The use of quantitative models in execution analytics and algorithmic trading

Michael Simmonds
Equity Quantitative Analytics

16 February 2007

LEHMAN BROTHERS

Outline

- Projects of Interest in Equity Quantitative Analysis
- Estimating Transaction Costs
- Producing a Transaction Cost Aware Trading Strategy
- Conclusion

Equity Quantitative Analytics Group

- Group based on the trading floor.
- Not a "research" group so does not provide explicit buy/sell recommendations.
- Provides advice on executing trades, market microstructure, cost measurement and portfolio management, Provides important tools for the trading floor and clients such as algorithmic trading and risk liquidation facilitation.
- Principally populated by PhDs in Maths, Physics, Engineering and Economics.

Transaction Costs

A key element of the investment process

- Actual investment performance is the result of both
  - The investment alpha
  - The execution costs
- It is widely recognised that these execution costs can reduce substantially the notional return of the investment strategy
- As a result, good insights into the drivers of execution cost can prove hugely valuable to the investment process both at the portfolio construction and at the trading stages

Introduction

The Implementation Shortfall framework

- The Implementation Shortfall methodology was introduced by Perold (1988) to measure transaction costs
  - It defines the Implementation Shortfall as the difference between the actual portfolio return and its paper return
  - The Implementation Shortfall can be broken down into
    - Explicit costs: Commissions, custody fees, taxes
    - Implicit costs: Market impact, opportunity cost
- PRISE model = Pre-Trade Implementation Shortfall Estimation model
- PRISE focuses on implicit costs
- It produces market impact estimates with confidence intervals

The Data

Where it all began: The orders and executions database

- Over the last few years, Lehman Brothers has built one of the largest global algorithmic trading franchise
- In the process, we built a very substantial algorithmic trading orders and executions database
- Our impact modelling work leverage this dataset
  - Breadth: Global database
  - Size: 2 years + trade history; 8 M+ worked orders
  - Uniqueness: Database of worked orders with whole sequence of resulting child orders
  - Caveat: Lehman Brothers specific data; the model forecasts only the impact generated by and the skill of our executions

The Data

A carefully constructed data sample

- A carefully constructed data sample
  - VWAP and with Volume worked orders
  - Fully executed orders with no restrictions
- A range of explanatory factors considered: Spread, trade size, trade rate, volatility, order duration
- Carefully constructed ex-ante predictions of explanatory variables (for eg trade rate)
- Intraday profile of variables taken into account

The Data

A carefully constructed data sample

Intraday estimates can be quite different from the daily averages

- Eg Glaxo SmithKline

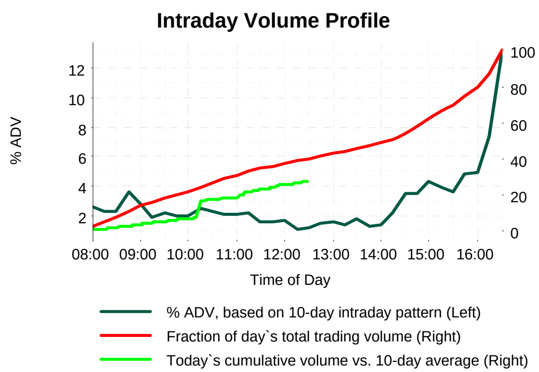

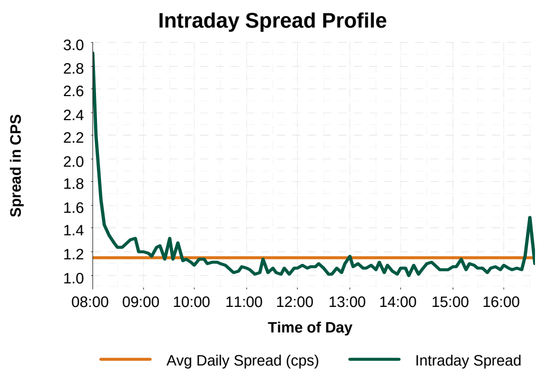

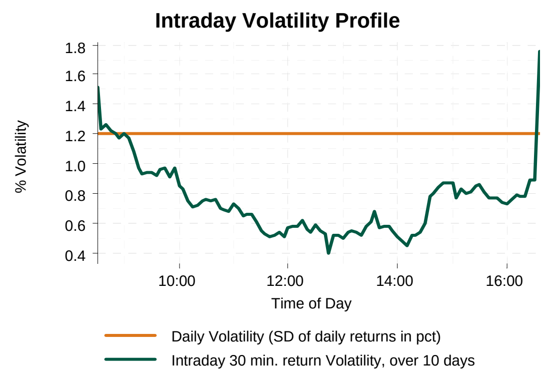

- Volume is low in the morning, increases after US open
- Spread is high on / after the morning auction and lower after 10am
- Volatility is high in the morning, lower at lunch time and higher again in the afternoon (with local peaks for US numbers and open)
- Up-to-date market and single stock profiles are available on TAPAS

The Data

Characteristics of the European data

- Full day or intraday trades
- 88% of the orders are VWAP orders; 12% with Volume orders
- The typical VWAP order is small: median value is $30k, median trade rate is 0.7%
- The typical with Volume order is larger: Median trade rate is 11%
- The vast majority of orders is in the 0.01–50% trade rate range
- The typical order is short: Median order duration is half hour
- Median VWAP impact is 2.2bps
- Median with Volume impact is 8bps

