

Fragmentation, Price Formation, and Cross-Impact in Bitcoin Markets

Jakob Albers*

Department of Statistics, University of Oxford
Oxford, UK
jakob.albers@merton.ox.ac.uk

Mihai Cucuringu

Department of Statistics and Mathematical Institute, University of Oxford, Oxford, UK
The Alan Turing Institute, London, UK
mihai.cucuringu@stats.ox.ac.uk

Sam Howison

Mathematical Institute, University of Oxford
Oxford, UK
howison@maths.ox.ac.uk

Alexander Y. Shestopaloff

School of Mathematical Sciences, Queen Mary University of London
London, UK
a.shestopaloff@qmul.ac.uk

August 21, 2021

Abstract

In light of micro-scale inefficiencies induced by the high degree of fragmentation of the Bitcoin trading landscape, we utilize a granular data set comprised of orderbook and trades data from the most liquid Bitcoin markets, in order to understand the price formation process at sub-1 second time scales. To achieve this goal, we construct a set of features that encapsulate relevant microstructural information over short lookback windows. These features are subsequently leveraged first to generate a leader-lagger network that quantifies how markets impact one another, and then to train linear models capable of explaining between 10% and 37% of total variation in 500ms future returns (depending on which market is the prediction target). The results are then compared with those of various PnL calculations that take trading realities, such as transaction costs, into account. The PnL calculations are based on natural *taker* strategies (meaning they employ market orders) that we associate to each model. Our findings emphasize the role of a market's fee regime in determining its propensity to being a leader or a lagger, as well as the profitability of our taker strategy. Taking our analysis further, we also derive a natural *maker* strategy (i.e., one that uses only passive limit orders), which, due to the difficulties associated with backtesting maker strategies, we test in a real-world live trading experiment, in which we turned over 1.5 million USD in notional volume. Lending additional confidence to our models, and by extension to the features they are based on, the results indicate a significant improvement over a naive benchmark strategy, which we also deploy in a live trading environment with real capital, for the sake of comparison.

*Denotes first author; the rest of the authors are listed in alphabetical order.

Contents

1. Introduction	2
2. Background	4
2.1 Historical Developments and Market Specifications	4
2.2 Data and Preprocessing	7
2.3 Remarks on Execution Cost	8
3. Features and Preprocessing	9
3.1 Base Features	9
3.1.1 Orderbook Imbalances	10
3.1.2 Trade Imbalances	12
3.1.3 Past returns	13
3.1.4 Mean Divergence	14
3.2 Eliminating Nonlinearities	15
3.3 Optimal Time Horizon Selection	19
4. Network Effects in Bitcoin Markets	21
4.1 Leader-Lagger Network	21
4.2 Accounting for Trading Realities	23
5. Cross-Sectionally Combining Information	28
5.1 Baseline Approach	28
5.2 LASSO Regression Approach	30
5.3 Meta Features Approach	34
5.4 Comparisons of Models	37
6. Market Making Experiments	40
6.1 Strategy Specification	40
6.2 Methodology and Experimental Setup	45
6.3 Results and Comparison with Benchmark	47
7. Conclusion	49

1 Introduction

Cryptocurrency markets have seen an explosion of trade volumes over the past year, as the price of Bitcoin soared to an all time high of over 60,000 USD in April 2021, sparking a great deal of interest both from the broad public and academics alike. The trading landscape for Bitcoin in particular has matured considerably over recent years, with an ever-greater proportion of trade volume occurring in complex derivatives rather than spot (fiat for Bitcoin) markets. With the rise of derivatives volumes came an accompanying flurry of academic works that investigate their role in the price formation process. For instance, [14, 15] examine the price impact of a set of popular unregulated derivatives exchanges, [24] performs a similar analysis, and [13] surveys market microstructural differences between spot and futures markets. Further work such as [29] provides insights on the way in which the proliferation of derivatives has exacerbated price jumps and thus contributed to greater volatility.

