (the broker, trader, or algorithm could have split the order better) as well as with larger trades (the order could have been grouped with a larger flow of orders to be executed in block if such trading capability is offered). This important point makes our approach different from the RPM, which is built on a combination of both the number of trades and volume.
- Volumes traded before at a better price allow one to assess whether the broker, trader or algorithm has been too patient.
- Volumes traded after at a better price allow one to assess whether the broker, trader or algorithm has been too aggressive.

EBEX relies on the elements above to measure the quality of execution as part of a peer group review and to identify whether the broker, trader or algorithm has implemented the execution too aggressively or too slowly.

This approach is built on two indicators. The Absolute EBEX indicator measures the quality of execution in a peer group review. The Directional EBEX indicator identifies whether the broker, trader or algorithm has implemented the execution too slowly or too aggressively. In other words, the first indicator assesses the quality of execution itself while the second indicator brings information about why the quality of execution is as observed.

At this stage, we can identify similarities as well as differences between the EBEX framework and the RPM. Both involve peer group review and compare the average execution price of the trade to all contemporaneous market activity. Although they have the same philosophy, the EBEX framework makes it possible to gain quick insight not only into the final trading performance (the absolute EBEX indicator) but also into the possible justification of the performance (the directional EBEX indicator), thereby providing a direct measure of the quality of the market timing. This point is very relevant, as we know that pre-trade analysis should help intermediaries predict likely market conditions and define an appropriate trading strategy when they have discretion over how they fill the order. Getting the best price in the market may not always be a realistic objective, but execution at fair market prices certainly is. What is relevant for an investor then is to know where, with respect to the industry average, the orders executed by his intermediary fall.

VI. A New Framework: the EBEX Indicators

2. Detailed Presentation of the Indicators

Our two fundamental indicators rely on the same method and are easy to compute and interpret. We describe in depth how they are built and highlight their multiple advantages over the benchmark comparison approach. Next, we present another related indicator whose purpose is to associate the intermediary performance score with a cost for the end investor. We finish by describing how our indicators can be adjusted to account for specific trading instructions and constraints.

(1) Absolute EBEX

The absolute indicator of EDHEC Best Execution for an order is the difference between one and a ratio between the aggregate volumes traded at prices better than the average trade price obtained for the order divided by the order size and the aggregate volumes without consideration of price divided by the size of the order. As illustrated in figure 13, the measurement window for the ratio corresponds to the period over which the intermediary has discretion to execute the trade, that is, from the time the broker receives the order (release time) to the next market close.³⁵

Figure 13: Time window for the Absolute EBEX



*35 - In the event of specific instructions from the investor, the measurement window ends at the specified time the intermediary must stop trading.
36- S_i can be simplified but is shown for clarity.*

The absolute EBEX indicator for order i is calculated (first for a buy, then for a sale) as follows:

$$EBEX_{abs,i} = 1 - \frac{\sum_{n=1}^N V_{n,day}^{P < AP_i}}{(S_i)} \div \frac{\sum_{m=1}^M V_{m,day}}{(S_i)}$$

$$EBEX_{abs,i} = 1 - \frac{\sum_{n=1}^N V_{n,day}^{P > AP_i}}{(S_i)} \div \frac{\sum_{m=1}^M V_{m,day}}{(S_i)}$$


In both equations, each element is defined as follows:

- $EBEX_{abs,i}$ is the absolute best execution indicator for order i during the day
- day is the interval between the time the broker receives order i and the next market close
- S_i is the size of order i ³⁶
- AP_i is the average trade price obtained for order i
- N is the number of trades at a price better than AP_i during the time interval
- $V_{n,day}^{P > (<) AP_i}$ is the size of trade n at a price higher (lower) than AP_i during the interval day
- M is the total number of trades during the time interval day ; $M \geq N$
- $V_{m,day}$ is the size of trade m during the time interval day

Given the way it is built, the absolute EBEX indicator can only take values between zero and one. Interpretation is thus greatly facilitated, as shown in figure 14 below.

VI. A New Framework: the EBEX Indicators

Figure 14: Interpretation of the Absolute EBEX indicator



The absolute EBEX indicator is superior to any other performance measure based on benchmark comparison because it provides a meaningful metric that is consistent over time and across securities. Two simple examples illustrate this advantage.

Example 1: Suppose two orders for which the benchmark comparison to the VWAP is 15bp in stock A and 40bp in stock B. Both executions are less favourable than the VWAP but one could be tempted to believe that the execution quality is higher in stock A, since 15bp is lower than 40bp. In fact, this way of thinking is misleading, as the difference may be caused by differences in market conditions, trading patterns, or stock prices. So, even if order size is comparable, the values themselves are not. By contrast, if we observe an absolute EBEX of 0.2 in stock A and 0.75 in stock B, we can conclude that the trader did a bad job in stock A but an excellent job in stock B. This means that the trading performance was better in stock B than in stock A, although the trader has a better benchmark score in stock A.

Example 2: Suppose a single stock for which the benchmark comparison to the VWAP is 40bp today but was 30bp yesterday. Although the stock is identical, the values

still depend on market conditions as well as price volatility. If the price range was less than 40bp yesterday, a value of 30bp indicates that the order got one of the least favourable prices of the day. Furthermore, if the price range today is much larger than it was yesterday, a value of 40bp implies better trading performance, even if the deviation from the VWAP is higher. So, benchmarks do not make comparison easy at all. By contrast, if the absolute EBEX is 0.65 today and was 0.45 yesterday, we can conclude that the quality of execution is better today.

The second advantage of the absolute EBEX is that it makes possible evaluation of intermediary trading performance consistence. This evaluation can be done easily by using a sufficiently large sample of trades executed by the intermediary and analysing the distribution of absolute EBEX. Traders who exhibit little deviation in performance measures are consistent while traders who often get extreme scores are gambling or taking risks. Moreover, comparative analysis across several intermediaries may easily be conducted based on absolute EBEX distributions wherein the score of each intermediary is defined as a risk-adjusted trading performance computed as the mean value of EBEX divided by its standard deviation. We illustrate this approach with the following example.

Suppose two brokers who executed approximately the same number of orders on similar stocks. According to their respective absolute EBEX distribution, broker A exhibits an average absolute EBEX of 0.65 for a standard deviation of 0.45 while broker B provides an average EBEX score of 0.55 for a standard deviation of 0.33. These figures

VI. A New Framework: the EBEX Indicators

suggest that, on average, broker A performs better than broker B (0.65>0.55) but that he is also less consistent (0.45>0.33).

As a result of what we have just seen above, the absolute EBEX indicator may be readily used to set a clear performance objective for any intermediary. Accordingly, an active trader will be expected to attempt to provide an average EBEX as high as possible with the lowest standard deviation. For example, any investor could request his intermediary to reach a defined average absolute EBEX to justify the use of an active market timing strategy. This performance objective would be very relevant because the absolute EBEX indicator cannot be gamed, as it involves post-trade peer group analysis.

37 - Or any other intermediary.

Given these advantages, the absolute EBEX method offers tools that comply with MiFID requirements. Indeed, under MiFID, traders must demonstrate that they have executed at the best possible price for their customers while taking into consideration potential multiple liquidity pools. For this purpose, it is necessary to compare each trade with the universe of all the trades referring to the same security and executed on all available trading venues. As our approach involves peer group analysis and theoretically includes as many trading venues as required, it is a suitable and relevant tool that should play an important role in the post-MiFID environment.