A Few Empirical Facts

Individual trade impacts are very noisy

Mean Impact = 3.9bps   Std Dev Impact = 21bps

- Impact distribution is highly leptokurtic

Impact Histogram

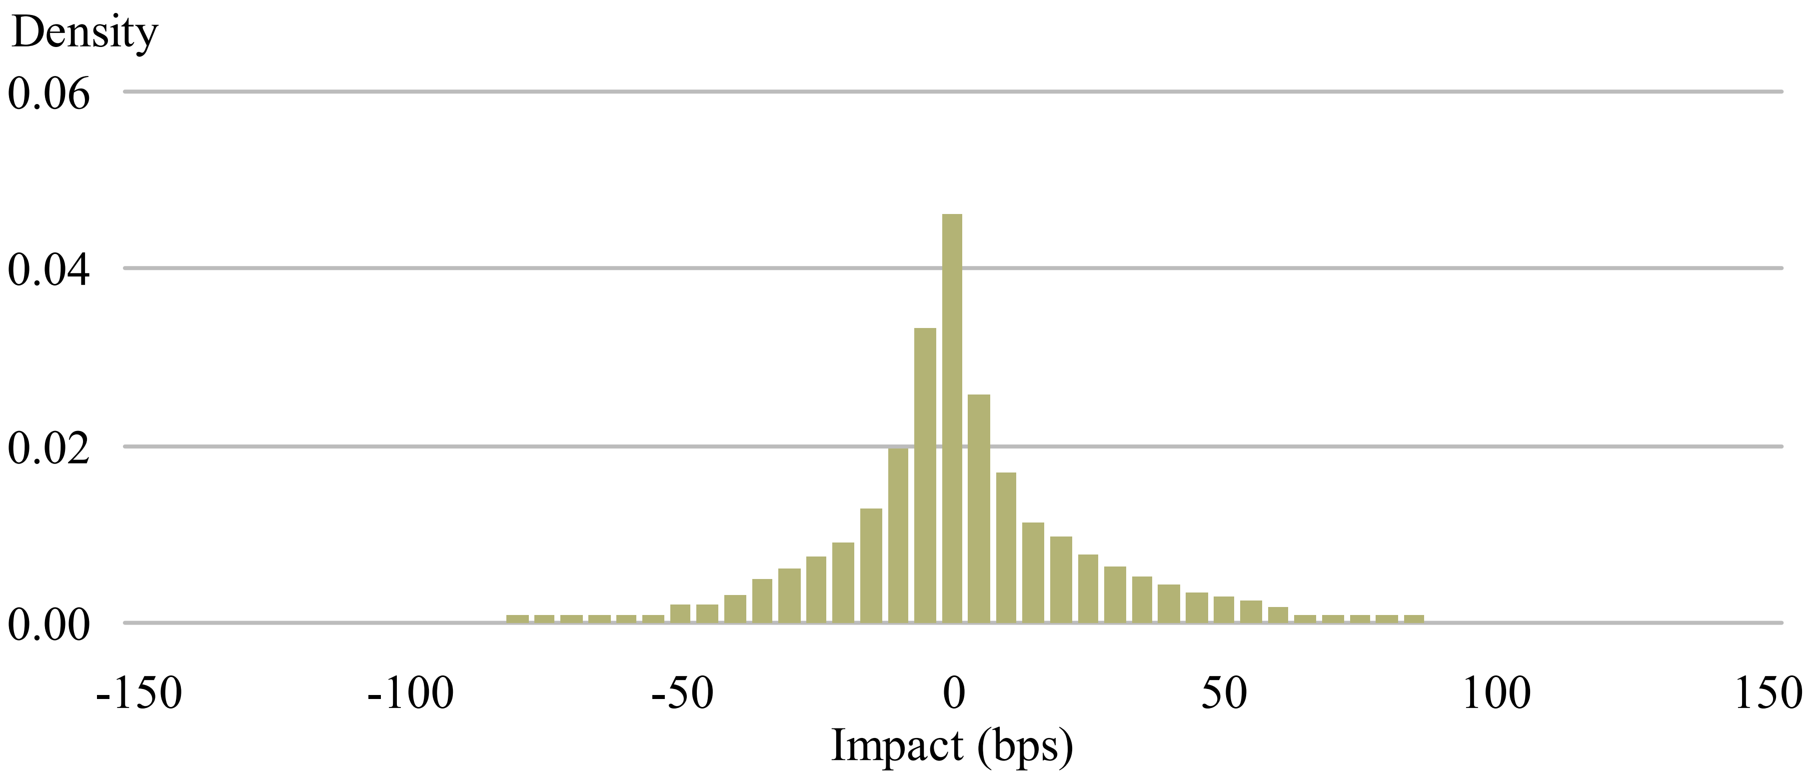

A Few Empirical Facts

Impact standard deviation increases with trade period volatility

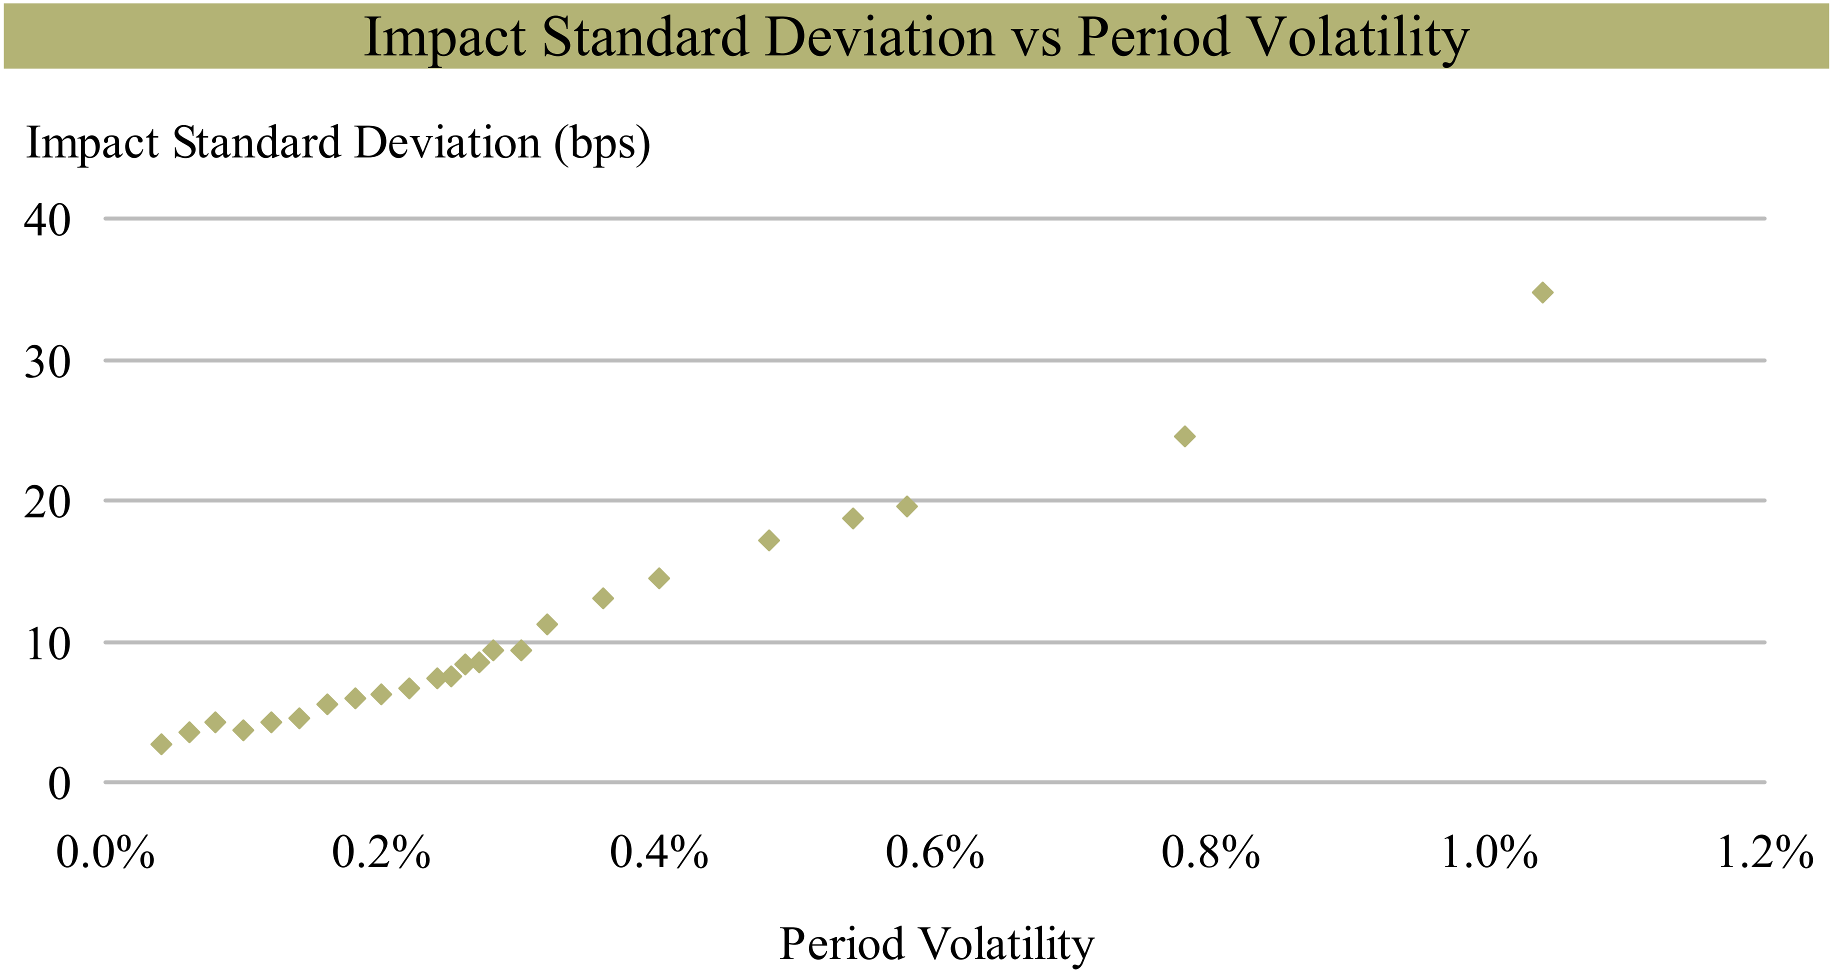

A Few Empirical Facts

Impact trajectory has a clear pattern

- Impact is high at the outset
- Marginal impact decreases with time
- Impact trajectory is reminiscent of a power law

Intra-Trade Impact Profile – 45mns Orders

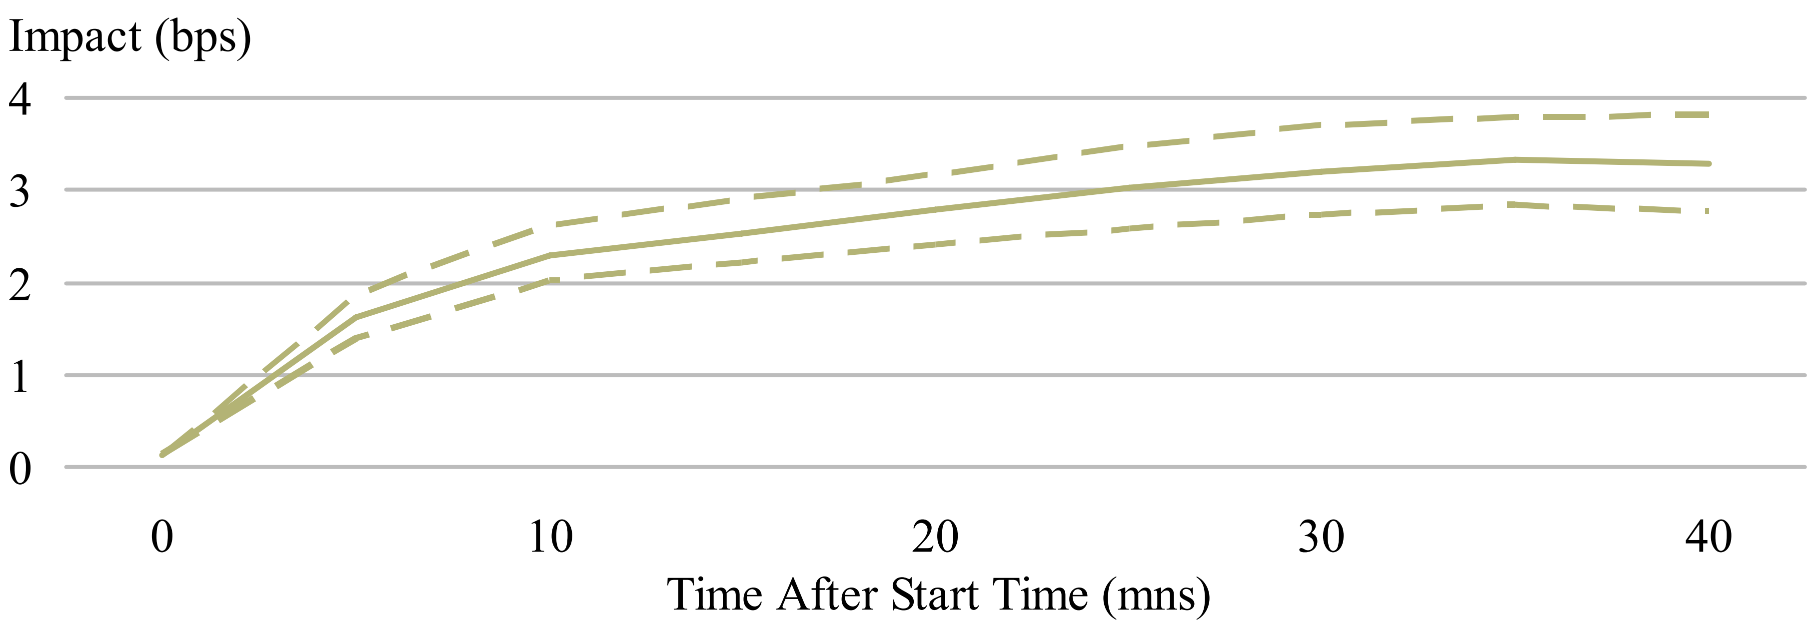

A Few Empirical Facts

Post-trade impact trajectory also has a clear pattern

- Impact reverts after the end of the trade
- It eventually settles to a level higher than 0
- This graph shows the existence for worked orders of a permanent and a temporary impact often mentioned in the block trade literature

