

Not what, but who

A WFE white paper on market-data use cases

May 2021



Not what, but who

Theory, practice and policy in the pricing of equity-market transaction data – a WFE¹ white paper

Introduction

There has been a small but very promising improvement in discussions about the market in equity transaction data. It comes from looking at the reality on the ground, rather than trying to make the reality ‘fit’ a theoretical template.

More specifically, the key breakthrough has been the first recognition of the huge difference that exists between the various uses to which exchanges’ data is put, reflecting major differences between types of user.² This is particularly true of low-latency data that tracks what is happening in real-time, capturing data about the state of demand to buy or sell shares as it evolves.

The distinction between use types may not sound that significant. This white paper addresses the question of why it is. It also maps out areas that would be fruitful for further research, and which are currently ignored, especially by those campaigning for regulatory action that is not in fact justified.

By way of background, the analysis in this paper builds on a fundamental point about the creation of data by exchanges: that it is an active and specialist role that requires expertise, care and thoroughness, underpinned by investment in *all* the functions that make an exchange a trusted and innovative venue for trading and reliable source as to the price of listed assets.³ In turn, when that data is made available to various types of participant – all the way through to end-investors – on terms that properly reflect their respective interests, the arrangement feeds back into liquidity in a virtuous cycle. This is to the advantage of all participants (and – provided the benefits do not instead accrue to venues that use the exchanges’ data to increase their *own* volumes – of the exchange itself). As a practical consequence of this, exchanges act as an originator of what are in effect multiple types of data product (or versions thereof), including a wholesale, low-latency product that costs more to produce and which only some participants use, as well as the simpler products that the broader market prefers. The product/version targeted at wholesale users (especially ‘alternative’ venues) is used in trading, which brings in to dealers multiples of the revenue that exchanges earn from data. (*See page 5 for details.*)

Accordingly, this paper provides a framework in which it is possible to think in a proportionate manner about specific, individual approaches to licensing market data; and to facilitate constructive discussions among all stakeholders. Given the importance of public markets to social welfare, understanding why the market in data works the way it does clearly matters.

The structure of this paper

We have structured this paper as follows. We begin with an introduction to the concept of a variety of use cases (participants) in the market place and therefore in relation to equity market data. (The type of data we are focusing on is the sequence of share prices and possible future prices.) We explain the impartial role of exchanges, not just as creators of data but also of the whole microstructure, within which participants operate and data has value. Responsibility for an effective microstructure entails being prepared constantly to innovate.

Further, we note that exchanges distinguish carefully *between* use cases, such that each user type pays a rate that is fair, given the type of price-data product they license. We contrast the impartial exchange role – which is consistent

¹ The [World Federation of Exchanges](#) (WFE) represents members operating over 250 exchanges, responsible for clear and honest information about the value of \$100 trillion of assets.

² We refer to the recent IOSCO [consultation](#), on which more below.

³ [‘Market data pricing’, July 2019](#)

with the promotion of social welfare – with the rivalrous competition between market participants to make money from trading (noting that some of these participants may also be competing with exchanges for execution business). As we note, exchanges make modest amounts from data, when compared with the amounts made by traders using data as their key input.

We build on this analysis by looking more closely at the time-frame within which various consumers of data find it profitable to do so. Thus, market-makers use instantaneous data in the business of facilitating trading by others; whereas at the other extreme, an academic conducting research may come back to the information years later; while TCA or algorithm development, for example, though they may involve equally extensive and granular data as a market maker, can be a profitable activity without needing to consume data the very microsecond it comes into existence.

We also talk about the fact that – with the exception of aggregator-vendors who charge a mark-up on the data that exchanges create – market traders license data primarily to inform trading strategies rather than to sell on to others. In such a business model, more data means a more complete picture, making each data point more valuable than the last.

By way of conclusion, we look at the implications for the pricing of data – observing that there is no market failure but rather the opposite: a very effective form of capital market infrastructure. If access to data was in the control of any one set of participants, then outcomes would inevitably be skewed, to the detriment of that overall marketplace. (Markets other than lit equity suffer from just this problem.) Exchanges are in fact unique in being aligned with the interests of the marketplace as a whole, with competition authorities empowered to intervene if they should ever deviate from this.

Comments and questions on this paper may be sent to Richard Metcalfe, Head of Regulatory Affairs, World Federation of Exchange: rmetcalfe@world-exchanges.org

The usefulness of use cases

That there are distinct and disparate use cases for trade data should not really come as a surprise to anyone. It is well known that ‘the marketplace’ is not homogenous in terms of participant types and that the various types of user may in fact have divergent interests. But somehow, deliberately or otherwise, when it comes to discussions of the market for market data, this message has been lost. In any case, it is key to understanding that market.

