09 ——————— 19

Delivering
Algo Performance
Through Enhanced
Market Simulation

A Bluepaper from Simudyne

Simudyne ——— Reality, re—engineered

Introduction

Sell-side institutions today are facing a highly competitive marketplace that is being constantly reshaped by client demands, new technologies and changing regulatory requirements.

With research being unbundled from execution as a result of global regulatory measures, and with low-latency connectivity to trading venues increasingly seen as a standard offering, quality and performance of trading strategies and algorithmic models becomes an important service differentiator to buy-side customers.

Achieving competitive performance levels requires rigorous testing and performance evaluation on realistic market data. Yet firms often find themselves limited by their ability to backtest their algos beyond available data; they struggle with overfitting, multiple testing bias and the problems associated with using limited (real) data in backtesting – or they may be unable to calibrate their algos quickly enough to respond to client requests.

These limitations restrict a firm's ability to qualify for inclusion in the 'algo wheels' operated by their clients: the lists of preferred (and approved) broker trading strategies they are able to use.

This bluepaper - based on a survey of 15 sell-side executives involved in algo trading - evaluates current challenges around trade execution and algo testing in simulated market environments and explores their current operating procedures.

It also discusses how a new approach to developing a highly granular, dynamic representation of the markets using agent-based simulation (ABS) techniques could help reduce bias and focus on dynamics of interest, such as stylized facts, allowing firms to optimise execution system performance as well as the ability of their algorithms to attract and retain buy-side liquidity.

2.0

Facing a dynamic, competitive environment

Competition among sell-side firms is fierce. New regulation has altered the way firms interact with their clients, the markets and each other. Regulatory changes and advances in technology are resulting in ever-faster changes to market structures. The result is that the idea of static trading strategies is dead and sell-side firms need a strategy to evolve and develop at pace.

Traditional differentiators have lost their potency. Research has been unbundled from execution by the EU's MiFID II regulation, a stance that is being emulated by regulators across the globe. Meanwhile, crippling costs and the laws of physics have stifled the impact of ultra low-latency connectivity as a service differentiator for all but a very few deep-pocketed specialists. In this environment, sell-side firms can find it difficult to differentiate themselves from competitors in the eyes of potential clients.

The 'secret sauce' embodied in firms' trading strategies and algorithmic models represents a rare differentiator that can help prime brokers and other sell-side firms stand out from the pack. Indeed, firms now see attracting client order flow as the main purpose of their investments in algorithmic trading, with 62% of survey respondents citing it as a key driver of their algo development activities.

Regulatory compliance (cited by 25% of respondents) remains a priority, despite the January 2018 implementation date of MiFID II having passed.

" —— What did we spend the most time on over the past six months? It's probably regulatory compliance. But what is my priority? It is getting more additional order flow."

This wide-ranging regulation altered the competitive landscape for firms in many ways, but it impacted the use of algorithms explicitly. For one thing, MiFID II requires firms to test their algos before using them in production model, fearful of a repetition of the so-called Flash Crash in 2010, in which the Dow Jones Industrial Average plummeted by almost 1,000 points, losing 9% of its value – only to recover it within 36 minutes. The use of spoofing algorithms took at least some of the blame for the crash; regulators responded by tightening the rules around the use of untested algorithms. MiFID II also raised the bar on best execution, introducing more stringent

obligations than its predecessor on sell-side firms to provide evidence of adherence to their published best execution policies. By extension, investment management clients now have an obligation to disclose best execution information to their end-customers, increasing transparency about how funds have been invested and transaction costs.

As a result, institutional trading desks frequently turn to transaction cost analysis (TCA) to measure costs, optimize trades and ensure best execution compliance. In turn, this has led to the emergence of the algo wheel, software that buy-side firms use to list those broker-provided algorithms approved for use by their traders. The algo wheel itself uses its own algorithms and trade data to help decide which algorithm among those listed would provide the best result, based on a range of criteria.

The US Securities and Exchange Commission, meanwhile, published amendments to Rule 606 of its Regulation NMS last November that imposes further disclosure requirements on sell-side firms with respect to client orders. Under the new language, "a broker-dealer, upon a request of a customer who places a 'not held' order (e.g., an order in which the customer gives the firm price and time discretion), [must] provide the customer with a standardized set of individualized disclosures concerning the firm's handling of the customer's orders. The new disclosures will, among other things, provide the customer with information about the average rebates the broker received from, and fees the broker paid to, trading venues."