One of the main distinguishing elements of Bitcoin markets compared to traditional markets for stocks or bonds, aside from the much greater volatility for the former, is the considerable degree of *fragmentation* inherent in the Bitcoin trading landscape. This fragmentation stems from the existence of 5-10 highly liquid exchanges where the majority of trade volume takes place. Within each of these exchanges, the volume is typically further distributed across a number of different liquid instruments with slightly different specifications and properties. Exchanges operate independently from one another. They are subject to different regulatory environments depending on where they base themselves out of. They often use different server locations, resulting in cross-sectional arbitrages that can persist for at least as long as the time of information transfer between two distinct server locations. There is no cross-collateralization across exchanges, unlike in traditional markets where prime brokers play this role. That is, if a trader enters a long position on one exchange and an equally large short position on another exchange, he is market neutral but must still ensure that both positions are sufficiently collateralized in order to avoid a liquidation event on either exchange.

Summary of main contributions. The above factors, and others discussed in subsequent sections, altogether contribute to the aforementioned fragmentation of liquidity across exchanges. This type of environment raises a number of interesting questions one can pose. The main contributions of our work is to identify possible answers to the following questions.

1. Where does new price information tend to arrive first? In other words, which venue is most often the originator of price transmission?
2. What are the predominant directions in information flow?
3. In what way exactly is price on one venue affected by the arrival of new information on another market?
4. To what extent do time discrepancies in price impact enable price prediction on each venue? That is, can we leverage the cross-section of information from all markets to make reliable price predictions?
5. Can these predictions be leveraged to produce trading strategies in a natural way, which give large PnL values? Do we obtain predictive power sufficient to produce "alpha" in excess of transaction costs?

Our paper represents an addition to the rapidly growing body of literature on Bitcoin price formation in at least the following two ways. First, it is the first of its kind to comprehensively examine the consequences of the aforementioned fragmentation of the Bitcoin trading ecosystem. Our study encompasses the largest markets in terms of trade volume while, to the best of our knowledge all other studies focus on the impact of either a single market or a limited subset of markets. Often such subsets are made up of illiquid spot markets or for example CME Bitcoin futures markets, which have little bearing on the price due to their lack in trade volume. Second, unlike other studies, we use extremely granular limit orderbook data which allows us to produce predictive models rather than merely explaining contemporaneous returns.

Paper outline. To give the reader a preview of how this is achieved, we first lay out the structure of this paper.

- To begin with, Section 2 provides necessary background knowledge pertaining to the Bitcoin trading ecosystem. This involves descriptions of the predominant markets and their trade volumes, along with some relevant information on their historical development. We will also describe relevant facts on execution cost, before giving the reader an overview of the data sets that underlie the analyses presented in this paper. This includes a description of our method for its procurement and explanations of the processing steps we employed to render the data amenable to subsequent analyses.

- Next, in Section 3 we define a set of microstructural features based on orderbook and trades information that build on notions of order flow imbalance [20] and price differences. The features are designed to quantify and make precise the vague notion of "informational flow" mentioned in the above questions. We additionally undertake a number of preprocessing steps in order to render the features more suitable as inputs to the linear models that will be trained in later sections.

- In Section 4 we begin to reap the rewards of our prior work. The previously defined feature set will be employed to answer the first three questions listed above. That is, we will investigate cross-impact between markets and derive various networks that illuminate the lead-lag relationships present between pairs of markets. We will then examine how a market's leadingness or laggingness translates to trading realities by defining natural trading strategies and performing PnL calculations. This will involve a foray into the crucial role played by transaction cost in determining not only PnL but also a market's proneness to being a leader or a lagger.

- Section 5 is devoted to the last two questions from our above list. We will compare three different methodologies for developing powerful linear predictive models that make use of all available information. These methodologies will be compared along the two axes of explanatory power and PnL.

- In Section 6 we will address a follow-up question that arises from the results of the previous section. The strategies we initially define for Sections 4 and 5 employ only taker orders (market orders or limit orders that lead to immediate execution) and we will find that they yield particularly low PnL values on markets where the taker fee is large. But such markets typically also give large rebates to maker orders (passive limit orders). This motivates the question of whether we can devise a natural maker strategy on the basis of our linear models and whether these strategies yield positive PnLs on markets with large maker rebates. We will define such a strategy and test its performance on one market by means of a live trading experiment since, as it turns out, testing market making strategies on historical data is virtually impossible.