(2) Directional EBEX


Directional EDHEC Best Execution for an order indicates how the broker³⁷ could have traded over time to provide better execution. This indicator results from the combination of two sub-indicators that measure the volumes traded at better prices both before

and after that the trade was executed. The directional EBEX indicator for order i is computed as follows:

$$EBEX_{dir,i} = NBBEX_{i,j} - NABEX_{i,t}$$

$NBBEX_{i,j}$ and $NABEX_{i,t}$ are thus the components of the directional EBEX indicator. Their definition and computation are very similar; only the measurement window of reference differs, as shown in figure 15.

Figure 15: Time window for the Directional EBEX



$NBBEX_{i,j}$

$NBBEX_{i,j}$ stands for **N**umber of **B**efore-**B**etter **E**xecutions for order i over the time interval j . This component can be defined as a ratio between the aggregate volumes traded at a price better than the average trade price of order i divided by the size of order i and the aggregate volumes without consideration of price divided by the size of order i . This ratio is computed over the interval j , which goes from the time the broker receives order i (release time) to the time order i is completely filled (execution time).

The mathematical notations referring to $NBBEX_{i,j}$ are given below (first for sell orders, then for buy orders):

$$NBBEX_{i,j} = \frac{\sum_{n=1}^N V_{n,j}^{P > AP_i}}{(S_i)} \div \frac{\sum_{m=1}^M V_{m,j}}{(S_i)}$$

VI. A New Framework: the EBEX Indicators

$$NBBEX_{i,j} = \frac{\frac{\sum_{n=1}^N V_{n,j}^{P < AP_i}}{(S_i)}}{\frac{\sum_{m=1}^M V_{m,j}}{(S_i)}}$$

In both equations, each element is defined as follows:

- $NBBEX_{i,j}$ is the number of better executions for order i during the interval j
- j is the interval between the time the broker receives order i and the time order i is completely filled
- S_i is the size of order i ³⁸
- AP_i is the average trade price obtained for order i
- N is the number of trades at a price better than AP_i during time interval j
- $V_{n,j}^{P > (<) AP_i}$ is the size of trade n at a price higher (lower) than AP_i during the interval j
- M is the total number of trades during the time interval j ; $M \geq N$
- $V_{m,j}$ is the size of trade m during time interval j

$NABEX_{i,t}$

$NABEX_{i,t}$ stands for **N**umber of **A**fter-**B**etter **E**xecutions for order i over the time interval t . This component is a ratio between the aggregate volumes traded at a price better than the average trade price of order i divided by the size of order i and the aggregate volumes without consideration of price divided by the size of order i . This ratio is computed over the interval t , which starts at the time order i is completely filled (execution time) and ends at the market close of the day.³⁹

The mathematical notations referring to $NABEX_{i,t}$ are given here below, for sell orders and buy orders respectively.

$$NABEX_{i,t} = \frac{\frac{\sum_{n=1}^N V_{n,t}^{P > AP_i}}{(S_i)}}{\frac{\sum_{m=1}^M V_{m,t}}{(S_i)}}$$

$$NABEX_{i,t} = \frac{\frac{\sum_{n=1}^N V_{n,t}^{P < AP_i}}{(S_i)}}{\frac{\sum_{m=1}^M V_{m,t}}{(S_i)}}$$

In both equations, each element is defined as follows:

- $NABEX_{i,t}$ is the number of better executions for order i during the interval t
- t is the interval between the time order i is completely filled and the next market close
- S_i ⁴⁰ is the size of order i
- AP_i is the average trade price obtained for order i
- N is the number of trades at a price better than AP_i during time interval t
- $V_{n,t}^{P > (<) AP_i}$ is the size of trade n at a price higher (lower) than AP_i during the interval t
- M is the total number of trades during the time interval t ; $M \geq N$
- $V_{m,t}$ is the size of trade m during time interval t

Now that both components of the directional EBEX indicator have been presented, we can focus on how they can be interpreted to characterise the timing of the trade. This interpretation is very easy because

38 - S_i can be simplified but is displayed for the sake of clarity.

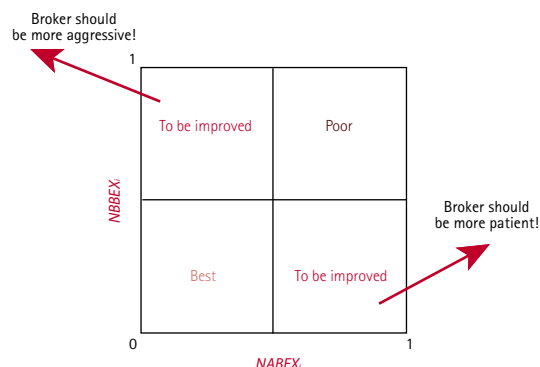
39 - In the event of specific instructions from the investor, the measurement window ends at the specified time the intermediary must stop trading.

40 - S_i can be simplified but is displayed for the sake of clarity.

VI. A New Framework: the EBEX Indicators

both range from 0 to 1, given the way they are built. Figure 16 contributes to an understanding of the interpretation.

Figure 16: Interpretation of the Directional EBEX components




NBBEX is close to zero when few traders execute at more favourable prices before order execution. In this case, the quality of execution may be said to be high. The intermediary did a good job because few traders obtained better prices. By contrast, NBBEX is close to one when many traders got more favourable prices before the order execution time. Hence, the quality of execution is low and the intermediary should have been more aggressive. Indeed, the intermediary would have had more opportunities to trade at better prices before the order was executed.

The way we can interpret NABEX is similar, except that we focus on what happens after order execution. NABEX is close to zero when few traders obtained more favourable prices after the execution of the order. The quality of execution is then high because few traders execute at better prices than the intermediary did. NABEX is close to one when many traders obtained better prices, indicating a poor trading performance. In this case, the intermediary did a bad job in the sense that more patience should have been shown. The intermediary would have

had more opportunities to trade at a better price after the order was executed.

A direct comparison of both indicators yields our directional EBEX indicator, interpretation of which is even easier. The goal of this indicator is to give information about how the intermediary could have traded over time to provide better execution. Given its construction, the simple difference between NBBEX and NABEX, the directional EBEX indicator can range from -1 to +1. Figure 17 (below) summarises the interpretation of the directional EBEX indicator.

Figure 17: Interpretation of the Directional EBEX indicator



The directional EBEX indicator exhibits a negative value when NBBEX is lower than NABEX. In this situation, we can say that the intermediary should have been more patient because more opportunities would have arisen to enable trading at a better price after the order execution. By contrast, the directional EBEX indicator is positive when NBBEX is larger than NABEX. In this case, the intermediary should have been more aggressive because more opportunities would have arisen to trade at a better price before the order execution.

The specific situation in which NBBEX is equal to NABEX corresponds to a directional EBEX of zero. This means that there were

VI. A New Framework: the EBEX Indicators

as many better executions before as after the execution of the order. In this case, the intermediary could have traded at any other time to provide better execution. This specific case can also refer to the outstanding situation where both NBBEX and NABEX tend to zero. This should mean that the intermediary chose exactly the right moment to trade and that his market timing was perfect. In this last specific situation, the zero directional EBEX is accompanied by a very high absolute EBEX score.