Post-Trade Impact Profiles

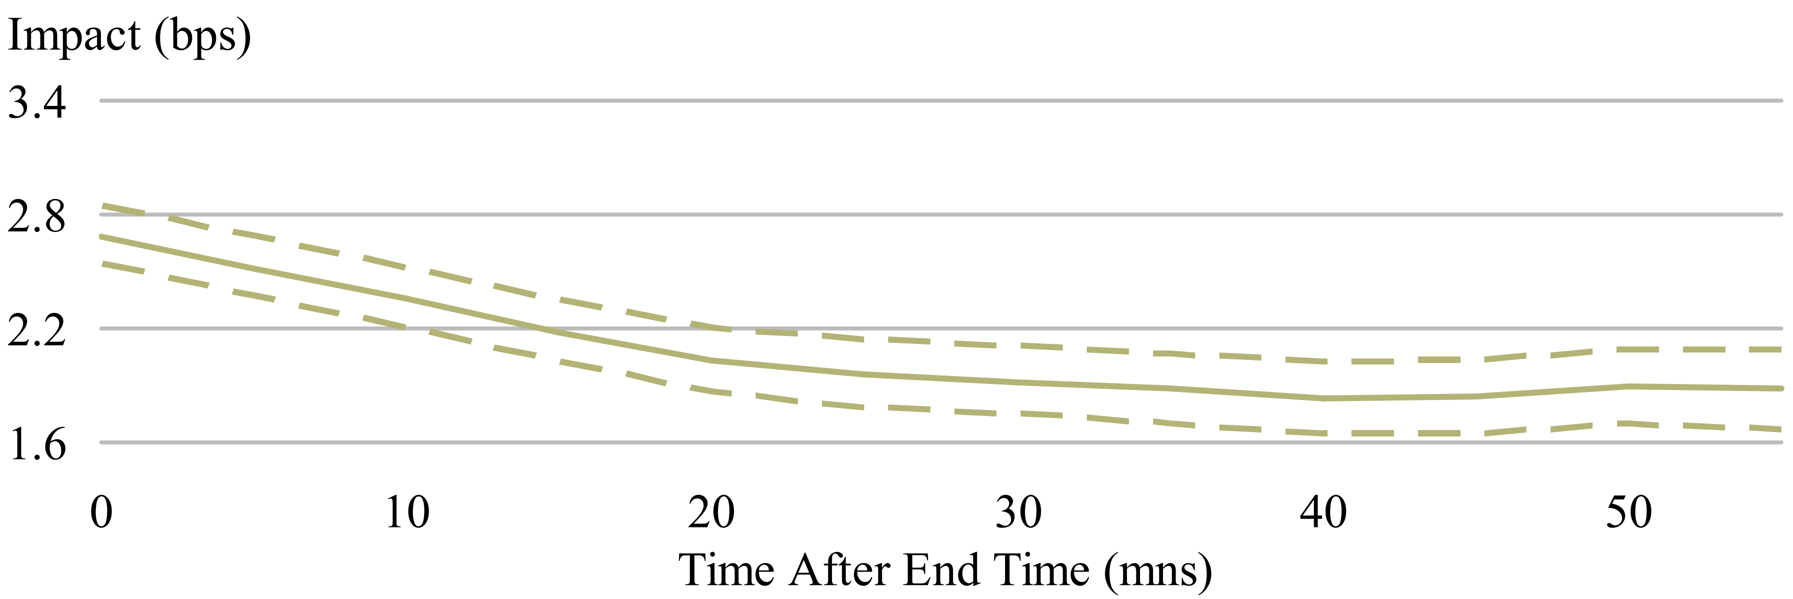

A Few Empirical Facts

Impact increases with trade rate; the relationship is noisy

Correlation = 19%

Impact vs Trade Rate – 1000 Bins

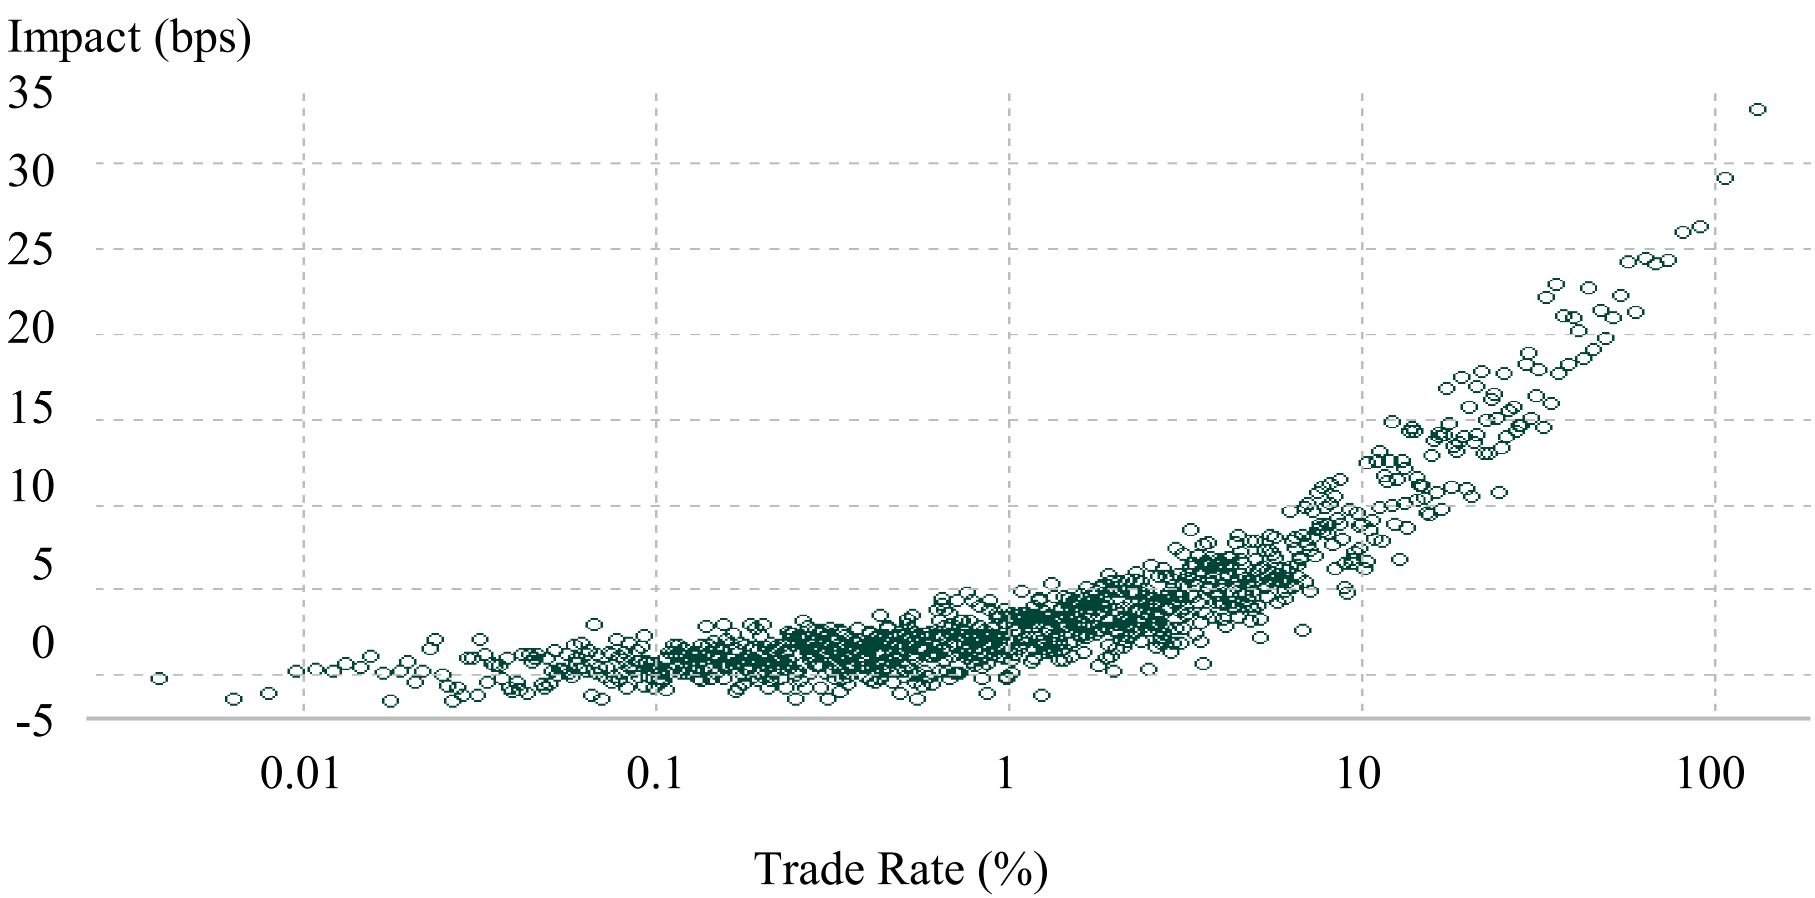

A Few Empirical Facts

…but extremely significant

p-value = 0

Impact vs Trade Rate – 100 Bins

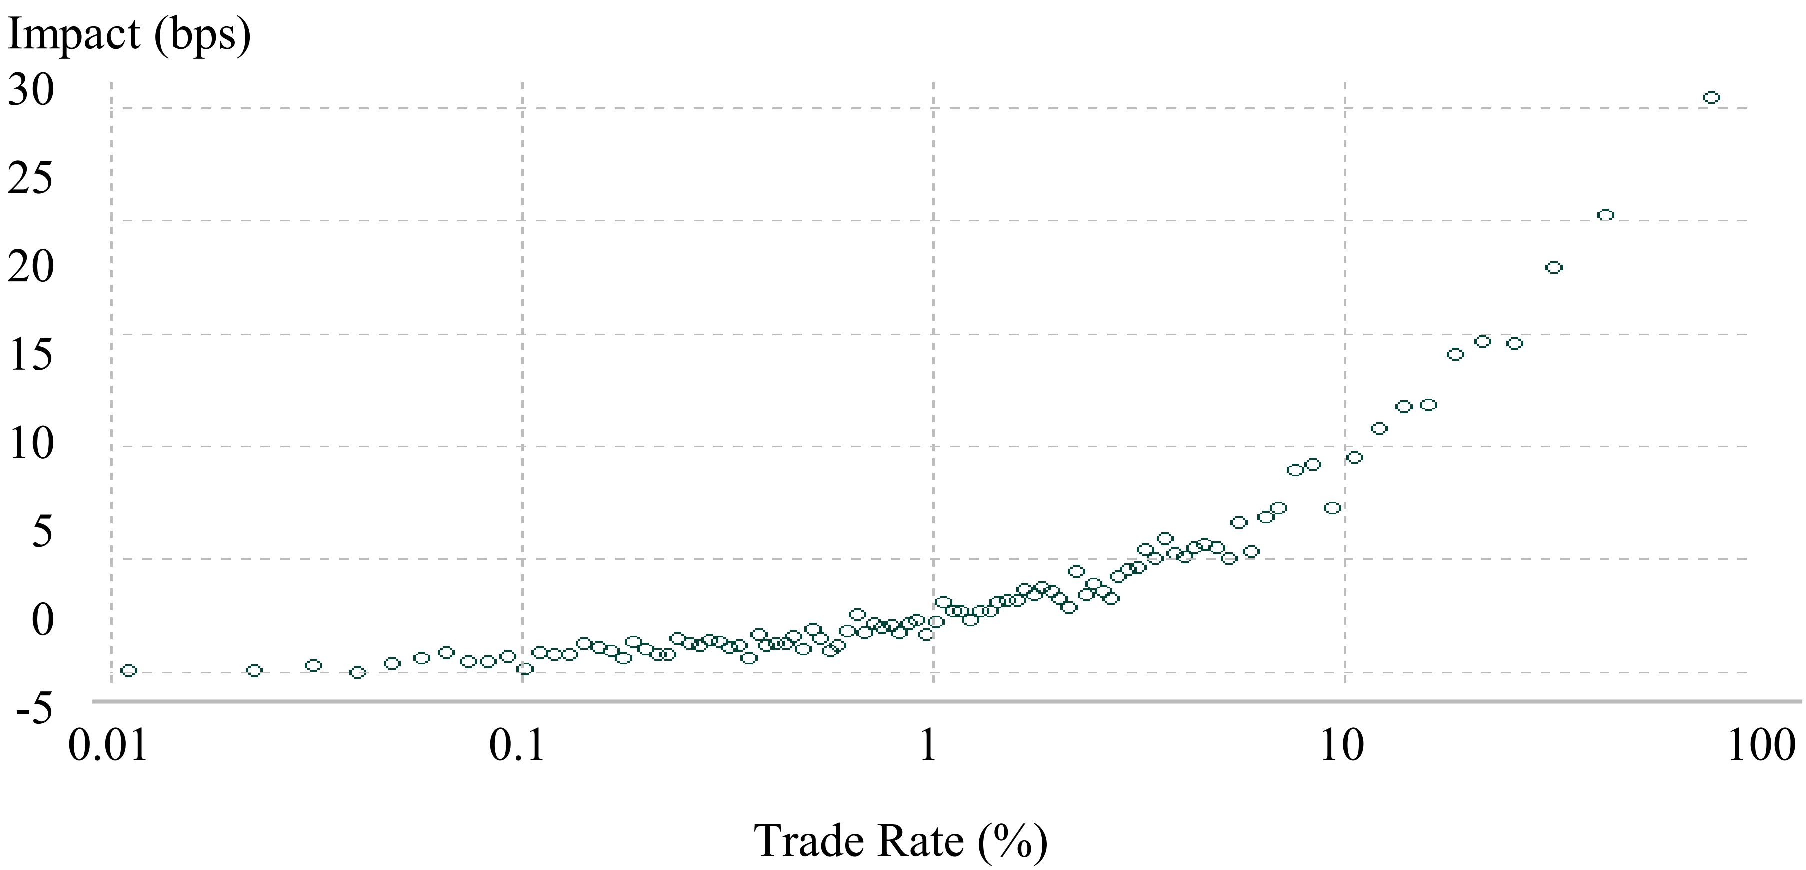

A Few Empirical Facts

Impact also has a clear positive relationship with period volatility

Correlation = 7.2% p-value = 0

Impact vs Period Volatility – 100 Bins

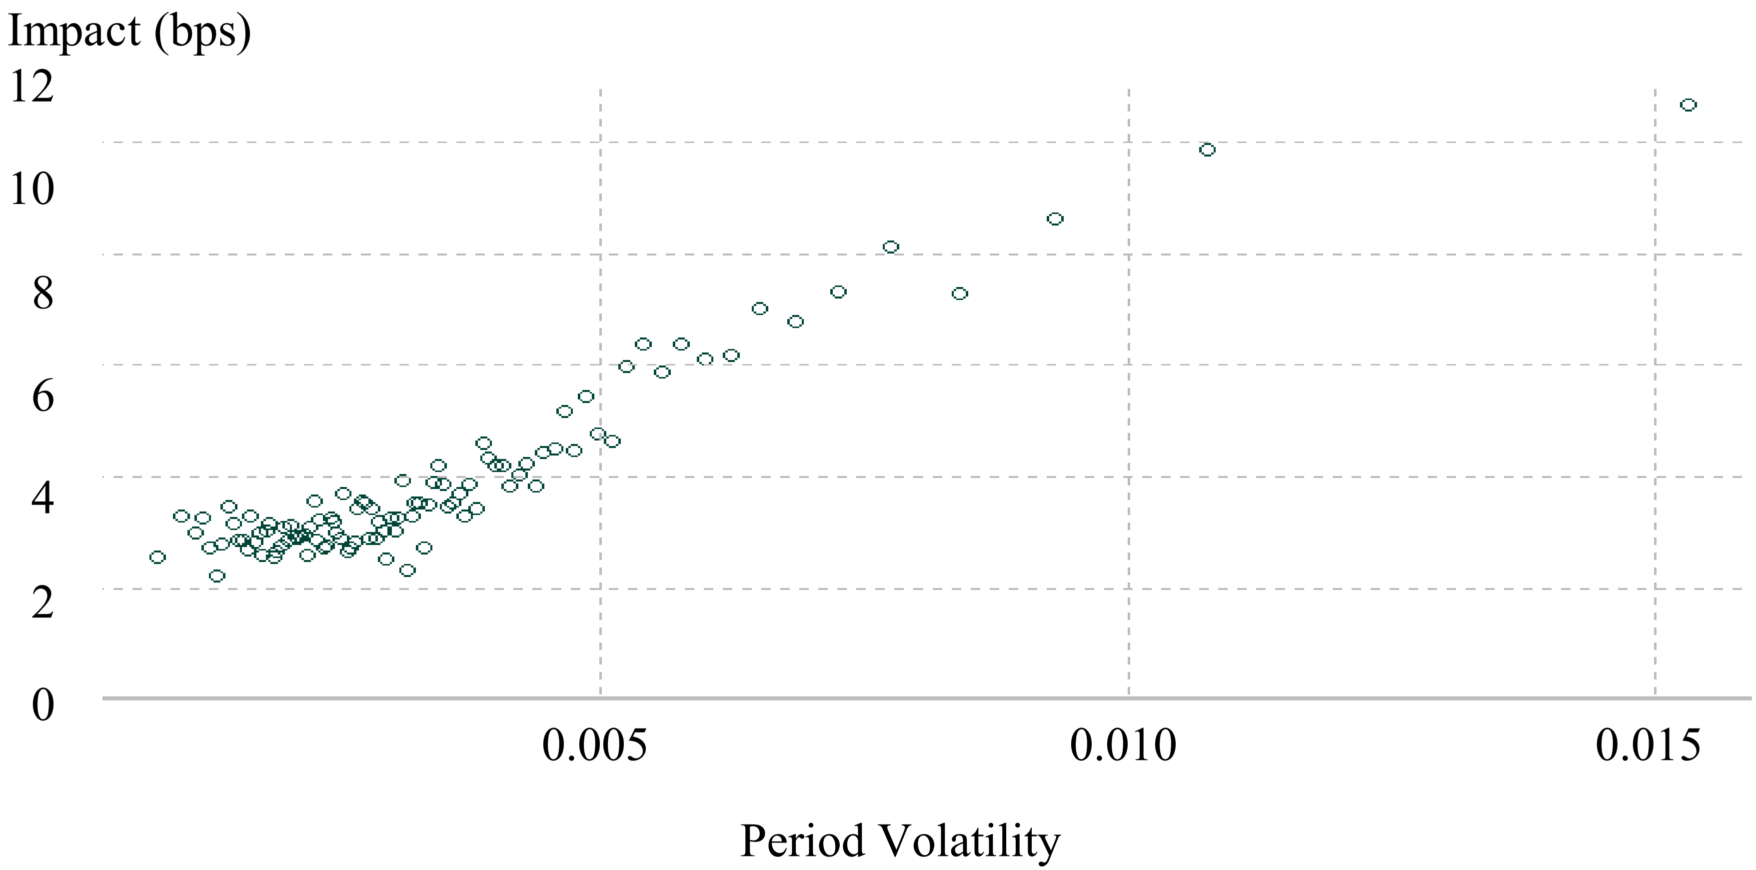

The Model

A third generation model

We leveraged this empirical analysis and previous impact modelling work to redesign our impact model

- First Generation Models
  - Total Cost = Fixed Cost + Size Cost
    - Fixed Cost = Alpha x Spread
    - Size Cost = Beta x Volatility x sqrt(% Expected Daily Volume)
  - The traditional inventory risk model (Grinold and Kahn) is a 1G model
  - A one day execution horizon is implied

The Model

A third generation model (cont'd)

- Second Generation Models addressed this flaw
  - Size cost = Beta x Volatility x sqrt(Expected Trade Rate)
  - Supports an arbitrary trade horizon
    - Impact vs execution risk efficient frontiers can be determined
    - An optimal horizon can be determined

The Model

A third generation model (cont'd)

- PRISE breaks away from 2G models in several ways
  - It is the result of a very detailed empirical analysis
    - Its structure is designed to reflect stylised empirical facts
    - It is calibrated on our extensive database
  - It supports intraday profiles
  - It breaks impact into permanent, transient and instantaneous terms
    - Different intra-trade time scalability properties of these terms
    - Different post-trade decay properties of these terms
    - Trade trajectory is also modelled
  - It supports arbitrary trade horizons and trade schedules