- Finally, in Section 7 we conclude by summarizing our findings and discussing future research direction.

2 Background

This section provides background information on various aspects pertaining to the structure of crypto markets. It is necessary to have a full contextual understanding, in order to be able to appropriately describe results in later sections.

2.1 Historical Developments and Market Specifications

Crypto markets have matured considerably in recent years. This can be seen by the increasing dominance of derivatives relative to spot volumes. Nowadays, average daily volumes on derivatives markets eclipse those on spot markets by a factor of greater than five, whereas up until (and including) the year 2018 spot markets had greater average trade volumes [5, 30]. On a spot market, traders transact the asset itself, while on derivatives markets traders exchange contracts whose value depends in some pre-specified way on the price of an underlying asset. For instance, on a BTC/USD spot market buyers and sellers trade “physical” Bitcoin for USD. On the other hand, on derivative markets whose underlying is BTC/USD, traders can buy or sell contracts whose value derives (hence the name “derivative”) from that of the BTC/USD spot price. Derivatives provide two main advantages for traders. First, they greatly enhance the expressability of traders’ opinions on the future price of, for instance, Bitcoin. Traders that believe the price of Bitcoin will decline in the future can use derivative contracts to profit from the anticipated decline by purchasing a contract whose value rises as the price of Bitcoin declines. The second advantage is that derivative contracts give traders the ability to use leverage, meaning that they can *leverage* a small amount of initial capital to increase their buying power. With the use of leverage comes the additional risk of forced liquidation. Simply put, with leverage, traders can take extra risk for the chance of extra reward, which substantially boosts capital efficiency. For more details on the mechanics of leveraged trading in crypto markets, we refer to [9]. It is generally to be expected that the aforementioned advantages of derivatives, namely improved capital efficiency and greater expressability of views on price development, give rise to a more efficient marketplace.

The easy access to leverage and its accompanying increase in capital efficiency can, however, also bear some disadvantages. This is particularly true for crypto markets, where the amount of leverage that many of the unregulated exchanges allow traders to employ is almost comically large when compared to traditional and regulated financial markets. It is not uncommon for exchanges to allow 100x leverage, meaning that traders of derivative contracts can increase their buying power by a factor of 100 [1, 9]. By comparison, regulated marketplaces of traditional assets such as stocks, bonds, and ETFs typically do not allow more than 5-10x leverage, with some prime brokers allowing leverage up to 20x on some exchanges (as in London and Tokyo Stock Exchanges). With such large leverage, however, comes considerable risk of ruin. For example, a trader who holds a 100x leveraged position will lose his entire initial capital if he experiences an adverse price move of just 1%. A common theme in crypto markets is that of *cascading liquidations*. To illustrate what this means, suppose for instance that a large number of traders hold highly leveraged long positions. This means that a relatively small price move down can cause their positions to be forcibly liquidated by the exchange, resulting in a great deal of unnatural selling. This artificial (and unintentional) selling pressure drives the price down further which might result in more sell liquidations, and so on. Extremely high leverage is akin to gunpowder in the powder keg, and a small initial price move acts like the match to light the keg, triggering a cascade of liquidations which greatly amplify the initial small price drop to an unnaturally large one. The historical price development of Bitcoin is littered with interesting case studies of this dynamic playing out in real life. For more details, we refer the reader to [29].

Our selection of markets. When we use the term “*market*”, we are referring to a pair of *exchange* and *symbol*, where a symbol on a given exchange denotes either a spot market (for instance BTC/USD) or a derivative contract (for example a quarterly futures contract).

New information marks its arrival on the markets as a whole in the form of traded volume. That is, when a trader possesses knowledge which she/he believes is not yet reflected in the current price of an asset, they will make a trade that contributes towards closing the gap between the asset’s current price and its estimated value. By this logic, if we compare two Bitcoin markets, the one with the larger average trade volume will have a greater importance to the Bitcoin trading ecosystem, since it is more often the venue on which new information initially arrives. Guided by this rationale, we choose as the basis of this paper those Bitcoin markets which have among the highest trading volumes. In Figure 1, we show a ranking of major Bitcoin markets by their average daily volumes over the course of February 2021. In the remainder of this paper, we will chiefly concern ourselves with the top 14 markets from this bar plot. Next, we provide further context on each of these markets along with their specifications.