When investors give their intermediaries discretion over how they implement their trades, they like to assess whether their trades are well timed. As we have just seen, the directional EBEX offers a straightforward and easy-to-interpret measure of the quality of the market timing that varies from -1 (too aggressive) to 1 (too slow) and makes it possible to gauge the intermediary's use of the discretion he has been granted. This point is relevant as pre-trade analysis should help intermediaries predict likely market conditions and define an appropriate trading strategy. However, to be really relevant, the directional EBEX assumes that the intermediary has full discretion about how he fills the order within a given time window and that he does not systematically execute the order at the end of the specified period. The first assumption is realistic since many brokers are often given the market close as a deadline. The second assumption depends more on the intermediary's commitment to best execution.

Much as with absolute EBEX, we may analyse the distribution of directional EBEX of an

intermediary to assess the consistency his market timing. Again, comparative analysis across several intermediaries may be done based on directional EBEX distributions wherein the score of each intermediary is a risk-adjusted performance computed as the mean value of EBEX divided by its standard deviation. An active trader will be expected to provide an average directional EBEX as close to zero as possible with the lowest standard deviation.

(3) \$EBEX

In addition to the two fundamental indicators that measure the quality of execution and the quality of market timing, we have developed a third indicator, known as \$EBEX, whose purpose is to indicate the cost to the investor of the intermediary's trading performance. In short, this indicator likens the absolute EBEX scores reached by intermediaries to an opportunity cost for their clients. The \$EBEX is computed trade-by-trade for a given absolute EBEX that is defined arbitrarily or, when possible, corresponds to the performance objective set by the investor.

Take, for example, an EBEX target of 0.75 specified by the investor. To determine the total cost implied by the quality of execution achieved by his intermediary, we focus only on the trades with an EBEX lower than 0.75 and we compute trade-by-trade the signed difference between the actual execution price and the price that would have yielded an EBEX of at least 0.75. We then multiply each price difference by its corresponding trade size and add up all the weighted price differences to get the \$EBEX indicator. To facilitate interpretation and allow meaningful comparisons, we express the \$EBEX in % of the total monetary

VI. A New Framework: the EBEX Indicators

volume executed by the intermediary. The \$EBEX can be viewed as an opportunity cost since it represents the loss of failing to achieve particular trading performance.

Suppose, for example, a broker who has executed 100 orders for a single investor with a total traded volume of €49,827,671. His distribution of absolute EBEX exhibits a median score of 0.50 while he was set a target of 0.75. About eighty of the 100 trades he completed missed the target. Based on these trades, the \$EBEX is equal to €104,586 and corresponds to about 0.21% of the total volume traded. Now suppose another broker who has also executed about 100 orders for the same investor with a slightly lower total volume of €43,405,588. He too has a median score of 0.50 for the absolute EBEX but his \$EBEX amounts to €182,546, about 0.42% of the total volume he executed. Using \$EBEX to compare the two brokers helps the investor to identify that, though the quality of their execution is similar, the second broker's underperformance is more expensive.

(4) Potential adjustments

When intermediaries have discretion over how they execute orders within a specified horizon, EBEX indicators based on all market activity traded in that horizon reflect their ability to work orders and seize the best available trading opportunities. However, when intermediaries are given limited discretion, the EBEX indicators computed over the entire trading horizon may lead to noisy performance measures, especially when the constraints set by the investor are not neutral. This is true especially when they force the intermediary to execute at times of the least or most favourable prices.

Constraints are generally dependent on either price or volume. Most correspond to predetermined core trading strategies that are applied manually or, more and more, automatically, as a result of algorithmic trading that has become the must-have in the last few years and delegates the scheduling and execution of an order to a computer programme.

We can categorise price constraints as explicit and implicit strategies. Explicit strategies promote trading at a given price or better. The most frequent targets are the closing price of the day (market-on-close orders) or the bid-ask midpoint at the release time (arrival price orders⁴¹) but the target may also be a price directly specified by the investor (€25, for example). Implicit strategies are mainly VWAP strategies and customised strategies. A VWAP strategy schedules the trade according to the historical average volume profile of the security: it breaks up the order into smaller lots and trades them at every x minutes depending on the volume that has traded historically in that interval. A customised strategy gives a trader who is not benchmarked the ability to stop or cancel the execution of an order during the day or the trading horizon. For example, trading in a particular security may be stopped once a particular price or percentage of traded volume is reached.

Volume constraints correspond to participation strategies that schedule the trade by market volumes. These strategies result in executing the order as a specified constant percentage of the actual market volume, regardless of price. The key feature here is that the duration of the trade is determined not by a timeframe, but by the market conditions of the security.

41 - Some algorithm providers refer to them as implementation shortfall orders.

VI. A New Framework: the EBEX Indicators

When intermediaries are given explicit price constraints, they have a specific target to beat but they still have discretion over how they will work the order to reach their objective. Since they do not have any restrictive timing constraints within the trading horizon, EBEX indicators based on all market activity traded in that horizon reflect their ability to time orders and seize the best available trading opportunities. By contrast, when intermediaries are given implicit price constraints or volume constraints, they have limited discretion over how they can work the order to reach the investor's objective. In fact, both implicit price constraints and volume constraints impose indirect timing constraints on intermediaries, who are then forced to schedule in a predetermined way the trade by the historical average volume profile of the security. In such situations, we need to avoid noise and distinguish the trading performance of the intermediary and the instructions given by the investor. For this purpose, EBEX indicators may be adjusted on both the specified strategy and the expected market conditions over the trading horizon as follows:

$$\overline{EBEX}_{abs,i} = \sum_{j=1}^T \frac{x_{i,j}}{X_i} EBEX_{abs,ij}$$

where j is the time interval, T is the total number of intervals within the entire trading horizon, $x_{i,j}$ is the quantity to execute in interval j according to the specified strategy for order i , X_i is the total size of order i and $EBEX_{abs,ij}$ is the absolute EBEX indicator of the lot $x_{i,j}$ computed over interval j . $\overline{EBEX}_{abs,i}$ is then the weighted absolute EBEX for order i that attempts to measure the global quality of execution achieved for this order while accounting for the times that trades are requested of the intermediary.

The example shown in figure 18, a VWAP strategy for an order of 120,000 shares, illustrates this approach. If eight intervals are considered over the trading horizon, the schedule of execution is predetermined to participate proportionally in the total volume over the horizon. The order is broken up into eight lots that depend on expected market conditions. Once the entire order is executed, the absolute EBEX indicator for each lot can be computed by comparing its average execution price and the actual market activity within the corresponding trading period. Next, the quality of execution for the entire trade is obtained by computing the weighted average of all the absolute EBEX scores.

Figure 18: Illustration of the weighted absolute EBEX for a VWAP strategy

Trading period	Expected volume		VWAP strategy	
	Buy	Sell	Executed shares	EBEX score
1	200,000	200,000	24,000	0.52
2	150,000	150,000	18,000	0.63
3	100,000	100,000	12,000	0.71
4	50,000	50,000	6,000	0.44
5	50,000	50,000	6,000	0.54
6	100,000	100,000	12,000	0.52
7	150,000	150,000	18,000	0.72
8	200,000	200,000	24,000	0.27
Weighted EBEX				0.53

3. Illustration of the Framework

To illustrate both the use and the interpretation of our framework, we present hereafter the findings that we obtained for a relevant sample of orders provided by a European investment firm. Before presenting and interpreting the results, we will describe the data that we used as well as the assumptions that we made to calculate our indicators.