These new disclosures are designed to help investors better understand how the broker-dealer routes and handles their orders and assess the impact of their broker-dealers' routing decisions on order execution quality. Such regulatory initiatives have altered the objectives of institutional clients when placing a trade, with consistency of performance a more important consideration than cost.

A majority (53%) of the sell-side executives interviewed for the survey said consistency of performance was the top priority for customers when selecting trading strategies or algorithms. This compared with a quarter (27%) who cited client choice based on TCA or cost. "Different clients have different goals," said one respondent. "However, I would say that generally clients are trying to improve their performance versus some TCA benchmark." About 20% of respondents said clients are looking to optimise performance on a specific trade.

3.0

Attracting and Retaining Client Order Flow: Getting on the Algo Wheel and Staying There

As a result of these changing dynamics on the buy-side, sell-side firms are facing a variety of challenges as they seek to attract new client business and order flow. Survey respondents gave almost equal weight to several challenges.

Obstacles To Attracting Buy-Side Order Flow

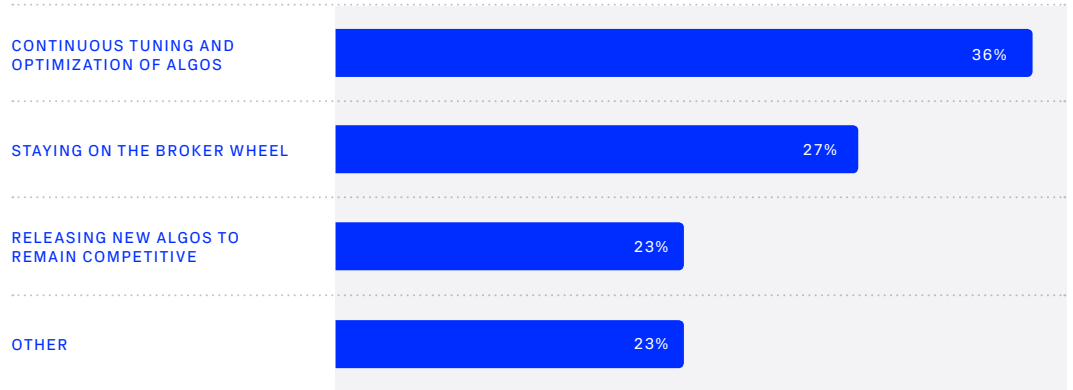|	
APPLICABILITY OF AN ALGO TO MEETING CLIENT OBJECTIVES	31%
PROVIDING EVIDENCE OF ALGO PERFORMANCE	23%
ACHIEVING BENCHMARK PERFORMANCE	23%
DIFFERENTIATING VS. OTHER SELL-SIDE ALGO OFFERINGS	23%

Some 31% indicated their biggest challenge was demonstrating the applicability of an algo to meeting client objectives. Another 23% said their biggest challenge was differentiating vs. other sell side algo offerings, whilst a further 23% said they were challenged to achieve benchmark performance under different market conditions.

However, 23% cited factors around providing evidence of algo performance. Said one respondent, "For us, proving our capability is the biggest challenge. It's about proving that our algos are better than others." Says another, "the hardest thing is getting the opportunity to demonstrate what it is that we can do. Regardless of how you pitch it, I'm not sure clients are particularly fluid in their choice of providers. So, we are really fighting for that opportunity."

Other respondents talked about the impact of regulatory change. One mentioned the new US SEC changes to Rule 606. Another said, "MiFID II, Brexit, take your pick. They just keep changing the regulations…this is our biggest impediment currently."

Once they've got on the algo wheel, retaining order flow is also very challenging for sell-side firms.

Obstacles To Attracting Buy-Side Order Flow

CONTINUOUS TUNING AND OPTIMIZATION OF ALGOS	36%
STAYING ON THE BROKER WHEEL	27%
RELEASING NEW ALGOS TO REMAIN COMPETITIVE	23%
OTHER	23%

Continuous tuning and optimization of algos was perceived as challenging by 36% of respondents and staying on the broker wheel is the biggest challenge for 27% of firms. Said one respondent, "It's about getting on the Algo Wheel. It's easy once you are on it. Most clients don't have enough flow to warrant an extended wheel, so there are relatively few slots available. If a client trades only 10 million shares, then it takes a lot of infrastructure to connect to brokers, so they minimize the effort. Then, how do they know they are getting best execution from the few brokers on their wheel?"