Figure 1: Average daily trading volume in billion USD per market in February 2021

Market Specifications. Each of the markets ranked in Figure 1 can be classified as either a *spot* market, a *futures* contract market, or a so-called *perpetual contract* market (sometimes also called a perpetual swap, or simply a perpetual). The latter market type was invented by the crypto exchange Bitmex in 2014, and, to the best of our knowledge, does not exist outside of the crypto trading world at the time of this writing (July 2021). From the previously mentioned three types of markets, the perpetual swap is the most popular among crypto traders and it has by far the largest trade volumes. Since its inception, the success and broad popularity of the perpetual contract has compelled many other exchanges to follow the blueprint originally laid out by Bitmex. Competing exchanges such as Binance, Huobi, Okex, FTX, Deribit, Bybit, and Kraken have either directly copied this derivative contract or launched a nearly identical product with some minor modifications. The combined trade volume across these perpetual contracts routinely exceeds 50 billion USD per day [10].

Given the large role this derivative product plays in the crypto markets, it is worthwhile giving some background on its mechanics. Loosely speaking, the perpetual swap can be described as a “perpetually expiring futures contract” (hence its name). It allows traders to enter leveraged long or short positions without having to pay much attention to pricing a futures premium, since the perpetual swap is designed to closely track the price of a spot market (or a basket of spot markets). The mechanism by which the close link to the spot price is achieved is the so-called “funding rate”. This is a payment exchanged every eight hours (in case of the Bitmex and most other perpetual contracts) between the longs and the shorts. The magnitude and sign of this payment depends on the deviation of the perpetual’s price from the spot price over the preceding eight hour time window. For instance, if the price of the perpetual over such an eight hour time window was, on average, greater than the spot price, then the longs must make a funding payment to the shorts. The size of this payment depends on what the average deviation of the perpetual from the spot market was. The more the perpetual strayed above the spot price, the larger is the cost incurred by the longs. The introduction of this additional cost to the longs (and, in this case, rebate to the shorts) incentivizes traders to short the perpetual contract and it disincentivizes traders to open long positions. These twin incentives, when acted upon, will cause the price gap between the perpetual and spot markets to close. The reverse case where the price of the perpetual over an eight hour time window strays *below* that of the spot market is of course completely analogous. In this case, the longs will receive a rebate while the shorts incur a cost. More details on how exactly the perpetual swap functions can be found in [1].

An important property of Bitcoin derivative contracts is the type of initial capital (also known as *margin* or *collateral*) with which the derivative can be purchased. This is typically either Bitcoin, USDT (also known as Tether), or another so-called *stablecoin*. A stablecoin, of which USDT is the most prominent example, is a cryptocurrency whose value is (meant to be) pegged one-to-one to the value of a US dollar. At the time of this writing in July 2021, the market capitalization of Tether is around 61 billion USD. This provides an indication of the value that crypto traders place on being able to transact and store wealth in a quasi-USD asset which is void of many of the inconveniences associated with transacting actual USD, such as having to deal with slow bank transfers, stringent identity verification procedures, and so on. For more details on stablecoins in general and Tether in particular, we refer the reader to [21]. The FTX perpetual is unique among our set of 14 markets in that it allows positions to be