VI. A New Framework: the EBEX Indicators

A brief description of the order sample is also included.

(1) Data and assumptions

Our sample consists of 4,542 Euronext stock orders over the three-month period from January to March 2005. These orders were fully executed, so each order may be viewed as one trade. We should point out that we have no information regarding the potential constraints attached to these orders. This information was unavailable when we collected these order data.

The order database contains the following information for each order/trade:

- ISIN code for identifying the stock
- trade direction—buy or sell
- traded volume and its average execution price
- time-stamp for the time the order is sent to the broker—release time
- time-stamp for the time of broker response with execution details—response time
- identification code for the broker who completed the trade—broker ID

To apply our method to this order sample, we used the Euronext public trade data over the same period. This data set contains the following information for each trade executed in the central order book:

- Euronext internal code for the stock
- time-stamp for the time the trade is executed—execution time
- price and the number of shares traded

Given the data at hand, we had to formulate four assumptions to compute our indicators.

1. No transmission delay: we assume that the time the order is sent to the broker is the time the broker receives the order and can

start working it. Similarly, we assume that the time the response is sent by the broker is the execution time of the order.

2. No consideration for split executions: we assume that the order is fully executed at the time the broker answers, without considering the execution time of multiple trades, if any, to entirely fill the order. We assume that the time-stamp response corresponds to the time the broker stopped timing the market and working the order.

3. No specific instruction set by the investor: we assume that the broker has full discretion over the timing of the trade to complete. As a result, we always compute our indicators using as a last point of reference the market close of the day of execution.

4. The Euronext order book is the market of reference: as explained above, we use only Euronext public intraday trade data to compute our indicators. Our peer group analysis does not take into consideration the entire universe of trades since we lacked information about trades executed off the Euronext central order book (block trades executed on the upstairs market and OTC trades).

(2) Analysis of the order sample


The sample contains orders that exhibit various characteristics in terms of direction, size, release time and execution time.

Figure 19 shows the distribution of buy and sell orders over the three-month period. January, with about 39% of orders, is the busiest month. By contrast, only about 27% of orders are released in March. 51.47% of sample orders are for sells, 48.54% for buys—a good balance. Although it does not appear here, we can add that a great majority of the orders are executed the

VI. A New Framework: the EBEX Indicators

day of their release. Only about 14% are executed later.

Figure 19: Distribution of buy (B) and sell (S) orders by month



To document order release and execution times, we divide each trading session into one-hour intervals. We show the distribution of orders by release time in figure 20 and by execution time in figure 21. As the figures show, most orders are released in the morning with a peak between 10:00 and 12:00. The bulk of the orders are executed in the 17:00 interval that includes the market close. At the time (2005), the market closed at 17:40 at the latest on Euronext and it could be surprising that some orders are executed after the close. This might be the result of off-market trades but it is more likely that they are the result of our assumption that the broker response time is the execution time.

Figure 20: Distribution of orders by release time





Figure 21: Distribution of orders by execution time




The distribution of orders by size is shown in figure 22. For a meaningful comparison of stocks, we categorise the size of each order with respect to the daily average volume (DAV) of the stock. We define DAV as the total traded volume in EUR in a day divided by the total number of trades completed on that day. We use the DAV calculated for the stock and the day to identify the following five size categories:]0; 0.5] for orders whose size does not exceed half the DAV;]0.5; 1] for orders whose size exceeds half the DAV but is still inferior to the DAV;]1; 5] for

VI. A New Framework: the EBEX Indicators

orders whose size lies between one and five DAV;]5;10] for orders whose size lies between five and ten DAV; and finally]10;] for orders whose size is larger than ten DAV. Figure 22 shows that the sample contains orders of every size, with a predominance of large and medium-sized orders.

Figure 22: Distribution of orders by size



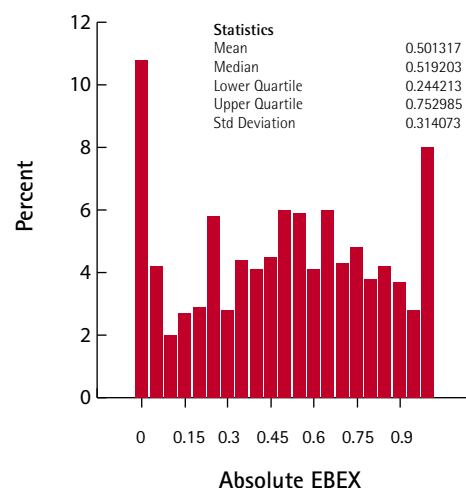
(3) Results and interpretation

Now that we have described the sample of orders, we can present the findings we obtain for our EBEX indicators. The objective here is to document both the transparency of the framework and the level of interpretation. We begin with an analysis of the overall quality of execution. In this case, orders can be aggregated at any level (direction, size, trend, etc.) and the interpretation of the absolute EBEX indicator is the most relevant. We then show the analysis done broker by broker. Here, interpreting the absolute EBEX with the directional EBEX as well as the \$EBEX is relevant.

(a) Overall performance

Figure 23 shows the empirical distribution of the absolute EBEX indicators computed for all the orders in our sample as well as some of the usual descriptive statistics. As the figure shows, the entire sample exhibits an average absolute EBEX of about 0.50, with a standard deviation of about 0.31. The median is about 0.52 and suggests relatively average overall execution since 50% of the orders have an absolute EBEX lower than 0.52.

Figure 23: Distribution of absolute EBEX for the entire order sample



Differences between buy (B) and sell (S) orders are shown in figure 24. As we may observe, the quality of execution looks slightly better for buy orders. Both the mean (0.52) and the median (0.54) scores are higher for these orders while standard deviation (0.31 instead of 0.32) is lower. Consistent with this result, the peak close to zero is also lower for buy orders: less than 10% against almost 13% for sell orders.