Algo wheels came in for quite a bit of criticism from survey participants. Says another, "There are lots of problems with them, to do with the way people implement them. Most of our clients don't have enough flow to actually use a broker wheel, but if they are, they start to get spurious results. So, you can be unlucky and be kicked out. For example, if they base their decision on a small sample size, and within that sample size it suddenly looks like you are doing badly."

"Another problem is setting yourself up correctly for whatever you are being measured against, but then being set up the wrong way for their measurement," the respondent continues. "The client can tell you one thing, but then you can misinterpret how their benchmarks work, or how their orders work, so you are being measured in a different way to how you thought you were going to be measured. There is a lot of luck in this, I'm afraid."

A further 18% said that releasing new algos to remain competitive with existing clients is their biggest challenge. "Everyone comes out with a bright and shiny toy, and then everyone wants to see how it works," says one respondent. "Whoever was with the client last, and presented them with a new product," gets the business, he says.

4.0

… Against the Backdrop of Expense and Complexity of Changing Regulatory Requirements

Woven through all of this are the many regulatory challenges that the execution desk faces, with algo performance striking at the heart of the impact of the market structure changes wrought by MIFID II and other regulations. It was clear from the survey results that the effects of these changes are creating extra levels of complexity for firms.

According to the survey results, MiFID II's RTS 6 is viewed as a significant challenge, cited by 36% of respondents. RTS 6 specifies the organisational requirements of investment firms engaged in algorithmic trading, including the testing of algorithms, the presence of risk and controls as well as governance processes, and controls and monitoring to prevent algorithmic tools from disrupting orderly markets. The unbundling of research and execution (9%) and the implementation of trade surveillance (9%) were viewed as important regulatory challenges.

Many respondents cited other challenges posed by the new regulatory landscape, among them:

• Managing for the changing liquidity landscape and calibrating algorithms accordingly. "Suddenly you now have certain orders that you cannot send to certain venues under certain circumstances," said one executive. "And then you have clients opting out of venues and all that stuff… the impact of all that choice has been huge."

• Keeping up with what will be required for Brexit, which will depend on the form that this takes. For example, a hard Brexit could impact where certain equities are traded.

• The US SEC's Rule 606 reporting. Says one respondent, "these requirements are just bad, they are more convoluted than they should be."

• Reporting requirements under best execution are burdensome for some respondents.

In summary, sell-side firms are faced with a variety of new challenges today.

"⸺ Suddenly you now have certain orders that you cannot send to certain venues under certain circumstances, and then you have clients opting out of venues and all that stuff… the impact of all that choice has been huge."

Although these challenges may have had their origins in regulatory changes in both the EU and the US, the business impacts have been enormous. For example, MiFID II has completely changed the market's approaches to algorithm trading as a result of the new algo testing requirements, more rigorous best execution rules and transparency requirements, and the emergence of the algo wheel. In this environment, sell-side firms are seeking to be more competitively nimble through new algorithms. However, these same regulations can make the launch of, or re-tuning of algorithms more operationally challenging than in the past.
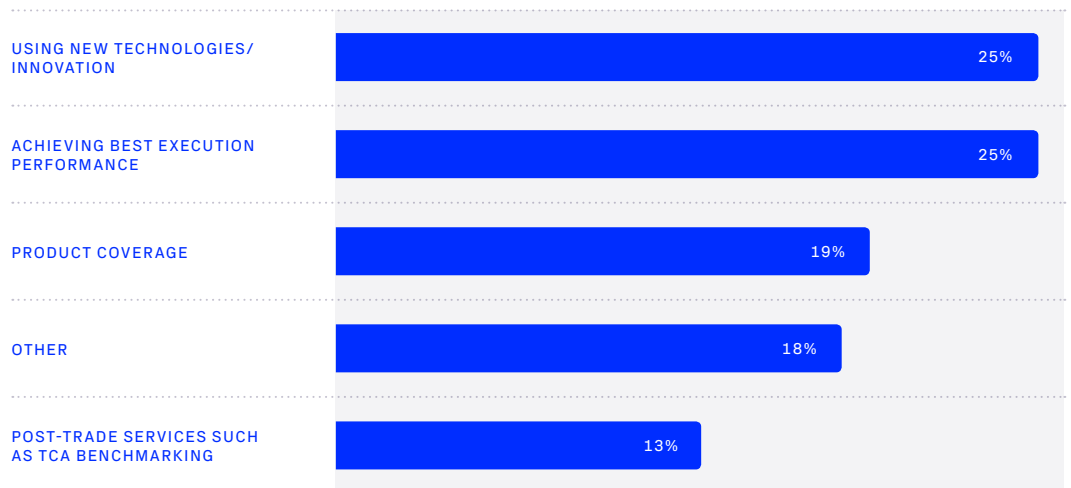
Understanding current operational limitations

Today, the pressure is on for firms to differentiate their algos from competitive brokers and execution desks, but this is not easy.