The term ‘use cases’ refers to the way that demand for data – and even one’s perspective on what data is relevant – can and does differ significantly by type of market participant and how they exploit it: retail/wholesale; intermediary/end-customer; trader/risk manager; day trader/buy-and-hold. For some, the only data that matters is transactional price-volume information that is timely; for others, an end-of-day mark. But, even when the focus is purely on transactional data, it matters in different ways to different users.

This insight into use cases, of which we give more examples later, opens up the real, underlying issue, because the role of data is interwoven with the way those types of participant interact with each other. This is consistent with the fact that it is impossible to consider trading and data separately. (Who [apart from vendors who on-sell] would be interested in licensing trading data *without* being interested in the practice of trading and closely related disciplines such as risk management? QED: they are a joint product of the exchange, rather than one being a by-product of the other.⁴)

Each exchange takes all these important intra-market subtleties into account in organising its ‘microstructure’ – essentially the terms on which price formation happens via fair and transparent trading on that exchange; and the data is published for all interested parties. It does this in a way that aims to ensure the sustainability of that price-formation process, for commercial use by wholesale traders⁵ but with that wider, end-user marketplace in mind as well; and as a result of all that (and *only* as a result of it), creating returns for the exchange itself. As is generally acknowledged, these activities in turn serve the purpose of increasing social welfare.

By contrast with the role and objectives of the exchange, *individual* participants very obviously have much narrower interests, which are ‘rivalrous’ vis-à-vis each other with regards to trading profitably (which data helps them to do very efficiently). Put simply, they compete – with other participants of the same type (eg, market makers competing for order flow) and, such being the nature of traded markets, potentially with other types of participants. (Clearly, venues do compete with each other but not with its customers, unless they also happen to be running alternative venues.⁶)

There are consequences that flow from this rivalry between market participants and we should recognise them. One is the way that such rivalry inevitably gives rise to increased demand for data, because of the commercial advantage it can give, especially when used on an industrial scale to help monetise bid-offer spread and profit from price trends. And with more than 30 *billion* microseconds in a typical trading day, there is no shortage of opportunities to make a turn using that data, if one uses more of it more quickly than others do or even can.

⁴ Ruben Lee, ‘What is an Exchange?’, Oxford Press, 1998

⁵ “He deftly explains that [the] characteristics of trading on exchanges create opportunities for specialists on the exchange floor.” Hester Peirce, SEC Commissioner (‘Alligators in Nirvana’, May 2019; paraphrasing Harold Demsetz in ‘The Cost of Transacting’ [Quarterly Journal of Economics, 1968])

⁶ Dealers like to act as non-transparent, quasi venues, eg, in fixed income. A complicating factor in the *equity* data debate is that, in some jurisdictions, exchanges face competition from technologically advanced trading firms that are permitted to use exchanges’ data to run venues, but without the same transparency and not just in large trades but increasingly in smaller ones that could easily be digested in lit public markets.

(To visualise in more intuitive terms how fine the time distinctions are within the slickest trading operations, recall that a microsecond is one millionth of one second – a distinction that is only really meaningful to a computer and way faster than one can blink an eye.⁷)

A second, rather more disturbing consequence of failing to examine where rivalrous interests exist is that one ends up with a grossly misleading picture as to the *exchange's* role and interests. The biggest market participants appear to be claiming that 1) it is they who are on the side of the end-investor; and 2) that all participants jointly have an adversarial relationship with the exchange. As per the explanation above of the exchange's role vis-à-vis the customer, such a characterisation is at odds with the truth. Exchanges positively *want* market participants to make use of the data they manufacture. (Venues do compete, but with each *other*, for transaction flow, rather than with their customers who take positions, through market-making or otherwise.)

Why does this matter? Because it confuses the issue and, more ominously, threatens to take us back to a world in which the conflicts inherent in the intermediary-customer relationship (on data and more widely) are not even acknowledged, let alone addressed properly. It is not so long ago, in market and regulatory evolution terms, that end-user access to markets and to (equity-market) data was in effect controlled by a small number of dealer firms who could dictate who saw what and when. This situation still prevails in fixed income and other asset classes that are predominantly traded OTC.