The Model

A richer framework

Example of Intra-Trade and Post-Trade Impact Trajectory

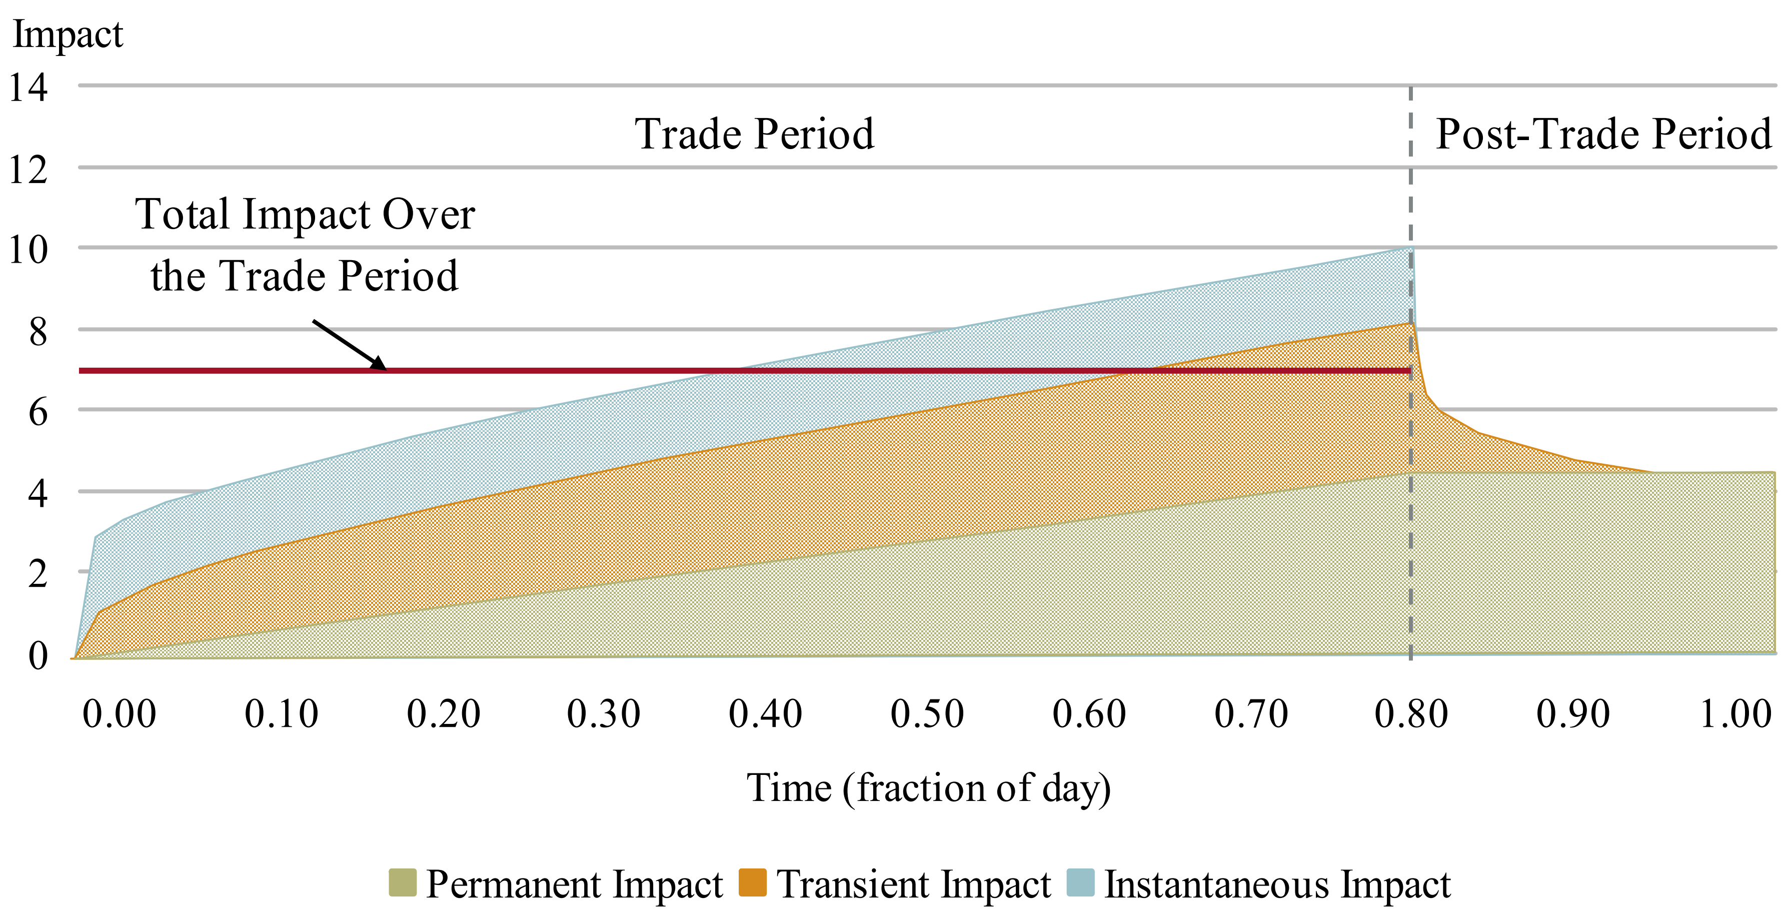

Permanent Impact

A richer framework

Example of Intra-Trade and Post-Trade Impact Trajectory

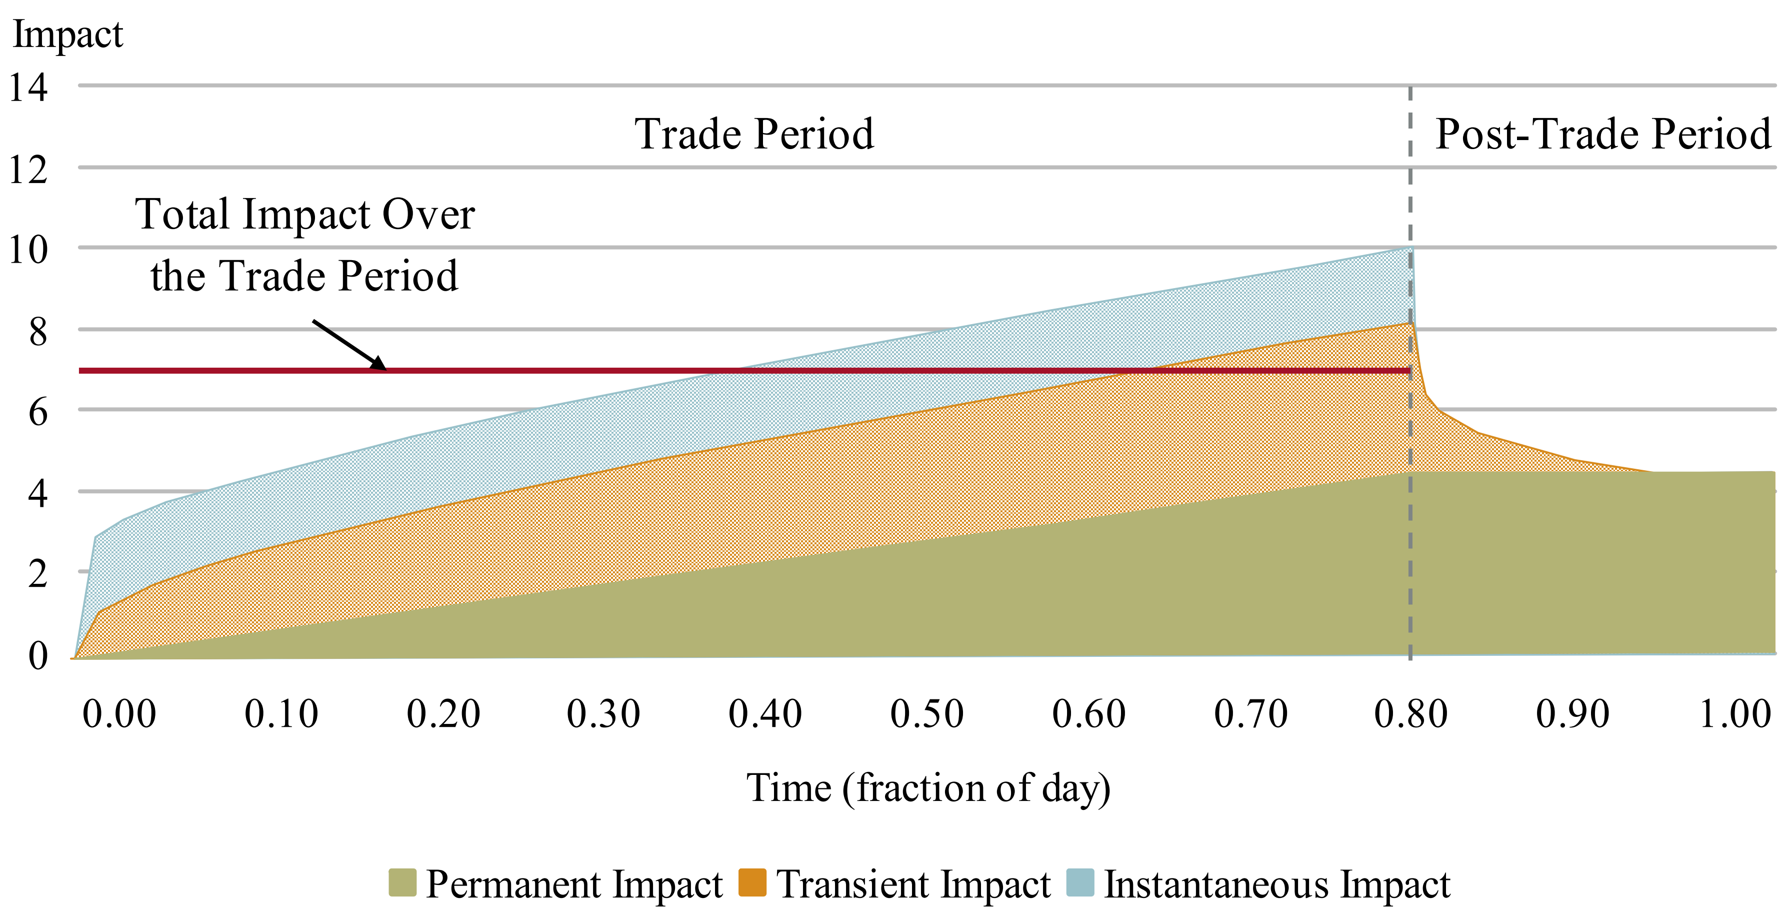

Transient Impact

A richer framework

Example of Intra-Trade and Post-Trade Impact Trajectory

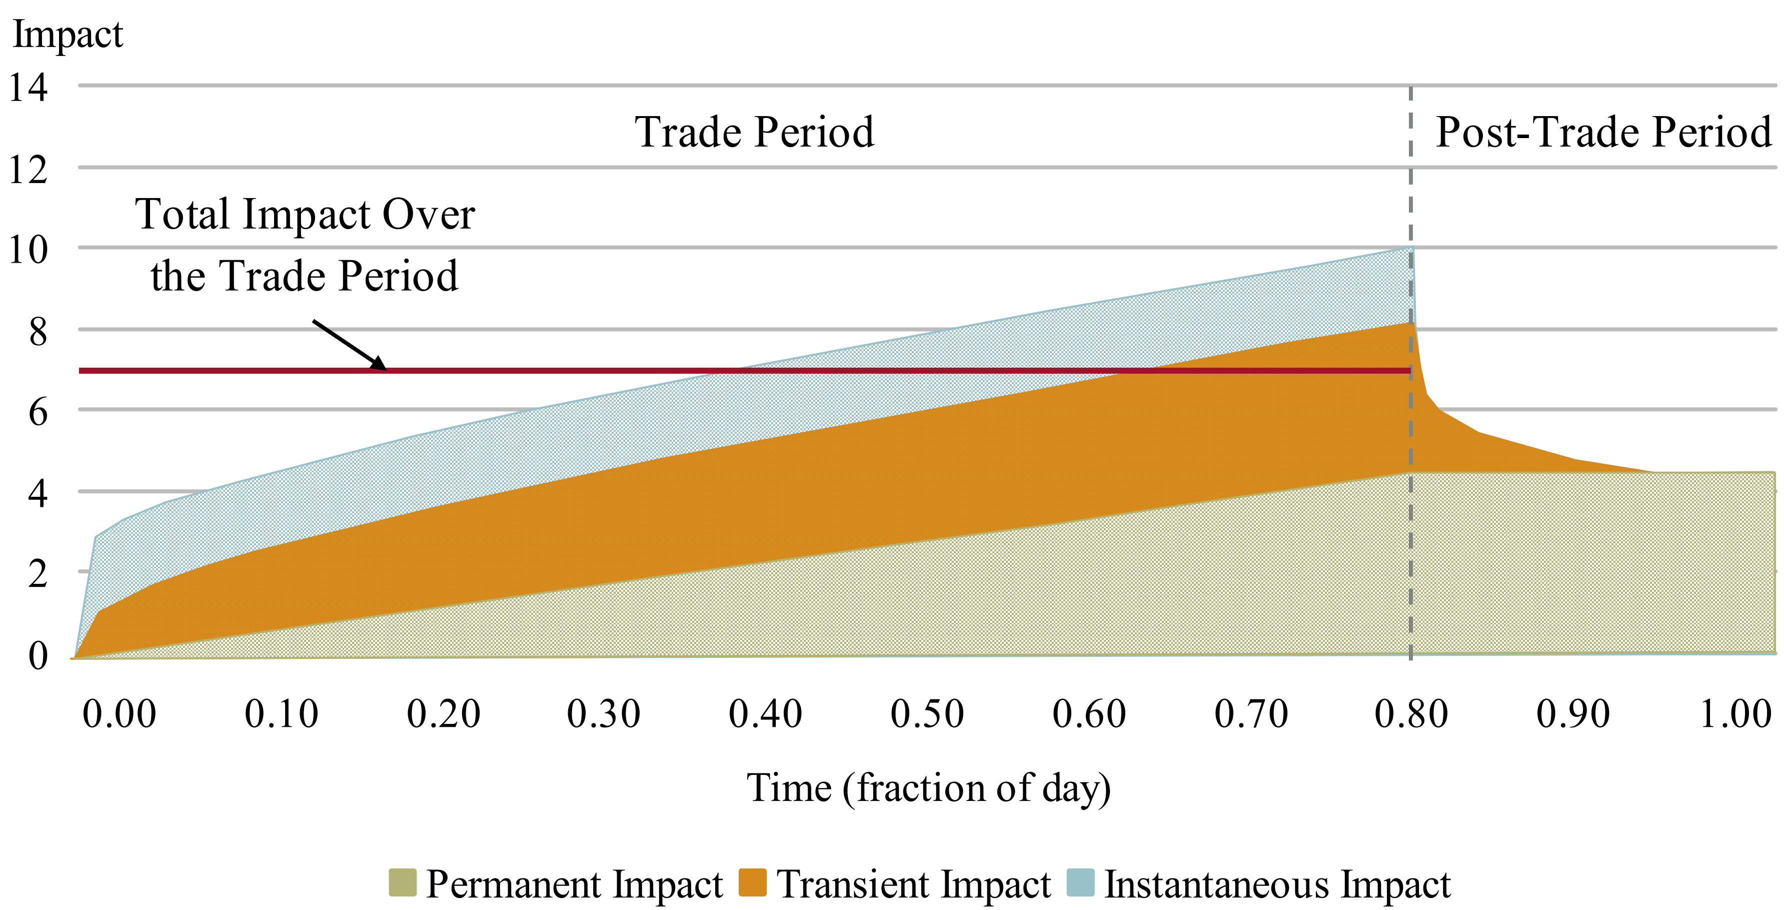

Instantaneous Impact