Table 1: Properties of markets under consideration in this paper

	Type	Exchange	Symbol	Max. Leverage	Margin	Expiration	Underlying
hbdm_BTC_CQ	Futures	Huobi	-	125x	BTC	Quarterly	BTC/USD
okex_BTC-USD-210326	Futures	Huobi	-	125x	BTC	Quarterly	BTC/USD
hbdm_BTC-USD	Perpetual	Huobi	-	125x	BTC	-	BTC/USD
bitmex_BTC/USD	Perpetual	Bitmex	-	100x	BTC	-	BTC/USD
deribit_BTC-PERPETUAL	Perpetual	Deribit	-	100x	BTC	-	BTC/USD
bybit_BTC/USD	Perpetual	Bybit	-	100x	BTC	-	BTC/USD
okex_BTC-USD-SWAP	Perpetual	Okex	-	125x	BTC	-	BTC/USD
ftx_BTC-PERP	Perpetual	FTX	-	101x	Cross	-	BTC/USD
binancecmfut_BTC/USD	Perpetual	Binance	-	125x	BTC	-	BTC/USD
binancefut_BTC/USDT	Perpetual	Binance	-	125x	USDT	-	BTC/USDT
okex_BTC-USDT-SWAP	Perpetual	Okex	-	125x	USDT	-	BTC/USDT
tbybit_BTC/USDT	Perpetual	Bybit	-	100x	USDT	-	BTC/USDT
thbdm_BTC-USDT	Perpetual	Huobi	-	125x	USDT	-	BTC/USDT
huobipro_BTC/USDT	Spot	Bybit	BTC/USDT	-	-	-	-

collateralized with various stablecoins, actual USD, Bitcoin, Ethereum and many other tokens.

The underlying index of the derivative contract is typically computed as a (weighted) average value of a basket of BTC/USD spot markets or BTC/USDT spot markets. It is worth pointing out that the BTC/USD price can deviate substantially from the BTC/USDT price under certain circumstances. For example, this can arise when investors lose confidence in the supposed peg of USDT to the US dollar, in which case the value of a USDT will decline. On the other hand, when there is particularly high demand for USDT (due to the conveniences associated with transacting quasi USD on a blockchain), investors might be willing to pay a premium for USDT. The exact composition of the baskets making up the index can vary from exchange to exchange. The BTC/USD indices typically include at least the Coinbase, Bitstamp and Kraken spot markets, while the BTC/USDT indices typically include spot markets from Binance, Huobi, and Okex. The precise mechanism via which the index price is computed from its constituents is generally also different between exchanges. For instance, some exchanges exclude the markets with the minimal and maximal price in the calculation, in order to achieve greater robustness of the index price and resistance to price manipulations [6]. The fact that trade volumes on derivative markets far exceed those on the spot markets which make up the derivatives' underlying indices can create opportunities for market manipulations. Colloquially speaking, the much larger volumes on derivative markets relative to spot markets is sometimes referred to as the "tail wagging the dog" phenomenon. These realities necessitate a certain robustness in the calculation of the index price in order to have resistance to market manipulations. To give a hypothetical example of such a manipulation, suppose the index of a futures contract was made up of just one illiquid spot market, where a relatively small trade can cause a big price move while the futures contract is extremely liquid. A trader might then get the idea to, for instance, open a massively leveraged long position in the futures contract (which can be done cheaply due to its liquidity) and manipulate the price of the underlying index upward shortly before the expiration of the futures contract. The trader could then profit handsomely from their leveraged long futures position when it expires at the artificially inflated price index price. These profits could be far in excess of the cost for performing this manipulation, which consists mainly of the execution cost associated with the "sloppy" manipulative buy orders on the illiquid spot market (which can be expected to be subject to high slippage cost). While this example was a hypothetical one, actual manipulations that resemble the above have actually occurred in reality. An interesting case study can be found in [4].

In Table 1, we list some key properties of the markets considered in the remainder of this paper. We see that the 14 most liquid markets from Figure 1 include two futures contracts, 11 perpetuals, and only one spot market. The most liquid futures markets tend to be the ones whose expiry happens at the end of a quarter. Huobi and Okex, the exchanges whose quarterly futures appear in our set of markets, both list a weekly expiration, a bi-weekly one, and a quarterly one. Okex additionally offers a bi-quarterly futures contract. In the remainder of this paper, we shall refer to binancefut_BTC/USDT as the "Binance USDT-margined perpetual contract" or simply the "Binance USDT perpetual". Similarly, binancecmfut_BTC/USD shall be referred to as the "Binance BTC-margined perpetual contract" or for short the "Binance BTC perpetual". Other markets will be described in the same fashion using their type, exchange, and when necessary, their margin currency.