Key Drivers of Algo Optimisation Efforts

USING NEW TECHNOLOGIES/ INNOVATION	25%
ACHIEVING BEST EXECUTION PERFORMANCE	25%
PRODUCT COVERAGE	19%
OTHER	18%
POST-TRADE SERVICES SUCH AS TCA BENCHMARKING	13%

Key drivers for those seeking to improve their algorithmic offerings include:

• Achieving best execution performance (25%)

• Using new technologies/innovation (25%)

• Product coverage (19%)

• Post-trade services such as TCA benchmarking (13%)

• Other (19%) – this included best execution reporting as well as a more personalized relationship with clients.

Survey respondents said they'd invested in several ways to their capacity to create algorithms, including people (PhDs from MIT and Stanford) and technology systems (co-located smart routers, high-performance market data stacks, etc.). Said one executive, "New tech is key. We've invested in this. We can now customize an algo in five minutes and make it bespoke to a specific client's portfolio." Talking clients and prospects through the physical capabilities of the sell-side firm is important.

Executives also spoke of the need to capture metrics that provide quantifiable evidence of what they are achieving on behalf of existing clients, and what they would accomplish for prospects if given the opportunity. There was also the additional challenge of getting these metrics in front of prospects. Metrics include best execution data, independent third-party TCAs, internally calculated TCAs, and information about access to liquidity. Said one respondent, "For some clients, it's really about education. For more sophisticated clients, it's about the customization of reporting."

Algo Pre-Production: Key Challenges

Sell-side firms are finding themselves extremely challenged in algo preproduction–across the range of issues. Different firms are challenged by different things. In the survey, the averages for each challenge answer option were nearly the same across all five possibilities. However, these close averages mask wide variability in the underlying answers. What was very challenging for one firm, another firm considers much less challenging. For example:

- 44% of firms rated of the cost of exchange data as not very challenging or challenging (a 1 or 2), while 55% rated it as challenging or very challenging (a 4 or 5).

- Half of executives rated modeling for new regulatory and structural challenges to the market as challenging or very challenging, while 20% viewed it as not challenging or very challenging.

- 40% said modeling against unusual market conditions was challenging or very challenging.

- 40% indicated that modeling macroeconomic shocks was challenging or very challenging.

- 50% said finding high quality data was challenging or very challenging, while 40% found this not challenging or very challenging.

While different firms seemed to be challenged by different things, adding to the complexity of the current operating environment, modeling was a common, significant and recurring theme. Many firms are finding modeling to be either challenging or very challenging within algo pre-production, and are concerned that this is holding them back competitively.

The methodologies that firms use to model market conditions to test the behaviors of their algos are fairly standard. Most use a combination of historical data and simulated data to put their algos through their paces. "We built a market simulator in pre-production to test all of our changes in UAT," says one trading executive. "We use it extensively for regression testing as well as for testing new features, to make sure they are delivering the expected behaviors and results." This firm also uses A/B testing to help generate results faster.

"We have our own development environment where we are able to run historic simulations so that we can put extreme situations into that modeling environment," said a trading executive. "You cannot model for every single market scenario, so it still requires some kind of controlled release process. There is no perfect modeling environment."

Indeed, executives recognize the limitations inherent in these modeling approaches. Said one executive, "You don't know what you don't know. No matter how comprehensive our testing is, there will always be scenarios that we didn't think of, or events in the real world that are not covered by the tests." Said another: "This is more of an art, I'd say." All of the executives interviewed were only able to adjust a limited number of parameters on the fly – parameters that had been

" ——— You don't know what you don't know. No matter how comprehensive our testing is, there will always be scenarios that we didn't think of, or events in the real world that are not covered by the tests."