VI. A New Framework: the EBEX Indicators

Figure 24: Distribution of absolute EBEX by trade direction

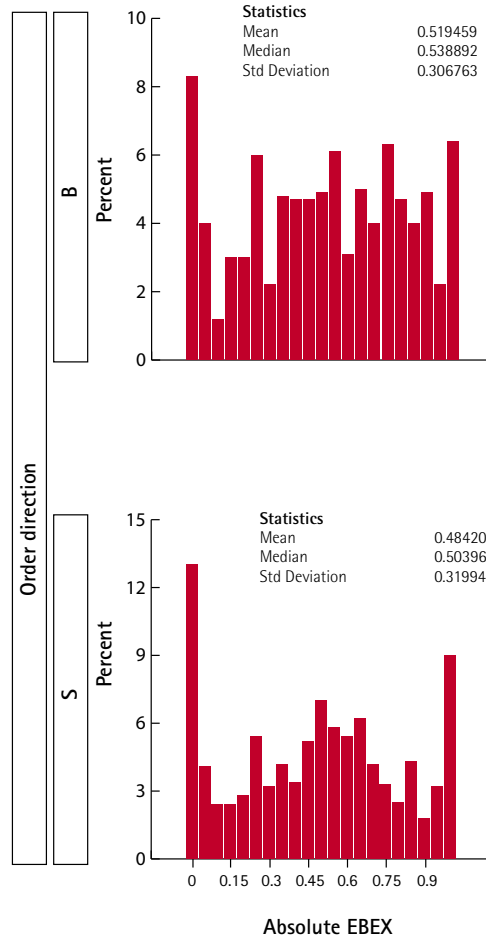


Figure 25: Distribution of absolute EBEX by execution date

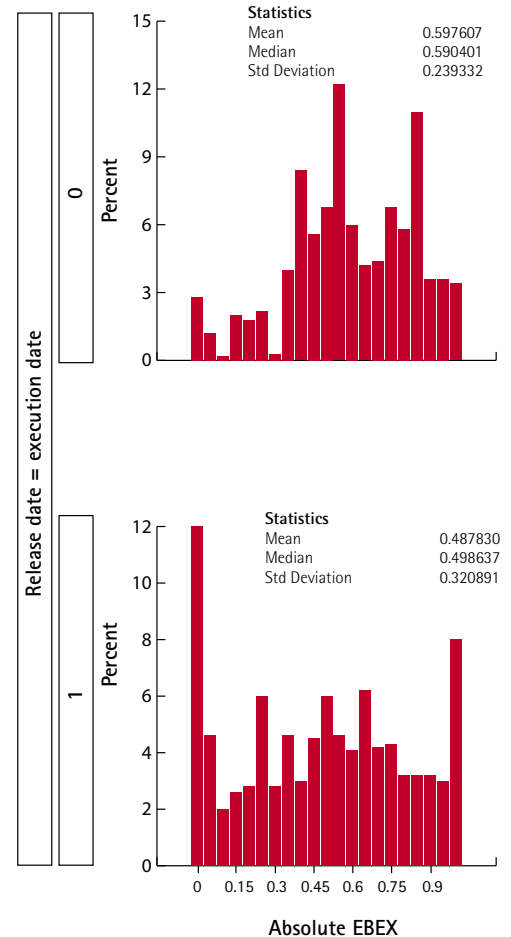



Figure 25 shows differences between the orders that are fully executed the day of their release (1) and those that are filled later (0). As the figure shows, the quality of execution appears overall to be better for the orders that are not entirely executed the day of their release. Both the mean and the median (0.59) scores are higher for these orders and standard deviation is lower.

Figure 26 exhibits the empirical distribution of absolute EBEX by the order size categories described earlier. No large differences appear at first sight but, when focusing on statistics (not shown in figure 26), both extreme order size categories exhibit the lowest average values for the absolute EBEX (0.47 and 0.48 respectively).

VI. A New Framework: the EBEX Indicators

Figure 26: Distribution of absolute EBEX by order size



(b) Broker performance

In our sample, more than ten different brokers are identified but three of them deal with more than 50% of the orders. For convenience, we report here the results for these three brokers only. In order to preserve anonymity, we will name them broker 1, broker 2 and broker 3. Figure 27 exhibits for each of them the number of trades completed as well as information about their size.

Figure 27: Distribution of orders by broker

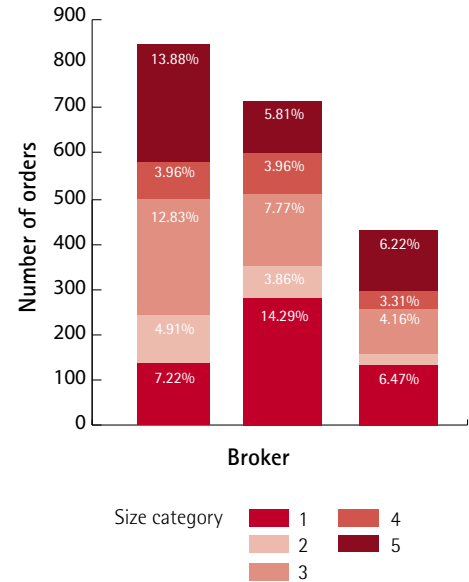
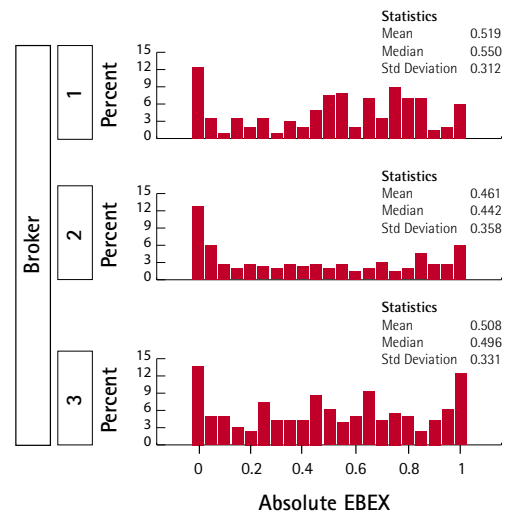


Figure 28 presents the empirical distribution of absolute EBEX by broker. Based on the median EBEX, we can directly conclude that broker 1 offers a higher quality of execution than broker 2 or broker 3. Broker 1 also appears to be more consistent than the two others (standard deviation of 0.31 instead of 0.35 or 0.33).

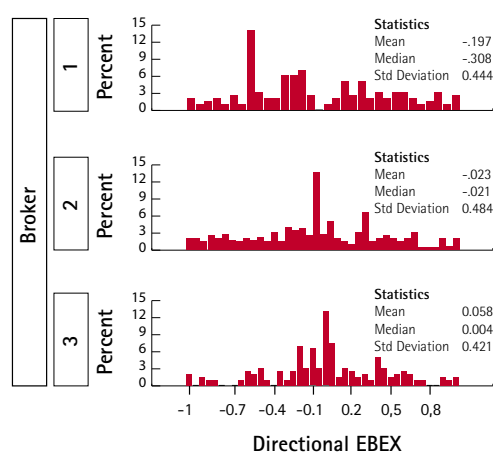
Figure 28: Distribution of absolute EBEX by broker



VI. A New Framework: the EBEX Indicators

In addition to absolute EBEX, we may look at the distribution of directional EBEX in figure 29 to analyse the brokers' market timing. For example, broker 1 exhibits a negative median directional EBEX meaning that he tends to be too aggressive. So, this broker could have reached a higher trading performance if he were more patient when executing orders.

Figure 29: Distribution of directional EBEX by broker



42 - Almost all data vendors as well as some trading venues compute and report in real time the daily VWAP.
43 - It indicates whether the trader received a higher or lower price than did the "average trader" of the day.
44- This stock was selected because it presents the largest number of trades to examine over the sample period (309 trades exactly).


(c) EBEX vs. VWAP

The common practice in the industry has so far been to use the daily VWAP to assess the trading performance, mainly because of its availability⁴² and easy interpretation.⁴³ However, the absolute EBEX indicator is superior to any VWAP score since it provides a meaningful metric that is consistent over time and across securities. We can illustrate this valuable advantage with some examples extracted from our empirical analysis for a given stock.⁴⁴

In figure 30, we plot both the absolute EBEX and the VWAP scores that we computed for all the trades executed on the selected stock. The VWAP score was first obtained by finding the signed difference between the execution price and the daily VWAP. Then

we divided this difference by the daily VWAP and multiplied the result by 10^4 to express the score in basis points. According to the benchmark comparison approach, a positive (negative) value indicates an execution that is less (more) favourable than the daily VWAP. Figure 30 is very interesting since it reveals quite dramatically the inconsistency of VWAP scores.