Fake volume in crypto markets. For the purposes of this study, we aimed to identify the set of the most liquid Bitcoin markets which account for the majority of Bitcoin trading volume. However, the task of determining what exactly these markets are can prove more challenging than one might initially suspect. The reason for this is the prevalence of fake volume. The study [2] by Bitwise from 2019 asserts that 95% of reported volume in crypto markets is “fake”. The precise meaning of the word “fake” in this context is somewhat ambiguous. It can for example simply mean that an exchange is reporting a trade volume value that does not correspond to the sum of volumes of actual trades processed in the matching engine. Volume can also be “fake” in the sense that an exchange is printing wash trades where the exchange itself is engaging in self-trading. In other cases, exchanges have been known for partnering with market makers and incentivizing them to self-trade with certain type of rebates and bonus payment.

The reasons for exchanges to fake their volume can be manifold. A new exchange might use it as a method of bootstrapping a user base. Traders tend to trade on exchanges that are used by many other people and which therefore have large trade volume. If an incorrectly reported volume figure convinces a user to start trading on that exchange, then the exchange has successfully gained a new customer. From the exchange’s perspective, the cost of such practices is a loss of credibility which can negatively impact its long term success.

In most cases, it is straightforward to spot fake volume because exchanges tend not to be very sophisticated about the methods they employ to achieve this. Some methods border on the comical, such as printing a trade at mid-price periodically. Another method used by some exchanges in the past was to report a 24th trade volume figure which is different from the one which can be obtained by summing up the volume of all trades printed on the exchange. For a more exhaustive list of the various ways employed by exchanges to fake volume, we refer the reader to [11] that presents a detailed study of the subject.

Seasoned traders generally know from experience where the real volume resides and are therefore able to steer clear of the many pitfalls of fake volume. An excellent resource in this regard is [10] which provides adjusted volume figures where reported volumes were purged of likely fake volume when applicable.

Our choice of markets for this study was made with the aforementioned caveats in mind, and on the basis of our prior and ongoing trading experience. We are confident to a high degree that the 14 markets which form the foundation of this study are not among those which are subject to fake volume. Moreover, any volume figures reported in this study are based on our own calculations using data on individual trades published by the exchange’s real time data feeds. We do not rely on values reported by the exchange, although we do find agreement between these figures and the ones we independently calculated.

2.2 Data and Preprocessing

Our data set is comprised of orderbook and trades data. Both of these sources of data were mined by subscribing to the exchange’s websocket API orderbook and trades endpoints for each of the markets under consideration in this paper. Detailed information on what exactly a websocket API is can be found in [12], though we will give a cursory description below. The data we obtained in this way was compressed and placed into file storage for later processing.

Websocket API endpoints generally provide the most granular crypto market data that one can procure. These endpoints push new data to the subscriber in real time. For trades data, this means that one receives information about all executed trades on the market to which one is subscribed. In the case of orderbook data, exchanges typically publish orderbook snapshots every 5-200ms depending on market activity. Note that we therefore do not have order-by-order information (typically referred to as Level 3 data), since exchanges batch changes over a short previous time window and then publish these updates as part of a new orderbook snapshot. The frequency at which new orderbooks arrive not only depends on market activity but it also varies from exchange to exchange. In some cases, there is even variation in the update frequency within an exchange when a market has multiple websocket API endpoints with slightly different properties. For instance, the Bybit BTC-margined perpetual offers a feed that publishes a new snapshot containing 25 orderbook levels every 20ms, while at the same time offering another slower feed which publishes updates every 100ms but provides information for the top 200 orderbook levels [3]. A common minimum update frequency used by exchanges (even when they do not document this) is around 20ms. That is, orderbook updates arrive at a frequency not higher than one per 20ms, even when market activity is extreme. Exceptions are Deribit and Bitmex where update frequencies are sometimes in the single digit milliseconds when market volatility is high.