'pre-set' in advance. For some firms, the number that they could adjust was very limited. "There are parameters that you never want to adjust on the fly. But there are other kinds of parameters that you need to adjust on the fly – for example, the passiveness or aggressiveness of the algo, to suit the client's wishes," said one executive. "Other changes and adjustments you need to test after you make them, to be sure they are working OK. You don't want to adjust them on the fly. So, you wait until after trading hours and then test them and roll them out."

" ——— This would allow us to test more quickly and ensure that our algos are performing without having to do as much running by, it would allow us to make changes more efficiently, more safely across our algo base."

Firms on average rated themselves as only middling in their ability to do this – scoring three on a scale of one to five. In short, executives recognize that to be competitive in this new, highly regulated environment, they must be more nimble when it comes to meeting the needs of clients and prospects.

This can include modifying algos to suit the approaches of clients, but at the moment most firms are only able to adjust a restricted set of parameters on the fly. Modeling any changes outside those parameters has to wait until after trading hours – a delay that could potentially lead to less positive outcomes when it comes to execution and cost that particular firm the client opportunity.

Taking a new approach to algo testing

In the survey, sell-side executives agreed that being able to adjust more parameters of their algos – having the confidence to be able to do that – on the fly would give them a significant advantage when it comes to improving execution quality. "This would allow us to test more quickly and ensure that our algos are performing without having to do as much running by," says one executive. "It would allow us to make changes more efficiently, more safely across our algo base."

Although it was difficult for respondents to estimate a financial benefit, a few were able to draw on their professional experience to do this. These respondents estimated that the financial benefit of such an approach could translate into substantial financial advantages – 0.05% to 1% improvement in the bottom line for a 1% improvement in fill rates. It's worth noting that some market leaders pour millions of dollars of investments in low-latency infrastructures to gain similar results in terms of improved fill rates.

It's worth noting that whilst there are already methods in place for algo optimisation for multiple parameters, they fail to link to or model market microstructure, making response to change slow and superficial. Their reliance on historical data to derive a market model makes the existing methods unfit for innovative approaches that want to look to the future.

7.0

Obstacles to optimizing algo performance

Survey respondents said they faced several challenges as they seek to ensure their algos succeed in attracting and retaining buy-side liquidity, allowing them ultimately to move up the broker algo wheel.

Any approach necessarily involves meeting the broker's need to demonstrate best execution, as required by regulations like MiFID II, and to adhere to business rules around transaction cost analysis (TCA). Algos also need to meet clients' market impact sensitivities, with sufficient transparency to allow buy-side customers to understand the impact of any given trading model.

More broadly, sell-side firms need to ensure their algos conform to MiFID II's RTS 6 testing specification, which aims to ensure that brokers' algorithms don't behave in a manner that can create disorderly trading conditions. The spec has been framed to ensure firms are able to detect and contain any potential behaviors that could spark a repetition of the 2010 Flash Crash.

Sell-side firms need to ensure their algos conform to MiFID II's RTS 6 testing specification, which aims to ensure that brokers' algorithms don't behave in a manner that can create disorderly trading conditions.

Designing a testing environment that allows firms to meet these requirements is challenging. Survey respondents outlined a few key pain points that preclude their existing set-ups from delivering what they need.

High on their list of issues was the difficulty in modeling their trading strategies for unusual market conditions. Respondents said it was impossible - using their existing processes - to model for every market scenario.

They complained that, due to their reliance on historical data, they were unable to factor in 'unknown unknowns', unexpected events thrown up by market circumstances that elude current forecasting capabilities. This also applied to the scope of the data they were using in their simulations, which in turn often limits the scope of events the simulation is capable of forecasting.

Respondents also reported difficulties in modeling for new regulatory and structural challenges. Keeping their simulation models up to date with emerging regulations is a challenge. This issue also translated into difficulty in developing algorithms rapidly enough to respond to a client inquiry or need.

From the survey, respondents indicated that better simulation capability would translate into more rigorous and expansive scenario testing. This points to the need for a highly granular, dynamic representation of a market that can help sell-side firms optimise the performance of their execution systems and algorithms in order to attract and retain buy-side clients. This can be achieved using agent-based simulation (ABS), which uses models that capture the structure and behavior of complex adaptive systems like financial markets down to each individual agent, like an exchange or a broker algorithm.

Why Agent-Based Simulation?