As mentioned, the crux is that differences exist among types of capital-markets participants, relative to *each other*. This important distinction, particularly in relation to transaction data, became a bit clearer in the consultation on equity market data launched by IOSCO in December 2020. That consultation explicitly acknowledges the range of participant types and asks who views what data as relevant, clearly inviting the comment that the utility of the data – and the use to which it is put – must depend on the nature of the entity using it. One answer is that what is relevant to a market maker is not and never will be core to someone conducting, say, end-of-day portfolio valuation.

Data fit for purposes

This brings us to the use cases we alluded to above. A quick glance at some examples reveals two things. The first is that there are significant differences in how quickly some participants want (or, as they say, 'need') to consume data, as compared to other data users. So while there are data users whose demand is for real-time data (eg, for trading), others want end-of-day data, eg, for valuation purposes. Within this range, the perhaps more significant distinction is between those who can and want to consume data most rapidly, pre-trade and, on the other side of the divide, the rest of the market.

The table below divides types of data users, in rough order of the time frame within which they typically find it profitable to consume data, that time frame running from the point at which the data is created to the point when the licensee has a commercial interest in seeing it.

⁷ Put another way, if microseconds were full seconds, then it would take 11.5 *days* to run up a million of them.

<u>Use case</u>	<u>Utility time frame</u>
Market-making / Proprietary trading	Pre-trade (real-time, sub-second – streaming full order book; premium information)
Dark pool / Some competing venues, eg, Internalisers	
Agency broking	
Day trading	Post-trade: 'real time' (from sub-second up to 15 minutes [delayed data])
EMS provision	
Risk Management	
Surveillance	
Index provision	
Media outlet	Post-trade: between seconds and hours (up to end-of-day)
Issuers	
Transaction-cost analysis (TCA)	
Value traders	
Portfolio / position valuation	Post-trade: end-of day
Academic	Post-trade: historic
Trading strategy / algo development	

As mentioned, the biggest difference is between those monetising bid-offer spread (which is an integral part of how market makers earn a living and serve the broader market) and those who are not. To be absolutely clear, we have no objection to anyone carrying out the former function, which is useful – if not essential – for the effective operation of markets and exchanges. But equally, it involves a different level of data demand and usage.

(The list in the table is not exhaustive but does cover major types of user. Please note also that we have not attempted to quantify how *much* data any one type of participant has a commercial interest in consuming, nor how granularity affects its value.⁸)

The current WFE paper is setting out this detail, against the backdrop of how exchanges manufacture data (which we have already explained in some detail, in previous [statements](#)). We also wanted to get at how this data-manufacturing process supports the exchange mission to run markets that are accessible to the public, and specifically how the exchange does so while taking into account the important, well documented, real-world differences in how market participants trade with each other.

Because it is impossible to understand the interests of various participants without looking carefully at how transaction data is used within each bit of the chain – rather than treating it incorrectly as an *un*-differentiated product that is the same to all users at all times (and for all purposes). That will only lead to even greater distortions and power imbalance than already exist, particularly in the current wholesale-retail interface.

⁸ By way of additional colour, one WFE member has analysed the US market and described a fuller range of users and why they use exchange data. See Chart 3 in: nasdaq.com/articles-21-02-25

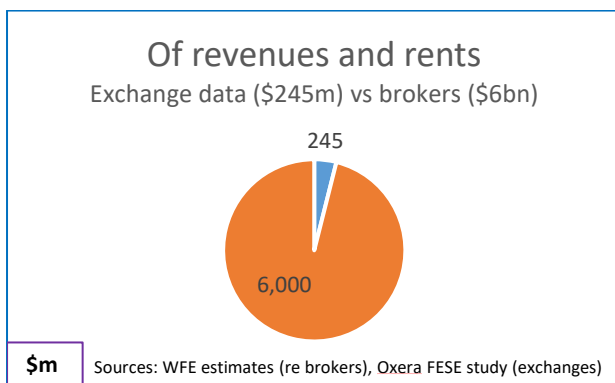
The ‘peculiar’ economics of data

To illustrate an important characteristic of the data product, contrast the market in ‘fast’ equity transaction data with that in goods – TVs, for example. Such goods will, in all but exceptional circumstances, follow what we might call a wholesale-retail ‘cascade’, with i) wholesalers buying from manufacturers and selling the product at wholesale price ‘x’ to retail outlets, who ii) sell to the ‘real’ consumer at the retail price. Yet in financial markets, wholesale traders may not in practice be sharing data with customers at all. And, even if they do, first and foremost they use it to fuel their trading business.