A richer framework

Example of Intra-Trade and Post-Trade Impact Trajectory

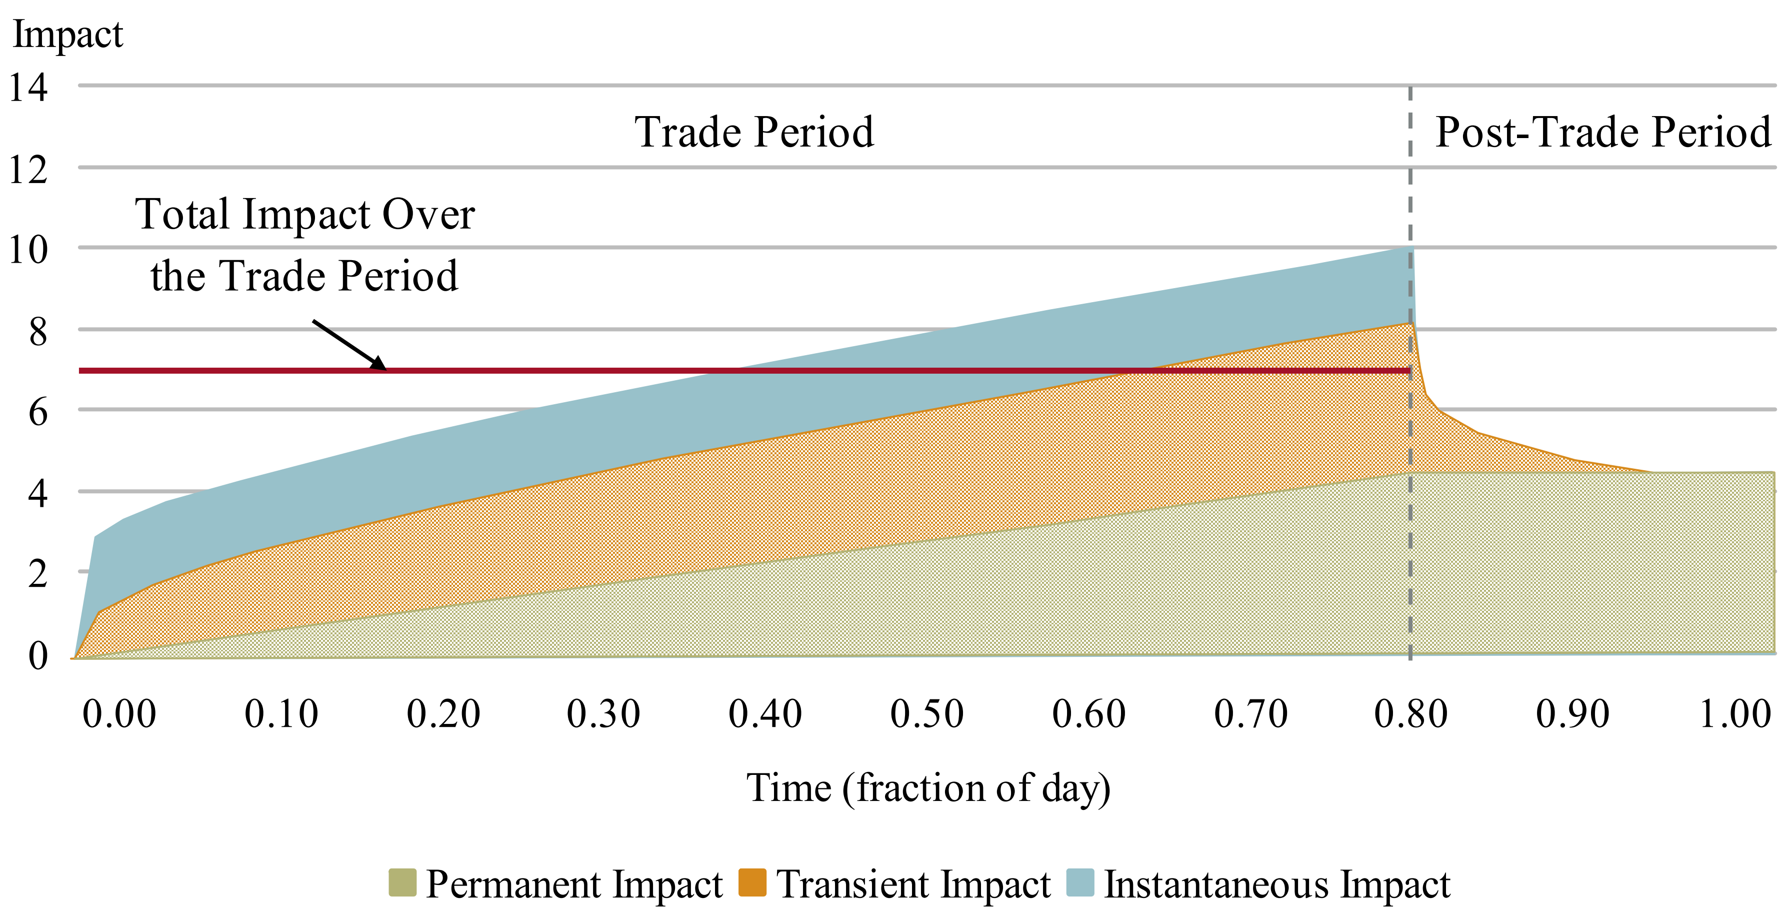

The Results

Excellent behaviour of the model

PRISE produces quantitatively accurate predictions of trade impact

Good average predictive power over the whole range of impacts

Predicted vs Actual Impact

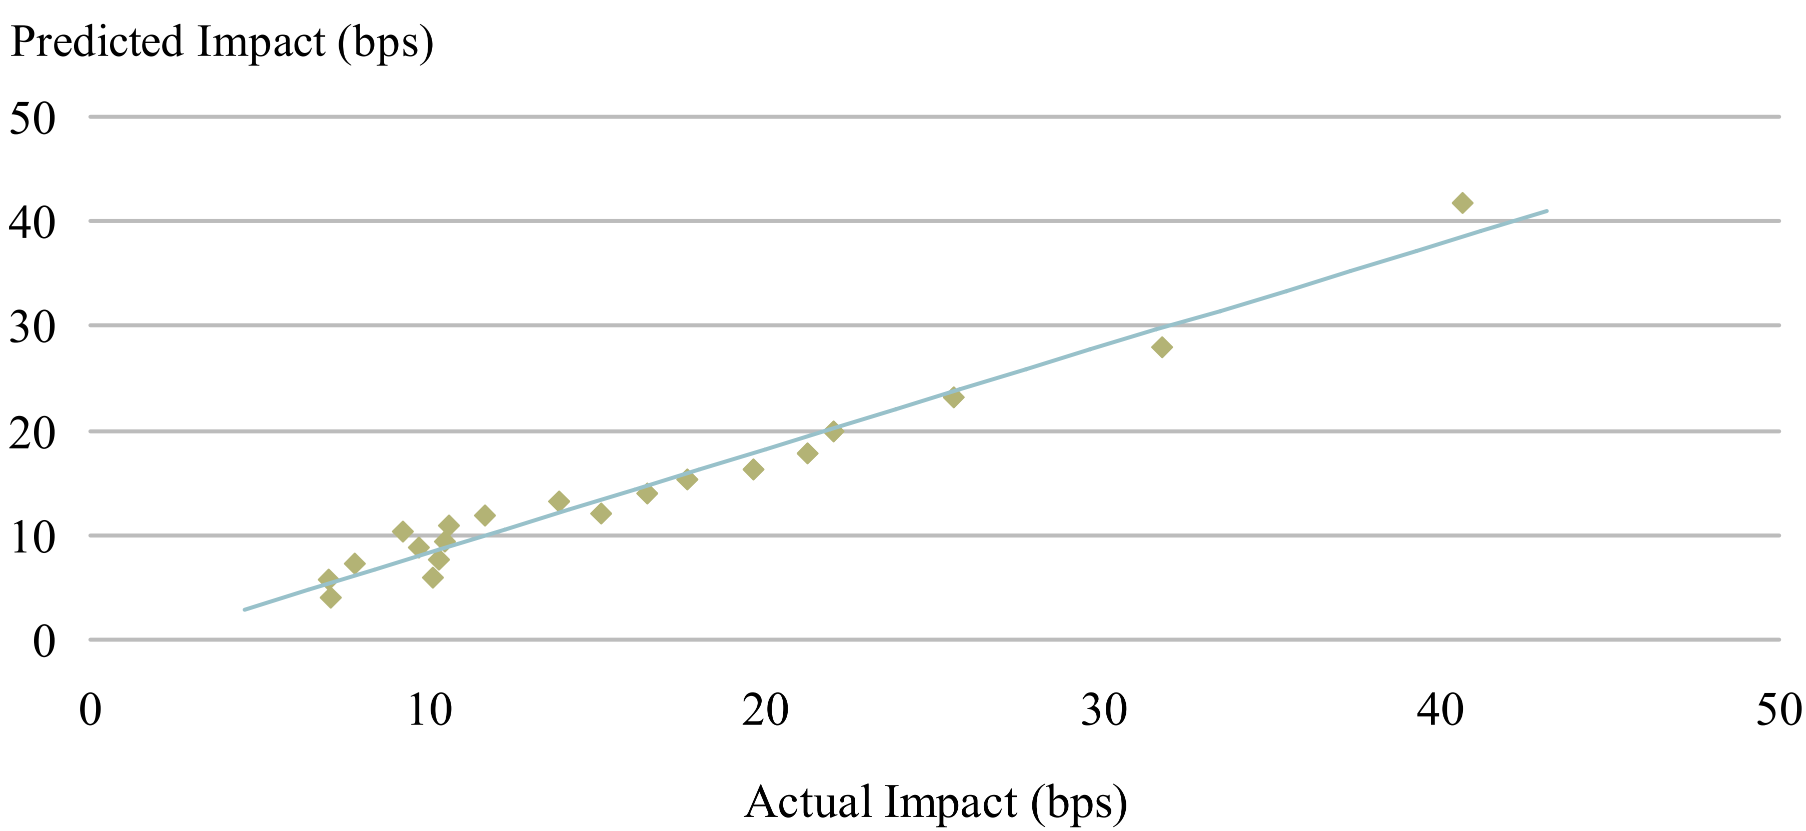

The Results

Excellent behaviour of the model (cont'd)