There is little uniformity across exchanges in the properties of their data feeds. That is, the number of orderbook levels varies from exchange to exchange (and even from market to market), the update frequencies are different, and the data comes in a different file format for each exchange. For each market in the set of 14 markets we consider in this paper, we receive between 25 and 75 orderbook levels. In our choice of the websocket API feed, we generally preferred speed over orderbook depth. A single day of orderbook data for all markets combined is over 30 gigabytes in size. Creating uniformity and comparability within our large and disorganized data set posed a major challenge.

Table 2: Lowest possible and default taker fees per market in basis points

	Lowest Possible Fee	Ordinary Fee
ftx_BTC-PERP	1.5	7
binancefut_BTC/USDT	1.53	4
binancecmfut_BTC/USD	1.8	5
huobipro_BTC/USDT	1.93	4.75
hbdm_BTC_CQ	2	4
okex_BTC-USD-210326	2.5	5
thbdm_BTC-USDT	2.7	4
okex_BTC-USD-SWAP	3	5
okex_BTC-USDT-SWAP	3	5
hbdm_BTC-USD	3.7	5
deribit_BTC-PERPETUAL	5	5
bitmex_BTC/USD	7.5	7.5
bybit_BTC/USD	7.5	7.5
tbybit_BTC/USDT	7.5	7.5

Furthermore, since the data is gathered asynchronously, timestamps of orderbooks generally differ. In order to more easily compare orderbooks across markets, we resample the data to a 50ms frequency using the last seen observation [8]. Missing values can appear when a market does not publish an orderbook update for, say, 100ms. In this case, we fill the missing value by propagating the last valid observation forward. For the sake of consistency, we also resample trades data to the same 50ms frequency. After applying this processing step, the buy (or sell) amount in a given 50ms window will be the sum of all buy (sell) trade amounts that occurred in this time window. The trade amounts can have different denominations on different markets. For example, for certain perpetual contracts, the value of one contract is 1 USD while on others it might be 100 USD or 0.001 Bitcoin. In order to have an “apples-to-apples” comparison between trade volumes across markets, we normalize all trade amounts to a USD amount, using the last seen Bitcoin price to convert amounts in Bitcoin to USD amounts.

Since we are working at a 50ms granularity (our resampling frequency), we shall define the trade price over a 50ms period as the volume-weighted average price computed over all trades (irrespective of side) which took place during this 50ms window. If no trade occurs during a time window, we fill the missing price value, as before, with the previous valid price observation.

With these preprocessing steps, successive samples in our data set are equally spaced in time (50ms apart), making cross-market comparisons easier. For each observation, we have one orderbook per market (potentially with differing numbers of levels), and information (buy amount, sell amount, average price) about a “meta trade” on each market. This is our data set in its final form, which will be the foundation for any further analysis.

2.3 Remarks on Execution Cost

The cost of execution of a taker order (one which leads to immediate execution) is composed of two different components. The first component is the taker fee charged by the exchange. Fees exhibit great variation across exchanges and even across markets within a single exchange. For instance, the taker fee for a derivatives contract tends to be, by and large, substantially lower than that of a spot market. A number of exchanges additionally offer tiered fee structures, where certain groups of traders receive lower fees than others. These fee rebates can be very substantial. Traders can ascend to a better fee tier by having high transaction volume. The volume cutoffs dictating a trader’s fee tier are different from market to market. Being in the top tier typically requires on the scale of 500 million to 1 billion USD of monthly notional volume. Certain exchanges (Binance, FTX, and Huobi) have issued their own cryptocurrency and offer additional fee discounts for traders that hold this cryptocurrency in their exchange wallet. The magnitude of this additional fee discount is usually proportional to the holdings in the exchange’s native cryptocurrency. An alternative (often undocumented and under-the-table) way of obtaining a preferred fee status is through the formation of strategic partnerships with exchanges. Particularly newly launched exchanges with little liquidity are typically interested in attracting liquidity, which they can achieve, for example, by partnering with market making firms whom they incentivize to provide the desired liquidity by offering various fee rebates. Table 2 displays the lowest possible (documented) taker fee that a trader in the top fee tier (who also owns a sufficient amount of the exchange’s native cryptocurrency, when applicable) receives.