What follows from this is that the economics of the market for market data must be considered according to that market’s own merits. This is quite different from the paradigmatic ‘widget’ which is distributed in a wholesale-retail ‘cascade’, ultimately reaching the end consumer. Data very clearly does not work like that. Intermediaries license authoritative, low-latency data to inform their decisions as to the price at which to sell or buy securities, in a business process in which data powers largely algorithmic trading strategies. While this may bring broad benefits, adding to liquidity in the marketplace as a whole, the immediate monetary value in this low-latency data is captured by the trading intermediary earning revenues.

Note also that, especially for the largest market participants, execution ‘costs’ in the form of exchange fees may in fact be zero or occasionally even less; while they make the most from execution, supercharged by their ability to exploit large quantities of data extremely rapidly. In other words, the perspective of a) participants and b) exchanges do not mirror each other. They are complementary but very distinct.

Consider the following pie chart, illustrating how much EU exchanges make from licensing data, as compared with how much brokers make from trading.



Exchanges do make some money from data. But the real income is made elsewhere, by trading firms. Revenues – even for the EMEA business of just a half-dozen broker-dealers – dwarf the revenues that EU exchanges make from data.⁹ From a global perspective, as the WFE noted in its February 2021 response to IOSCO, “The 10 major banks involved in global equity markets generated about \$47.6 billion in revenue in these markets in 2020.¹⁰ This stands in contrast to around \$2.2bn earned by the ten largest global exchange groups’ market data businesses [more

⁹ “Philippon’s striking finding is that the unit cost of financial services has barely changed over the past century.” Andy Haldane, Bank of England, November 2020 ‘[Seizing the opportunities from digital finance](#)’

¹⁰ WFE estimates, based on quarterly reports including Bank of America, Morgan Stanley, JP Morgan, Goldman Sachs, Citigroup, UBS, Barclays, Credit Suisse, Société Générale and BNP Paribas.

broadly¹¹] over the same period. These exchange revenues, annualised, furthermore represent just 6.8% of the revenues of the global market data industry.¹² Meanwhile, the proportion of investor fund management expenses attributable to [EU] buy-side market data spend is less than 0.001%.¹³”

Trading intermediaries (and competing dark venues) take the exchange’s ‘data-picture’ to inform spreads and positioning / warehousing. This can help facilitate trading activity and (depending on the degree of power that the intermediary has to decide where to route the trade) price formation. Others on the customer side then update their views on investments. Ultimately, data changes from a live, ‘in-play’ opportunity, to historical information. Please note that this is not to deny the benefit of what, to a market maker, may now be stale data. It can still benefit various parties in distinct ways: for example, to give the retail investor a sense of how the market is evolving; or to allow risk managers to devise a model. But let’s not get distracted or delude ourselves about the difference between

- i) the massive, heavily industrialised processes exploited (quite legitimately but also quite profitably) by those participants whose business model relies on the use of low-latency data (eg, market makers and HFTs) and who oil the wheels of traded markets by targeting the flow of micro-opportunities arising from the constant changing of The Price for share X; and
- ii) what a buy-and-hold, dividend-accumulating investor does when she buys on the dips and for the long term.¹⁴

It is an even more unusual property of data that, other things being equal, the more any technologically enabled wholesale market participant has of it, the more valuable it will tend to be to that participant. As far as we are aware, this property is completely ignored in the literature. (Certainly, you do not find any market makers drawing attention to the fact.)

To understand this characteristic of transactional equity data, especially the pre-trade variety, it is instructive to look again at the contrast with the economics of tangible goods. If I am a wholesaler with 10 tons of potatoes to shift, for well understood reasons, I will typically charge a *lower* per-unit price to a supermarket than the supermarket itself does to its retail customers, a kilo at a time.

With data, however, depending on which use case one falls into, the more of it one can license from an authoritative source (ie, an exchange), the *higher* the value per unit may be, because one has a fuller picture. (There is nothing inherently wrong with this, when one considers the structure of the market ecosystem, with market makers serving a wider user base.) As a market maker, if I can see the likely intentions of 1,000 participants, I can then judge much better how to pitch my bid-offer than if I only see, say, 12. In fact, the intuition – hard to quantify but following inevitably from the initial insight regarding the relative of value of one data point compared with one million – is that the value to a given participant of a certain type (ie, wholesale and technologically advanced) increases by more, for each additional data point they can see. The increase in value – albeit only for this user type – will be exponential.

Equal opportunity to access data means that anyone can build a model on this kind of data usage. It is just that in practice not all feel any incentive to do so, meaning that data is not used the same way across the ecosystem.