Good average predictive power over the whole range of trade rates

Similar results with a variety of other data samples and variables

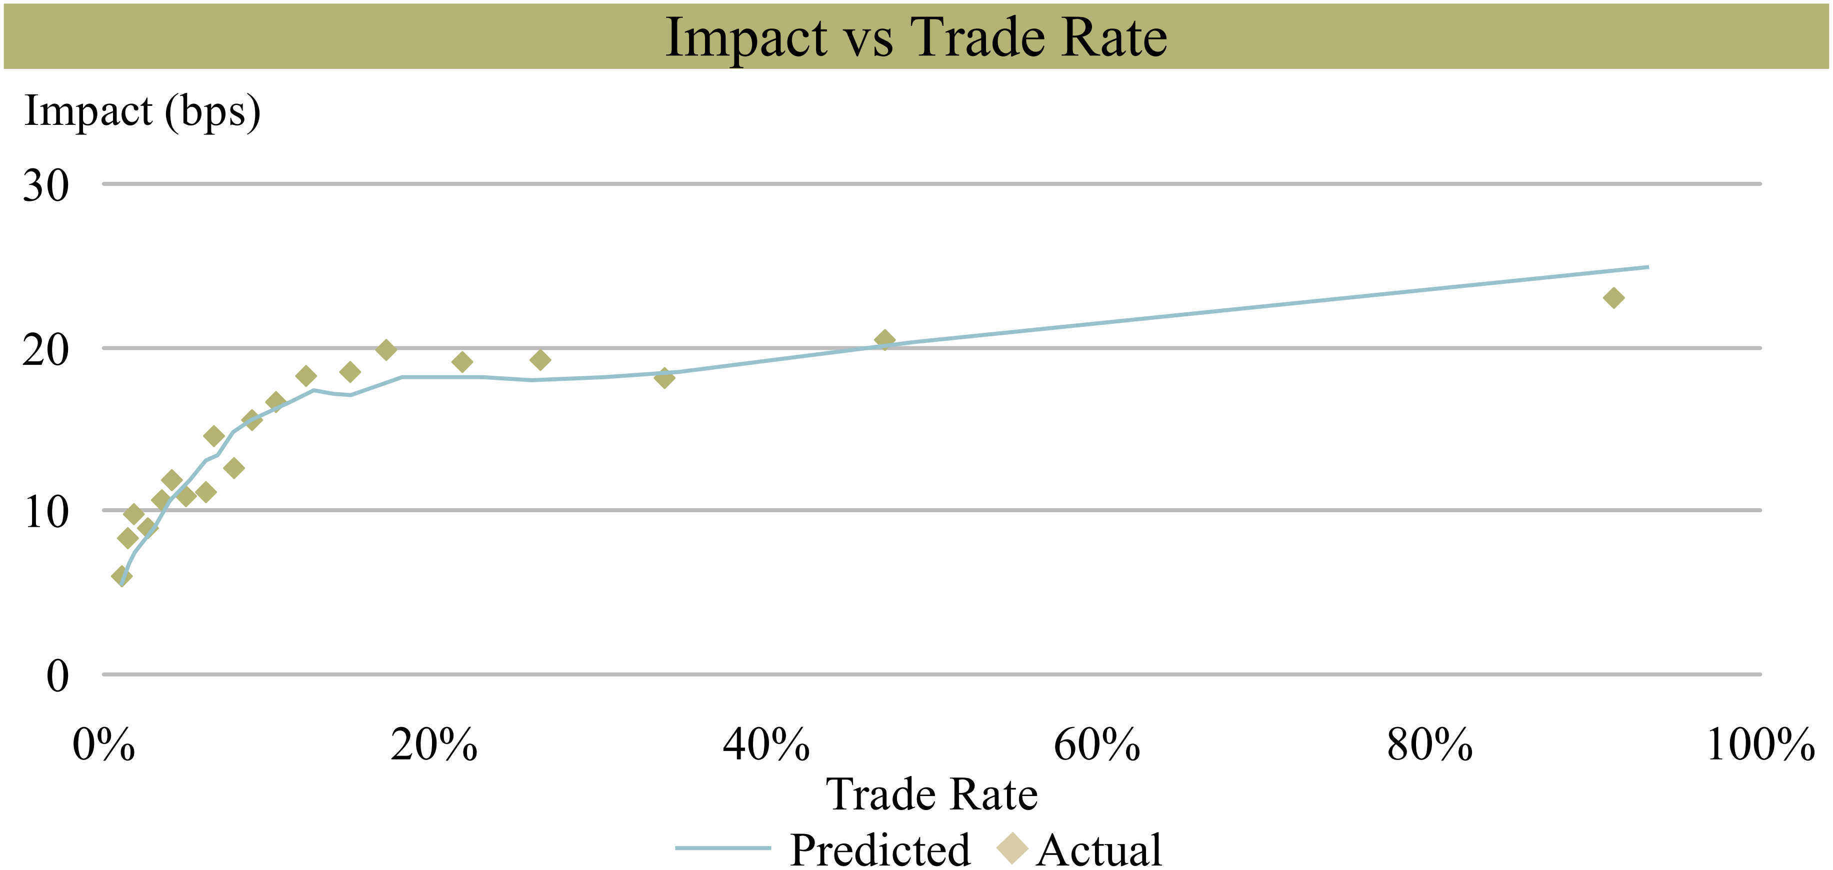

The Results

Impact vs spread

NYSE Out of Sample

Impact vs Spread (bps)

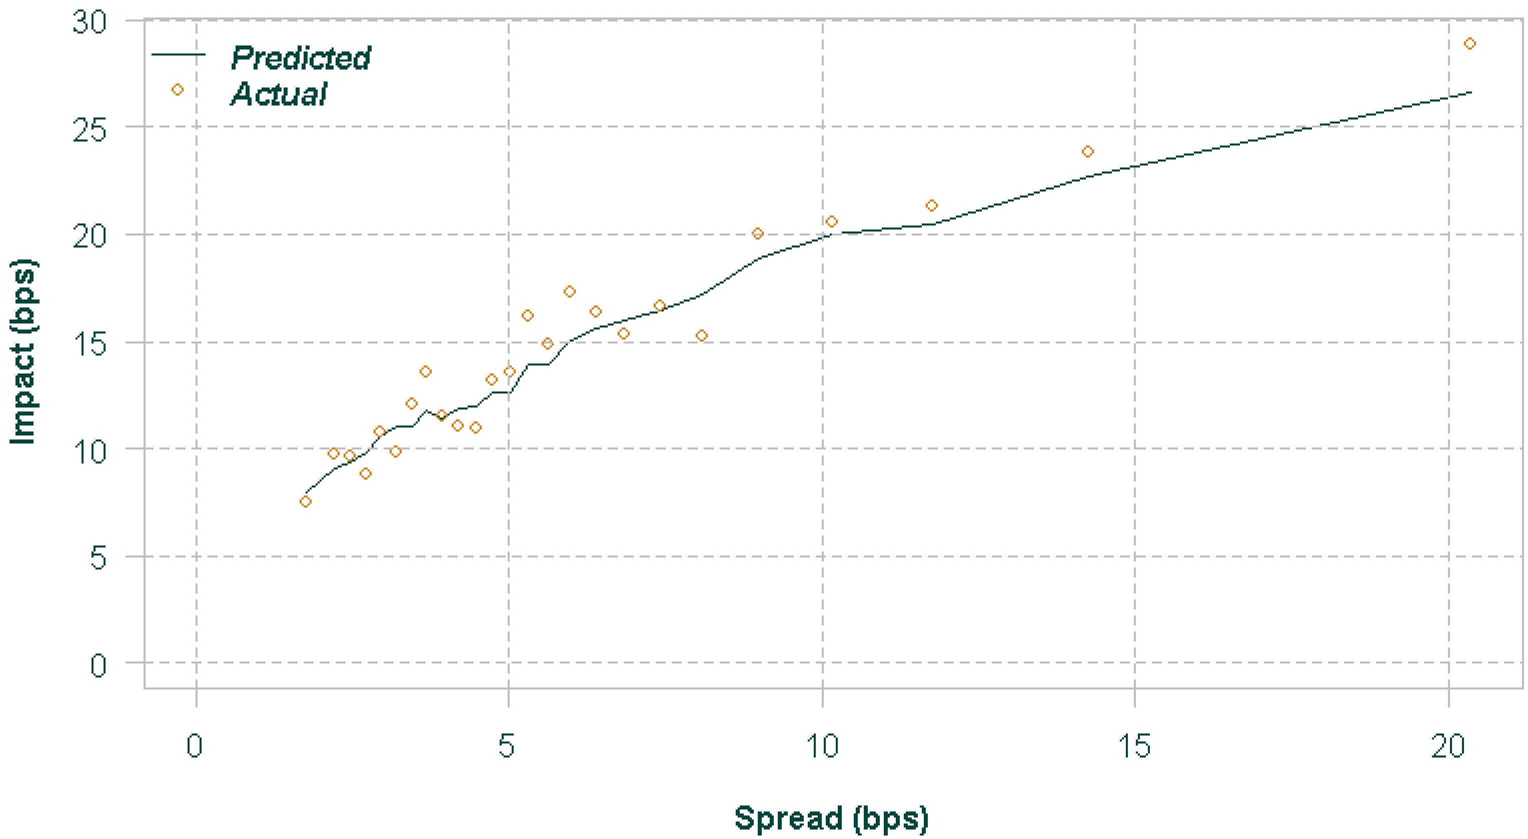

The Results

Example PRISE impact costs

- Estimated impacts of 1% ADV VWAP 30 minute trades
  - Exchange subsets of MSCI World
  - 30 November 2006

	Median Spread (bps)	Median Volatility	Median PRISE Impact (bps)
NYSE	4.6	21%	7.6
NASDAQ	4.3	28%	11.9
London	14.6	19%	5.9
Paris	9.5	18%	5.2
Tokyo	16.6	28%	12.5