Figure 30: Absolute EBEX vs. VWAP score



Ignoring extreme performances (bad or excellent), VWAP scores are globally centred on zero and take either positive or negative values whatever the corresponding absolute EBEX. For example, the absolute EBEX is close to 0.7 when only about 30% of all market activity was traded at better prices than the execution price of the trade. In such a case, the interpretation is clearly an above-average quality of execution. In figure 30, for this level of EBEX, we observe VWAP scores varying from -50bp to +40bp. When focusing on extreme performances, we identify the same phenomenon. For very bad performances, the absolute EBEX equals zero, while VWAP scores vary from about +95bp to -20bp. For excellent performances, the absolute EBEX reaches 1 while VWAP scores exhibit values ranging from about -90bp to +20bp. Depending on the prevailing market conditions as well as


VI. A New Framework: the EBEX Indicators

price volatility, VWAP scores do not make comparison and proper assessment easy at all. By contrast, the absolute EBEX is shown to be a consistent metric that takes into account contemporaneous market activity and volatility, allowing a direct assessment of execution quality in the prevailing circumstances.

In addition to the VWAP score inconsistency, figure 30 indirectly reveals that the daily VWAP should not be viewed as a "fair" price for any order. By nature, the daily VWAP is the average of all executions in a day, including both the opening and closing prices that represent a large portion of market activity. As shown by figure 31, for the selected stock, the daily VWAP is therefore often very close to those prices, which are not always appropriate references for trades executed during the continuous session.

45 - This window goes from the order release time to the market close of the day.


Figure 31: Daily VWAP vs. closing/opening price



One way of cutting the bias mentioned above is to compute a VWAP over a shorter interval than the day. In figure 32, we plot both the absolute EBEX and the TVWAP scores, where the TVWAP is calculated over the same measurement window as the EBEX.⁴⁵ Although slightly reduced, TVWAP scores still exhibit inconsistency because they indicate the deviation from an average

but do not deliver any information about all other market activity dispersion. This prevents direct assessments of the quality of execution.

Figure 32: Absolute EBEX vs. TVWAP score



Conclusion



Conclusion

In this guide to transaction cost analysis, we have attempted to cover the most meaningful tools and techniques available to investors, managers and intermediaries to properly assess the quality of execution and possibly provide an answer to the difficult challenges posed by article 21 of the MiFID.

Our Transaction Cost Analysis A-Z should be viewed only as a starting point and readers should actively refer to the publications and the authors mentioned in the pages of our guide. The literature on transaction cost analysis is not voluminous, but there is some, and it covers most aspects of the question. This guide therefore only provides an overview while attempting to clarify the various concepts involved.

The industry has begun embracing the tools and techniques described here and we are confident that current market developments and industry changes will fuel the need for more advanced techniques. In particular, we strongly believe the better formalisation of both execution performances and risk will open the door to an unlimited area for developing systematic trading algorithms whose aim is to achieve optimal execution plans; innovative value propositions offered by brokers, exchanges and obviously technology vendors cannot be far off.

Finally, we strongly believe that the role of back-office providers in the MiFID Best Execution chain is key, as they remain in a unique position to capture and process transaction data independently from venues and intermediaries on behalf of the investors or their representatives. It is only a matter of time before transaction

cost attribution is provided as a service, as portfolio performance already is. Current market conditions and the need to extract every single basis point of outperformance will contribute to this evolution. The EDHEC Risk and Asset Management Research Centre will remain involved in this evolution in the coming years and support the publication of research results obtained not only in the academic world but also by practitioners.

References



References

- Almgren, R., and N. Chriss. 2000. Optimal execution of portfolio transactions. *Journal of Risk* 3(2): 5–39.
- Boussema, M., A. Bueno and P. Sequier. 2001. Transaction costs and trading strategies: an empirical analysis on global equity markets. Working paper.
- . 2002. Transaction costs and trading strategies. In *Best execution: executing transactions in security markets on behalf of investors*. European Asset Management Association.
- Copeland, T., and D. Galai. 1983. Information effects and the bid-ask spread. *Journal of Finance* 38: 1457–69.
- Demsetz, H. 1968. The cost of trading. *Quarterly Journal of Economics* 82: 33–53.
- D'Hondt, C., and J.-R. Giraud. 2007. MiFID: the (in)famous European directive? EDHEC Position Paper (March).
- Domowitz, I., J. Glen and A. Madhavan. 2001. Global equity trading costs. Research paper, Investment Technology Group.
- Giraud, J.-R. 2004. Best execution for buy-side firms: a challenging Issue, a promising debate, a regulatory challenge. European Survey on Investment Managers' Practices (June).
- Harris, L. 2003. *Trading & exchanges: market microstructure for practitioners*. New York: Oxford University Press.
- Ho, T., and H. Stoll. 1981. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics* 9: 47–73.
- Karceski, J., M. Livingston, and E. O'Neal. 2004. Portfolio Transaction Costs at U.S. Equity Mutual Funds. ZAG commissioned study.
- Kissell, R., and M. Glantz (2003), *Optimal trading strategies*." New York, AMACOM.
- Munck, N. H. 2005. When transactions went high-tech: a cross-sectional study of equity trading costs in the light of more sophisticated trading systems. SSRN working paper.
- Nevmyvaka, Y., M. Kearns, A. Papandreou and K. Sycara. 2005. Electronic trading in order-driven markets: efficient execution. Working paper.
- Perold, A. F. 1988. The implementation shortfall: paper versus reality. *Journal of Portfolio Management* 14: 4–9.

References

About the EDHEC Risk and Asset Management Research Centre


EDHEC is one of the top five business schools in France. Its reputation is built on the high quality of its faculty (110 professors and researchers from France and abroad) and the privileged relationship with professionals that the school has been developing since its establishment in 1906. EDHEC Business School has decided to draw on its extensive knowledge of the professional environment and has therefore focused its research on themes that satisfy the needs of professionals. EDHEC is also one of the few business schools in Europe to have received the triple international accreditation: AACSB (US-Global), EQUIS (Europe-Global) and Association of MBAs (UK-Global).

EDHEC pursues an active research policy in the field of finance. The EDHEC Risk and Asset Management Research Centre carries out numerous research programmes in the areas of asset allocation and risk management in both the traditional and alternative investment universes.

The choice of asset allocation

The EDHEC Risk and Asset Management Research Centre structures all of its research work around asset allocation. This issue corresponds to a genuine expectation from the market. On the one hand, the prevailing stock market situation in recent years has shown the limitations of active management based solely on stock picking as a source of performance.

Percentage of variation between funds



Source: EDHEC (2002) and Ibbotson, Kaplan (2000)

On the other, the appearance of new asset classes (hedge funds, private equity), with risk profiles that are very different from those of the traditional investment universe, constitutes a new opportunity in both conceptual and operational terms. This strategic choice is applied to all of the Centre's research programmes, whether they involve proposing new methods of strategic allocation, which integrate the alternative class; measuring the performance of funds while taking the tactical allocation dimension of the alpha into account; taking extreme risks into account in the allocation; or studying the usefulness of derivatives in constructing the portfolio.