¹¹ WFE estimates, based on quarterly reports including: Cboe Global Markets, Deutsche Börse, Euronext, Hong Kong Exchanges and Clearing, Japan Stock Exchange, London Stock Exchange, Nasdaq, NYSE, SIX and TMX. These figures probably overstate the reality, given that some exchanges include a) asset classes beyond equity and b) information services beyond the scope of streaming market data. US exchange revenues include those from the SIP.

¹² Burton-Taylor, Financial Market Data/Analysis Global Share & Segment Sizing, 2020.

¹³ Oxera, Pricing of Market Data Services, An Economic Analysis, February 2014.

¹⁴ At time of writing, the US is scrutinising the practice of paying for order flow (PFOF). The existence of PFOF illustrates the value to some firms of seeing how other market participants are acting. (Those other participants may well undervalue the information they are giving to wholesale market participants.)

Again, the literature generally remains bogged down in other models of the world, into which it tries to force the data market. And again, some market participants appear to have no appetite to think otherwise – presumably because they have no incentive to do so.

Delivery costs, personal data and other

One way in which data is no different from other markets and where distribution/on-licensing does apply when vendor-aggregators step in. We will not delve into how much this convenience is worth but the economic point is clear and striking. Considering data procurement costs of large-scale consumers, the 2019 Oxera study suggests that fees levied by exchanges in the EU do not even represent 15% of the total. Unlike exchanges, other stakeholders in the provision of data services, do not have any transparent pricing scheme. They account by far for the largest portion of data-procurement costs. Seen in this context, exchanges are obviously the wrong target, conspicuous only because they have a transparency obligation.

At the same time, it is inevitably true that if usage of data goes up, then the aggregate spend will. A severe and chronic lack of public evidence exists in relation to the ‘increased cost of data’ claimed by many in the market.¹⁵

Incidentally, one can dismiss the idea that market data is akin to personal data of the sort appropriated and monetised by social-media companies when it really does or should belong to the individual. Market data is self-evidently *not* data about a person, her identity or characteristics, about her lifestyle choices; or indeed about anything except the overall market in shares. Attempts to compare the two show a serious failure to grasp the nature of *either* type of data – or possibly just a desire to misrepresent. In any case, equity transaction data is about a fair and transparent price-formation process delivering a picture that emerges from multiple interactions by diverse parties – not *about* those parties per se or whether they prefer a Volvo to a Cadillac. Market data is more like a combination of pixels that, only when combined by the exchange, create a picture¹⁶. Hence this perceptive comment by Charles Jones of Columbia Business School: “Data consumers buy... aggregated data not to view their own orders and trades but rather to see the overall state of the orders and trades in a market.”¹⁷

This is why an exchange’s data is also *un*-like certain online advertising boards. The likes of e-bay, effective though they may be at what they do, do not create a market in the sense that the WFE members understand the term. And they do not create a (let alone ‘the’) market price, because that would require them to reflect a balance of ever-evolving opinion, interest and activity.

And this is not just true of equity. In FX, where the notion of ‘the’ impartial yet authoritative price has historically been missing, there has been real value in integrating the information from a highly fragmented process, effectively performing the role that, in equity markets, exchanges perform. (To fill the information gap, a commercial data solution, which apparently prices according to demand, has emerged.¹⁸) But even in FX, there is a live question as to which venues deliver most value, in terms of the information they can generate. And we know what determines that value. Some capture more trading business (at least in some currency pairs) and their data has more value accordingly.¹⁹

¹⁵ The Oxera study cited above (in the pie chart) shows that unit costs for data have risen only very modestly.

¹⁶ A moving picture, to be precise.

¹⁷ [Understanding the Market for US Equity Market Data](#). Jones is Lear Professor of Finance. The work had financial support from an exchange. It also includes figures to back up its assertions, aka evidence. The WFE believes that anyone is free to sell equity market data, if they believe it to be of value.

¹⁸ New Change FX. While this is not an execution platform, it does inform execution, TCA and so on.

¹⁹ “ ‘The legacy venues can’t afford their volumes to drop much more because their market data will lose its value,’ says David Mercer, chief executive of LMAX Exchange, a rival platform.” [FT](#), 1st December, 2020.

So, what does this all add up to, in terms of how data will actually typically be priced?