Agent-based modeling (ABM) works by modeling micro-behaviors and interactions between different agents that lead to complex, macro-level outcomes. By understanding the interactions of behaviors of individual agents within system, users can gain a far deeper insight into the causes of macro outcomes.

Using agent-based simulation, firms can circumvent the issues presented by traditional approaches to backtesting, allowing them to train their algos for an unlimited range of trading regimes. While backtesting on real data commonly results in bias and overfitting, synthetic processes like ABM and stochastics can reduce this bias, allowing the firm to focus on feedback dynamics, which aren't supportable using simple historical time-series or standard stochastic processes.

By running thousands of simulations for every positive regime, the optimal decision path can be chosen by the algorithm every time. This can be run on the fly during the trading day, ensuring algos are optimised for all trading situations.

An agent-based simulation framework for algo testing allows sell-side firms to represent:

- The different trading strategies of market participants such as high frequency traders (HFT), investment managers or competing executions desks.

- The micro-structure of exchanges including the execution policies, order types, speed bumps or auction rules.

- The latency and locations of different exchanges or execution venues.

- The bank's own internal infrastructure, including smart-order routers, circuit breakers and trade processing critical to success.

Like an 'algo-gym', agent-based simulations can help ensure execution algorithms are ready to be deployed in as wide a range of potential futures as possible. Rather than training algos by looking purely at the 'surface' data - the 'real' historical data produced by exchanges and other trading venues - agent-based simulations can re-create the data-generating process, producing synthetic market data that is indistinguishable from a given real-world time-series.

Like an 'algo-gym', agent-based simulations can help ensure execution algorithms are ready to be deployed in as wide a range of potential futures as possible.

Having recreated the underlying system, these kinds of tools have the advantage of parameter adjustment, allowing users to construct any potential scenario they choose. By introducing multiple potential market dynamics, sell-side brokers can ensure their algos function under many types of stressed scenarios and are performant in the broadest range of eventualities. As a result, firms can:

- Explore extreme market events - These simulations are stochastic in nature and running them multiple times will generate a range of outcomes. By running the same simulation thousands of times, team can produce tail events, stressed conditions and market crashes. Specific scenarios can also be run in order to produce large amounts of data. For the first time, algos can be optimised to run in environments that rarely occur, such as in miniflash crashes.

- Accurately model market impact – Uniquely, agent-based simulation allows banks to run the same simulation thousands of times with or without an order. In comparing the results of the two categories of simulation it is possible to more accurately quantify the likely impact of a trade. This same process can be run hundreds of times with different execution strategies to determine which approach is the optimum for every individual execution order.

- Create scenarios with no historical precedent – With simulation, firms can create synthetic data that is calibrated and optimized against real data. Once the simulation faithfully creates market data, the underlying parameters are tweaked to produce novel data.

- Offer more transparency to clients and regulators – Agent-based simulation allows banks to demonstrate to regulators that they have tested their algos in a broad range of environments and prove that there is no way their algos can contribute to disorderly market conditions.

9.0

About Simudyne

The need for a new approach to algo testing is clear. The Simudyne team are the market leaders in agent-based financial market simulation for best execution. Partnering with Simudyne ensures your business is using the most advanced technology supported by the best advice that the world has to offer.

Simudyne's platform for optimization and back-testing process can deliver immediate results:

• Differentiate your market execution strategies and attract order flow by marketing a genuinely innovative approach.

• Move up the rankings on the broker wheel by optimizing your execution strategies and deploying them in a wider range of market environments.

• Test in a wide range of environments to create more robust algorithms.

Improve benchmark performance with highly optimised algorithms

Simudyne is a simulation technology company that offers organisations a new way to more effectively harness the power of agent-based modeling, AI and machine learning to test drive their decisions and drive growth. Simudyne's enterprise ready software is currently used by large financial institutions to quickly and efficiently simulate an unlimited number of future scenarios and measure their impact in a safe, virtual environment.

Contact

A:	St Michael's Alley, London, EC3V 9DS
E:	support@simudyne.com
W:	simudyne.com
TW:	twitter.com/simudyne
LI:	linkedin.com/company/simudyne

Simudyne is a rapidly growing technology business,
harnessing the power of advanced simulation,
to help organisations make radically better decisions.
Our efficient and scalable simulation platform allows
enterprises to create a virtual environment where they can
test drive their decisions, fail fast without consequences
and create solutions that drive growth.

Simudyne ———— Reality, re–engineered