An applied research approach

In an attempt to ensure that the research it carries out is truly applicable, EDHEC has implemented a dual validation system for the work of the EDHEC Risk and Asset Management

Research Centre. All research work must be part of a research programme, the relevance and goals of which have been validated from both an academic and a business viewpoint by the Centre's advisory board. This board is made up of both internationally recognised researchers and the Centre's business partners. The management of the research programmes respects a rigorous validation process, which guarantees the scientific quality and the operational usefulness of the programmes.

To date, the Centre has implemented six research programmes:

Asset Allocation and Alternative Diversification

Sponsored by SG Asset Management and Newedge

The research carried out focuses on the benefits, risks and integration methods of the alternative class in asset allocation. From that perspective, EDHEC is making a significant contribution to the research conducted in the area of multi-style/multi-class portfolio construction.

Performance and Style Analysis

Part of a business partnership with EuroPerformance

The scientific goal of the research is to adapt the portfolio performance and style analysis models and methods to tactical allocation. The results of the research carried out by EDHEC thereby allow portfolio alpha to be measured not only for stock picking but also for style timing.

Indices and Benchmarking

Sponsored by Af2i, Barclays Global Investors, BNP Paribas Investment Partners, NYSE Euronext, Lyxor Asset Management, and UBS Global Asset Management

This research programme has given rise to extensive research on the subject of indices and benchmarks in both the hedge fund universe and more traditional investment

About the EDHEC Risk and Asset Management Research Centre

classes. Its main focus is on analysing the quality of indices and the criteria for choosing indices for institutional investors. EDHEC also proposes an original proprietary style index construction methodology for both the traditional and alternative universes. These indices are intended to be a response to the critiques relating to the lack of representativeness of the style indices that are available on the market. In 2003, EDHEC launched the first composite hedge fund strategy indices.

Asset Allocation and Derivatives

Sponsored by Eurex, SGCI and the French Banking Federation

This research programme focuses on the usefulness of employing derivative instruments in the area of portfolio construction, whether it involves implementing active portfolio allocation or replicating indices. "Passive" replication of "active" hedge fund indices through portfolios of derivative instruments is a key area in the research carried out by EDHEC. This programme includes the "Structured Products and Derivatives Instruments" research chair sponsored by the French Banking Federation.

Best Execution and Operational Performance

Sponsored by CACEIS, NYSE Euronext, and SunGard

This research programme deals with two topics: best execution and, more generally, the issue of operational risk. The goal of the research programme is to develop a complete framework for measuring transaction costs: EBEX ("Estimated Best Execution") but also to develop the existing framework for specific situations (constrained orders, listed derivatives, etc.). Research also focuses on risk-adjusted performance measurement of execution strategies, analysis of market

impact and opportunity costs on listed derivatives order books, the impact of explicit and implicit transaction costs on portfolio performances, and the impact of market fragmentation resulting from MiFID on the quality of execution in European listed securities markets. This programme includes the "MiFID and Best Execution" research chair, sponsored by CACEIS, NYSE Euronext, and SunGard.

ALM and Asset Management

Sponsored by BNP Paribas Investment Partners, AXA Investment Managers and ORTEC Finance

This research programme concentrates on the application of recent research in the area of asset-liability management for pension plans and insurance companies. The research centre is working on the idea that improving asset management techniques and particularly strategic allocation techniques has a positive impact on the performance of asset-liability management programmes. The programme includes research on the benefits of alternative investments, such as hedge funds, in long-term portfolio management. Particular attention is given to the institutional context of ALM and notably the integration of the impact of the IFRS standards and the Solvency II directive project. It also aims to develop an ALM approach addressing the particular needs, constraints, and objectives of the private banking clientele. This programme includes the "Regulation and Institutional Investment" research chair, sponsored by AXA Investment Managers, the "Asset-Liability Management and Institutional Investment Management" research chair, sponsored by BNP Paribas Investment Partners and the "Private Asset-Liability Management" research chair, in partnership with ORTEC Finance.

About the EDHEC Risk and Asset Management Research Centre

Seven Research Chairs have been endowed:

Regulation and Institutional Investment

In partnership with AXA Investment Managers

The chair investigates the interaction between regulation and institutional investment management on a European scale and highlights the challenges of regulatory developments for institutional investment managers.

Asset-Liability Management and Institutional Investment Management

In partnership with BNP Paribas Investment Partners

The chair examines advanced asset-liability management topics such as dynamic allocation strategies, rational pricing of liability schemes, and formulation of an ALM model integrating the financial circumstances of pension plan sponsors.

MiFID and Best Execution

In partnership with NYSE Euronext, SunGard, and CACEIS Investor Services

The chair looks at two crucial issues linked to the Markets in Financial Instruments Directive: building a complete framework for transaction cost analysis and analysing the consequences of market fragmentation.

Structured Products and Derivative Instruments

Sponsored by the French Banking Federation (FBF)

The chair investigates the optimal design of structured products in an ALM context and studies structured products and derivatives on relatively illiquid underlying instruments.

Financial Engineering and Global Alternative Portfolios for Institutional Investors

Sponsored by Morgan Stanley Investment Management

The chair adapts risk budgeting and risk management concepts and techniques to the specificities of alternative investments, both in the context of asset management and asset-liability management.

Private Asset-Liability Management

In partnership with ORTEC Finance

The chair will focus on the benefits of the asset-liability management approach to private wealth management, with particular attention being given to the life cycle asset allocation topic.

Dynamic Allocation Models and New Forms of Target Funds for Private and Institutional Clients

In partnership with Groupe UFG

The chair consists of academic research that will be devoted to the analysis and improvement of dynamic allocation models and new forms of target funds.

The EDHEC PhD in Finance

The PhD in Finance at EDHEC Business School is designed for professionals who aspire to higher intellectual levels and aim to redefine the investment banking and asset management industries.

It is offered in two tracks: a residential track for high-potential graduate students who will hold part-time positions at EDHEC Business School, and an executive track for practitioners who will keep their full-time jobs.

Drawing its faculty from the world's best universities and enjoying the support of

About the EDHEC Risk and Asset Management Research Centre

the research centre with the most impact on the European financial industry, the EDHEC PhD in Finance creates an extraordinary platform for professional development and industry innovation.

Research for Business

To optimise exchanges between the academic and business worlds, the EDHEC Risk and Asset Management Research Centre maintains a website devoted to asset management research for the industry: (www.edhec-risk.com), circulates a monthly newsletter to over 200,000 practitioners,

conducts regular industry surveys and consultations, and organises annual conferences for the benefit of institutional investors and asset managers.

The Centre's activities have also given rise to the business offshoots **EDHEC Investment Research** and **EDHEC Asset Management Education**.

EDHEC Asset Management Education helps investment professionals to upgrade their skills with advanced risk and asset management training across traditional and alternative classes.

Industry surveys: comparing research advances with industry best practices

EDHEC regularly conducts surveys on the state of the European asset management industry. They look at the application of recent research advances within investment management companies and at best practices in the industry. Survey results receive considerable attention from professionals and are extensively reported by the international financial media.