The implications for pricing

Exchanges should price data in the purest way possible – according to demand. This is why the ‘hungeriest’ users of data should pay more. They consume not just larger quantities but, ideally, faster and with more granularity, as part of a lucrative business model. The literature supports this, noting the social welfare gains that arise when a) pricing is done according to price sensitivity of various customer groups; while at the same time b) ensuring broad availability of data and c) financing ongoing investment by exchanges to ensure the relevance and resilience of their offering, including market-data distribution.²⁰

At the same time, exchanges’ approach is entirely consistent with their incentives to promote the functioning of the equity marketplace overall. That in turn ensures that their own business is sustainable, including any future investment needs.²¹

As regards the financing of continuous investment, what matters is not just historic or running costs but the ability to remain competitive in future. The WFE believes it is important not to forget the challenging business landscape for exchanges, which:

- i) has persisted for decades now, through a mix of competitive forces, regulatory change and technological development;
- ii) shows no sign of easing and can reasonably be expected to remain challenging;
- iii) has increasingly accommodated business that run *off-exchange* venues, while piggy-backing on what exchanges do²².

The exchange *is* data – and it is clear that exchanges charge reasonable amounts for what is in practice a premium product, because it is the ultimate reference (‘gold standard’) in terms of timely availability and reliability.

Even when one takes into account the business of running the exchange, one is still left with a pivotal question, which is what the fair price is to a particular type of user. Certainly, there should not be one price for all. Set a single price and one creates (hidden) subsidies. Some could pay more than is fair – much more, because they will never be able to do 15bn round trips in a day, the way a computer algo can.

A retail investor currently pays a very low fee for consuming real-time market data. Setting a single price, irrespective of the usage type, would make access to real-time market data unaffordable to retail investors and amount to them subsidising those that use data to the maximum. (It could also interfere with exchanges’ ability to innovate in a way that is sensitive to the differing demands of different user types.)

Moreover, because the subsidies are (or would be) hidden, one achieves exactly the opposite effect to what users of data claim – to make the market in that data function better. A single price distorts the market. Market makers sometimes argue in favour of a single price, as though data was an undifferentiated product – the same at the point of delivery, no matter when one uses it or in what quantity. As we have seen from the discussion of the nature of data, such market-maker arguments do a disservice to open and honest debate.

²⁰ Lee, ‘What is an exchange’; and Carl Shapiro & Hal Varian, ‘Information Rules’, Harvard Business Press, 1998

²¹ Much is made some exchanges operating for-profit. It is hard to think of a factoid *less* relevant to correctly pricing market data. Any exchange will consider whom to charge for what, to promote an active market.

²² The outcome in the EU is that the second largest execution venue today is a ‘Systematic Internaliser’. See ESMA [Annual Statistical Report](#), November 2020.

As a principle, a market that has demonstrated no ‘market failure’ is one in which data is priced fairly, in line with the use(s) to which the data is put – and when and in what quantity. It should be viewed in the context of the overall market for equity trading, for which data is a valuable input and which itself functions more than adequately (at least on lit markets – we cannot say the same for markets where excessive fragmentation means wholesale participants can make even more money from data).

The WFE believes it is better – in fact necessary – to look at data in the context of competitiveness across the equity trading business as whole. And in extremis, competition authorities do have the powers to assess this and the actions of any individual exchange. But this ought by rights to consider data as part of an overall offering.

Closing thought

This paper makes a modest contribution to a long-running but remarkably poorly framed debate. The consideration of use cases promises a way forward, because it at last looks at data ‘in the wild’, as it relates to actual practice in equity trading, rather than according to theoretical dogma.

In the end, just because some market participants find utility (benefit) in the data, it does not make data provision a utility *service* in the sense of public property. The market continues to function – better than ever, in terms of overall cost to end-customers and it is exchanges that are there to make sure that continues to be the case.

Therefore, before any discussion about interfering with the pricing of exchange equity-market data, it is important to check whether there are truly any market failures or barriers to entry for venues that want to compete. As far as we can see this, this is not currently the case.

Exchanges perform a unique role, balancing the interests of all that participate as well as having regard to the exchange’s own viability as a business operating in a competitive and rapidly evolving environment. Data is much more than a side issue in this. As we have tried to make clear in this paper, the exchange’s role in ensuring that the right data reaches the right people, on the right terms and at the right time and price, is crucial to the capital formation in which public markets perform a central role. Included in this is an understanding of how consumption of data by some wholesale participants can profitably facilitate services to the broader market. Regulation (for instance, price controls) where there is no market failure and where a delicate balance of interests exists will harm the dynamics of innovation from which the marketplace as a whole benefits.