Applications

A range of key applications

Transaction Cost Estimates are Crucial to all Stages of the Investment Process

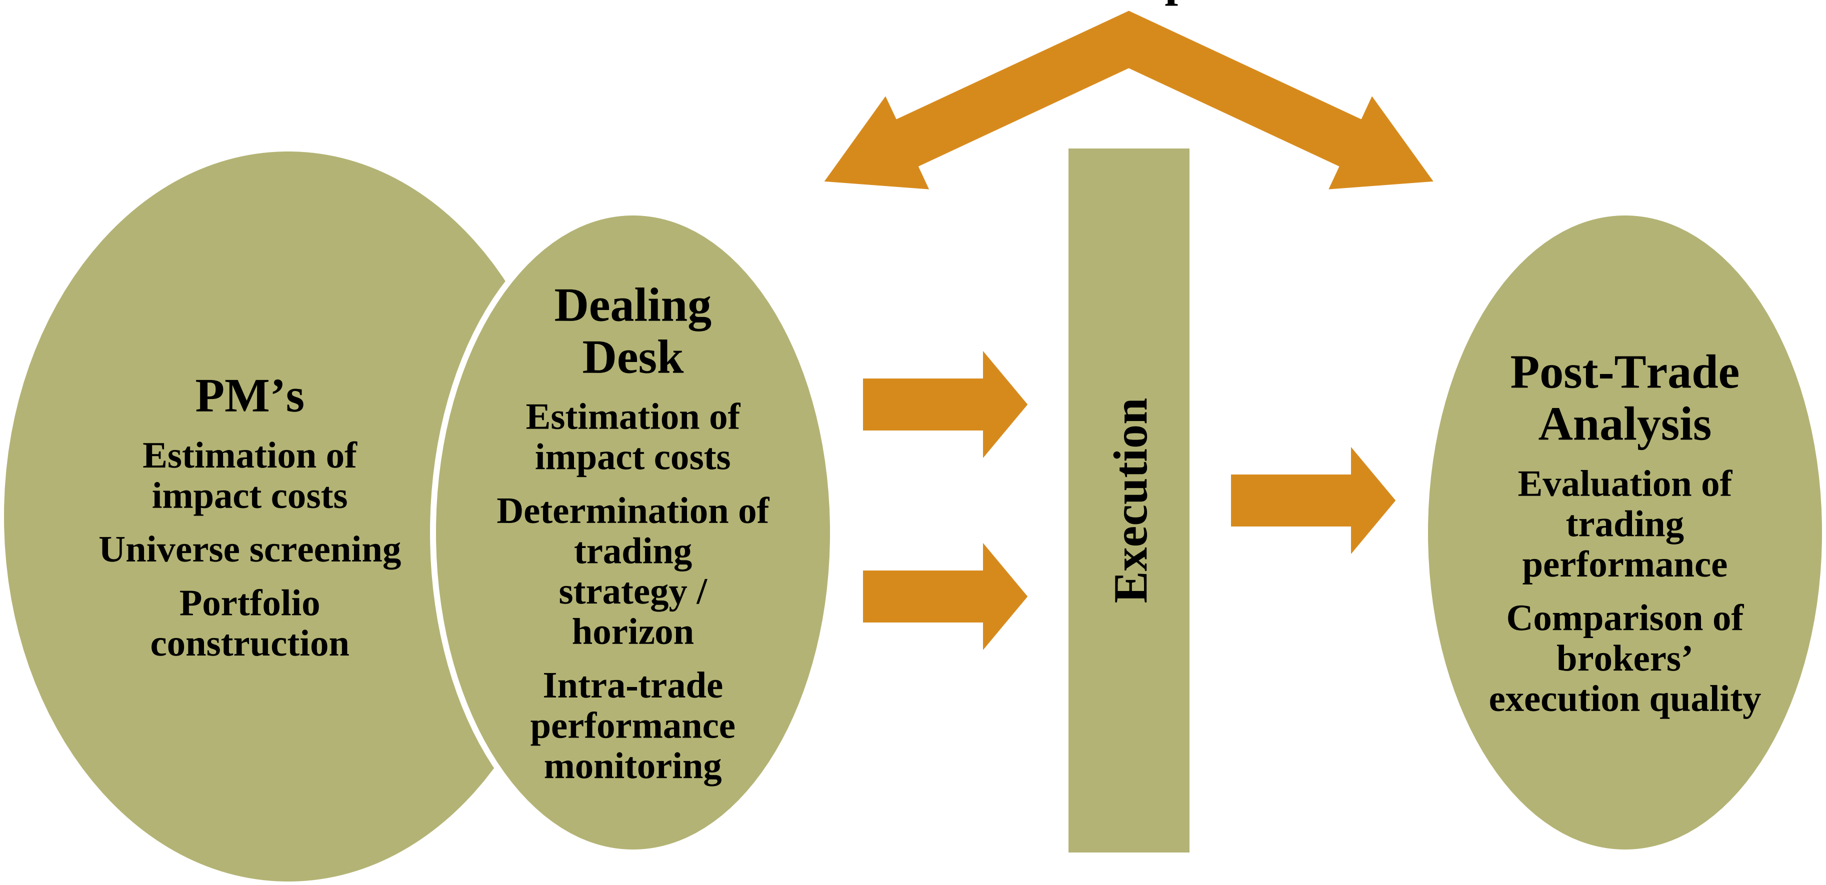

Application to Trading

Elementary Look

- The longer ones trade an order of fixed size on a prescribed stock the smaller the cost.

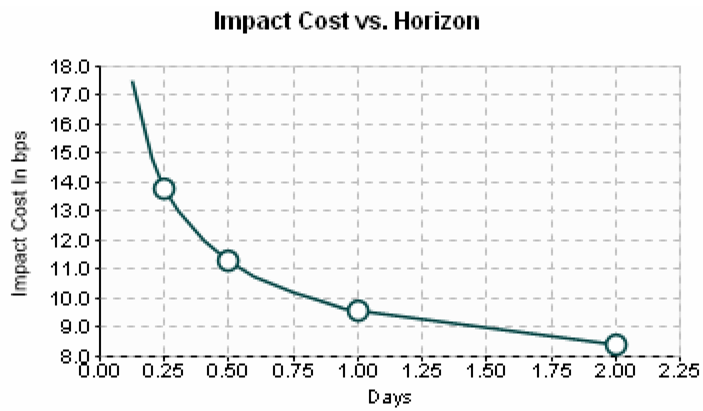

- Example: 10,000,000 shares of Vodafone. Cost reduces by about 25% when trading over ¼ of a day to 1 day.

- However risk increases substantially: In this case almost doubles from 40.2bps to 72.5bps!

Targetting Strike Price I

Mean-Variance Optimisation

- Can consider this as an minimisation of the following function:

  PRISE (Stock Specific, Order Size, Trading Horizon) +

  λ Risk (Stock Specific, Trading Horizon)

- Accepting that order size and stock are specified this is an optimisation of trading horizon. This would be simple (and provide a potentially analytical solution) for a very simple risk estimation ($\sigma T^{1/2}$) and simple cost model.
- Neither are the case though!
- However this can be optimised very quickly using commercially available quadratic optimisers: even when good quality cost models (PRISE) and intraday risk profiles.
- A further complication is that traders want (potentially qualitative) input to the optimisation process.

Targetting Strike Price II

Trader Input

- Allow the trader to control the value of λ (in qualitative ways) while bounding this by sensible values. With that in mind we introduce an idea of specifying “aggressiveness”.

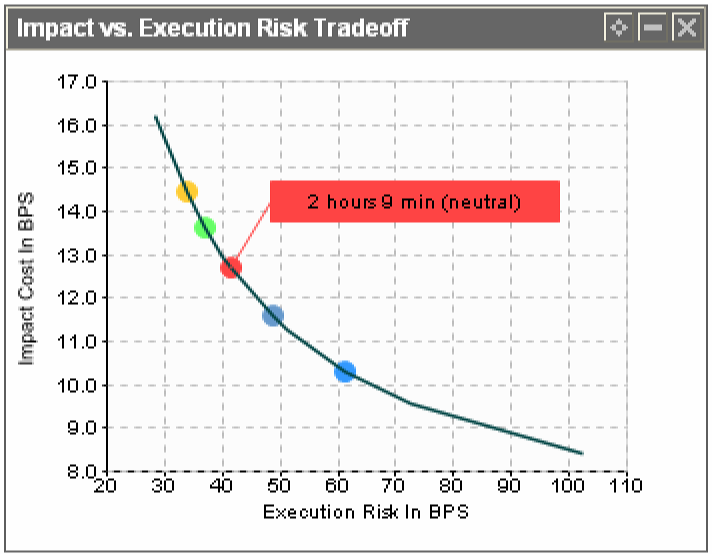

Impact Cost Over Time

Horizon	Cost (bps)	60% Conf. (bps)
1 hour	16.2 +/- 28.4	23.4
1 hour 24 min (very aggressive)	**14.5 +/- 33.7**	**23.0**
1 hour 42 min (aggressive)	**13.6 +/- 37.0**	**23.0**
2 hours	12.9 +/- 40.2	23.1
2 hours 9 min (neutral)	**12.7 +/- 41.6**	**23.2**
2 hours 55 min (patient)	**11.6 +/- 48.6**	**23.9**
1/2 day	11.3 +/- 51.3	24.3
4 hours 38 min (very patient)	**10.3 +/- 61.2**	**25.8**
1 day	9.6 +/- 72.5	27.9
2 days	8.4 +/- 102.5	34.4

- If we are aggressive then we are anticipating something we believe is neglected from the model that will increase volatility or spread or price explicitly and hence we need to increase the relative value of cost. If we are passive then we believe the converse.

Targetting Strike Price III

Trader Input

- Furthermore do not need to package as a "trading strategy" but can provide analytical products and let the trader route his own orders to simpler strategies (VWAP/WithVolume).

Comparison of Projected Execution Statistics										
	Risk and Cost Estimates				Strategy Parameters					
Strategy	Total Cost		Total Risk		Cons. Rate	Start Time	Duration (T+0)	Complet. Time	Complet. % (T+0)	Resid %
	bps	cps	bps	cps						
VWAP (Now to Close)	12.3	0.4	47.5	1.4	13%	14:46	1 hr 44 m	CLOSE T+0	100%	0%
WithVolume (10%)	11.4	0.3	54.3	1.6	10%	14:46	1 hr 44 m	09:40 T+1	77%	23%
TargetStrike Patient	11.9	0.4	49.7	1.5	12%	14:46	1 hr 44 m	08:35 T+1	91%	9%
TargetStrike Neutral	13.0	0.4	42.6	1.3	16%	14:46	1 hr 44 m	16:30 T+0	100%	0%
TargetStrike Aggressive	14.0	0.4	37.9	1.1	20%	14:46	1 hr 29 m	16:15 T+0	100%	0%

TargetStrike

Advantages and Disadvantages

- Do not need to study the stock.
- With not trade "too thinly" exposing oneself to opportunity risk or trade "too rapidly" exposing oneself to sharp price moves.
- Aims to minimise VWAP Slippage insomuch as is possible while minimising Strike Slippage.

- May take multiple days to trade (stock borrow costs, overnight risk consequences).
- With portfolios can provide intra-trade exposure to risk factors. For example 25,000,000 USD Long-Short for British Telecom/Ericsson. Would trade British Telecom over the day and Ericsson in 2 hours 16 minutes. So are exposed to Telecoms for more than 2/3 of a trading day.

Conclusions

- Study of market microstructure and mathematical optimisation can provide a practical tool to:
  - Trade Cheaply
  - Trade Stocks/Programs where the trader has little experience of the stock.
  - Intelligent Post-Trade Analysis
  - Understand Liquidity Constraints when constructing Portfolios

Further Study

- Continue to Improve Intraday Liquidity Modelling
- Correlated/Portfolio Transaction Cost Model
- Improve/Adapt for Non-Electronically Traded Markets
- Continue to Improve Volume Prediction

Disclaimer

This document is for information purposes only and it should not be regarded as an offer to sell or as a solicitation of an offer to buy the securities or other instruments mentioned in it. Any security or market data that is mentioned or used in this document is intended for illustrative purposes only. No part of this document may be reproduced in any manner without the written permission of Lehman Brothers Inc. We do not represent that this information is accurate or complete and it should not be relied upon as such. Opinions expressed herein are subject to change without notice. The products mentioned in this document may not be eligible for sale in some states or countries, nor suitable for all types of investors; their value and the income they produce may fluctuate and / or be adversely affected by exchange rates, interest rates, or other factors. Additional information will be provided upon request.

Lehman Brothers Inc and / or its affiliated companies may make a market or deal as principal in the securities mentioned in this document or in options or other derivatives based thereon. In addition, Lehman Brothers Inc, its affiliated companies, shareholders, directors, officers and / or employees, may from time to time have long or short positions in such securities or in options, futures, or other derivative instruments based thereon. One or more Directors, Officers, and / or Employees of Lehman Brothers Inc or its affiliated companies may be a Director of the issuer of the securities mentioned in this document. Lehman Brothers Inc or its affiliated companies may have managed or co-managed a public offering of securities for any issuer mentioned in this document within the last 3 years.