Recent industry surveys conducted by the EDHEC Risk and Asset Management Research Centre

The EDHEC European ETF Survey 2008 sponsored by iShares

The EDHEC European Investment Practices Survey 2008 sponsored by Newedge

EDHEC European Real Estate Investment and Risk Management Survey 2007 sponsored by Aberdeen Property Investors and Groupe UFG

EuroPerformance-EDHEC Style Ratings and Alpha League Table

The business partnership between France's leading fund rating agency and the EDHEC Risk and Asset Management Research Centre led to the 2004 launch of the EuroPerformance-EDHEC Style Ratings, a free rating service for funds distributed in Europe which addresses market demand by delivering a true picture of alpha, accounting for potential extreme loss, and measuring performance persistence. The risk-adjusted performance of individual funds is used to build the Alpha League Table, the first ranking of European asset management companies based on their ability to deliver value on their equity management. www.stylerrating.com

EDHEC-Risk website

The EDHEC Risk and Asset Management Research Centre's website makes EDHEC's analyses and expertise in the field of asset management and ALM available to professionals. The site examines the latest academic research from a business perspective, and provides a critical look at the most recent industry news. www.edhec-risk.com

About CACEIS

CACEIS is a banking group dedicated to asset servicing. Equally owned by the groups Crédit Agricole and Caisse d'Epargne/Banque Populaire, CACEIS services institutional clients and asset managers worldwide. With €2.3 trillion under custody and €1.1 trillion under administration, we are France's premier investor services provider and a leading global player.

Through a network of offices across Europe, North America and soon Asia, CACEIS delivers high quality services covering depositary/trustee - custody, fund administration, transfer agency and corporate trust.

Considerable expertise in Alternative Investment servicing makes CACEIS one of the leading service providers for these sophisticated funds. We also offer a range of specialist services including cross-border fund distribution support, designed to assist clients in achieving their international business development goals.

CACEIS

1 place Valhubert

75013 Paris - France

Tel.: +33 (0)1 57 78 00 00

www.caceis.com



... and climbing.

A global player in asset servicing...

Offering leading value in investor services demands constant evolution. At CACEIS, our strategy of sustained growth is helping customers meet competitive challenges on a global scale. Find out how our highly adapted investor services can keep you a leap ahead.

CACEIS, your comprehensive securities servicing partner.

www.munster-bbn.com

Custody-Depositary / Trustee
Fund Administration
Corporate Trust
CACEIS benefits from an S&P AA- rating

www.caceis.com

caceis
INVESTOR SERVICES
solid & innovative

About NYSE Euronext

NYSE Euronext (NYX) operates the world's leading and most liquid exchange group, and seeks to provide the highest levels of quality, customer choice and innovation. Its family of exchanges, located in six countries, includes the New York Stock Exchange, the world's largest cash equities market; Euronext, the Eurozone's largest cash equities market; Liffe, Europe's leading derivatives exchange by value of trading; and NYSE Arca Options, one of the fastest growing U.S. options trading platforms.

NYSE Euronext offers a diverse array of financial products and services for issuers, investors and financial institutions in cash equities, options and derivatives, ETFs, bonds, market data, and commercial technology solutions. NYSE Euronext's nearly 4,000 listed companies represent a combined \$28.5 / €18.4 trillion (as of May 31, 2008), in total global market capitalization, more than four times that of any other exchange group. Liffe offers an unrivalled range of global futures and options products, which are made available to customers worldwide on its state-of-the-art trading platform, LIFFE CONNECT®.

NYSE Euronext
Cannon Bridge House - 1 Cousin Lane
London EC4R 3XX - United Kingdom
Tel: +44 (0) 207 623 0444

39 rue Cambon
75039 Paris Cedex 01 - France
Tel.: +33 (0)1 49 27 10 00

www.nyx.com



FACTS YOU CAN TRADE ON:

**WE SET THE EBBO
94% OF THE TIME,
WITH THE MOST
TRANSPARENCY.**

The most liquidity, the clearest choice. At NYSE Euronext, you'll find the largest, deepest pool of liquidity in the European market. As transparent as we are deep, NYSE Euronext allows investors equal access to information and instant visibility on the depth of the order book. And, with our TradeCheck product, you can analyze and optimize your trade executions. Our Liquidity Provider models ensure a permanent flow of liquidity and optimum trading conditions for all markets. Today, there's a new spirit of competitiveness within the European market. Look no further than NYSE Euronext, the place with the most efficient price, speed and transparency available.

nyx.com/equities

 **NYSE Euronext.**

This advertising is solely intended as information and does not constitute any investment advice or an offer, solicitation or recommendation to acquire or dispose of any investment or to engage in any transaction. Persons wishing to trade products available on NYSE Euronext markets or wishing to offer such products to third parties are advised, before doing so, to check their legal and regulatory position in the relevant territory and to understand the related risks. NYSE Euronext™, Euronext®, NYSE Arca Europe, SmartPool are registered marks of NYSE Euronext. ©2008, NYSE Euronext. All rights reserved. EBBO is the abbreviation for European Best Bid and Offer. Figures through Q3 2008 on Euronext central order book on all European securities.

About SunGard

With annual revenue of \$5 billion, SunGard is a global leader in software and processing solutions for financial services, higher education and the public sector. SunGard also helps information-dependent enterprises of all types to ensure the continuity of their business. SunGard serves more than 25,000 customers in more than 50 countries, including the world's 50 largest financial services companies.

Asset Arena is SunGard's global suite of solutions for front-to-back-office investment support, serving asset managers, third-party administrators, insurers and pension fund managers. Available as an integrated solution set or as individual offerings, Asset Arena covers the entire investment management life-cycle, including trading, compliance, investment accounting, portfolio management and performance measurement and attribution.

An integrated, front-to-back market data and historical database solution for storing and managing real-time and high-volume time series data, FAME is used by leading institutions in the financial, energy and public sectors, as well as third-party content aggregators, software vendors and individual investors. FAME provides investment professionals and intermediaries with a variety of real-time market data feeds, a Web-based desktop solution, application hosting, data delivery components, and tools for performing analytic modeling.

For more information visit:

www.sungard.com

Contact:


Elias Nechachby

elias.nechachby@sungard.com

Anders Nilsson

anders.nilsson@sungard.com

When you're satisfied, so are we.



SUNGARD ASSET ARENA

Investment
Management
Solutions

At SunGard we work hard to keep customer satisfaction at the forefront of our minds. With Asset Arena we provide investment management solutions for the global community. We combine innovative thinking with the technology to support smoother and more efficient workflows. To keep up with the ever changing complexity of the investment universe. To remain one step ahead of the competition.

More than \$18 trillion in assets are processed world-wide with Asset Arena.

To learn more please email assetarena@sungard.com.

www.sungard.com/assetarena

©2008 SunGard

Trademark information: SunGard, the SunGard logo, and Asset Arena are trademarks or registered trademarks of SunGard Data Systems Inc. or its subsidiaries in the U.S. and other countries. All other trade names are trademarks or registered trademarks of their respective holders.

EDHEC Risk and Asset Management
Research Centre
393-400 promenade des Anglais
BP 3116
06202 Nice Cedex 3 - France
Tel.: +33 (0)4 93 18 78 24
Fax: +33 (0)4 93 18 78 41
E-mail: research@edhec-risk.com
Web: www.edhec-risk.com