Notice that these heavy fee rebates for certain “VIP” traders lead to a wealth concentration whereby VIP traders, who presumably have some measure of wealth by virtue of the requirements they must have met to be a member of the best fee tier, are capable of turning greater profits (or trading profitably at all) in their trading since they incur far less fee cost than, for instance, a novice trader would.

The second component of the execution cost can be broadly described by the term “spread”. Since the aggressor at least crosses the bid-ask spread, the minimum spread cost is at least the bid-ask spread at the time of execution. If the size of their order is greater than the liquidity available at the first orderbook level which is aggressed against, the trader additionally “walks the book”. That is, part of her/his execution occurs at one or multiple prices which are worse than the top quote. In Figure 2, we provide an illustration of the size of the spread cost incurred by an aggressor. In this plot, the set of blue bars represents the median difference in bpts between the minimum and the maximum price of a single order, while the red set of bars represents the median difference in bpts between the best price (minimum for buy orders, maximum for sell orders) and the average price of a single order. The sample consists of all orders of size > 100k USD between February 15–27, 2021. It is interesting to note that there seems to




Figure 2: Price impact of taker orders of size > 100k USD.

be an inverse relationship between a market’s VIP taker fee and the price impact (in other words, spread cost) of moderately large to large orders (greater than 100k USD). Writing

$$\text{execution cost} = \text{taker fee} + \text{spread}, \tag{1}$$

it appears that the two terms in this equation balance one another out, in a sense. If a market has a high taker fee, the orderbook appears to be more liquid and hence the spread cost is small. On the other hand, a market with a low taker fee has little liquidity at or near the top of the book, and hence the spread cost is large. We will later see variants on this theme reappear, when we perform PnL calculations that take fees into account.

3 Features and Preprocessing

3.1 Base Features

The fundamental microstructural quantity that determines price movements in a market that employs a limit orderbook is, broadly speaking, *order flow* which comes in two forms. The first one is aggressive, where traders submit market orders or limit orders that lead to immediate execution. The second one is passive, where market participants submit limit orders that do not lead to immediate execution or where such orders are canceled. Our objective in this subsection is to develop a set of base features that build on order flow, and which achieve high efficacy for certain tasks tackled later in this paper, such as predicting future returns or characterizing cancellation behaviour of HFT market makers. We shall discuss our base features one by one, and provide a motivation and justification for their inclusion. In later subsections, we discuss ways of transforming and optimizing the base features for the tasks that lay ahead.

3.1.1 Orderbook Imbalances

It is well established in the literature that the shape of the orderbook has significant impact on the distribution of future returns [16, 18–20]. The intuitive reason for this is straightforward: if the liquidity on the bid side far outweighs that of the ask side, the volume required for a price move up is much smaller than the volume required for a price move down. Thus, if we assume that the arrival of a sell order is not much more likely than the arrival of a buy order, the probability of a future price move up is larger than that of a price move down. Clearly, similar reasoning applies to the dual case where the orderbook shape takes the opposite form.

The empirical evidence cited above provides ample justification for the inclusion of an orderbook imbalance feature in our set of features. However, the question remains of the precise notion of orderbook imbalance that is most appropriate to our setting.

The classical notion of orderbook imbalance is given by

$$\frac{b-a}{b+a} \in [-1, 1], \quad (2)$$

where a and b represent the top of the ask and bid liquidity, respectively. However, this quantity suffers from an obvious drawback, namely that it only takes into account the liquidity at the two tops of the book. In other words, information from deeper orderbook levels is discarded. This is particularly problematic for markets where the liquidity at the top of the book tends to be very small and instead concentrates more strongly around deeper orderbook levels.

Extensions of the classical orderbook imbalance [31] typically try to account for liquidity on deeper levels by means of a weighting scheme which replaces the quantity a in Eqn. (2) by $\sum_j w_j a_j$, where a_j denotes the volume on orderbook level j and $w_j \in \mathbb{R}_{>0}$ is a weight. The quantity b is of course replaced in an analogous manner. The weight w_j is usually constructed in a way which reflects the probability of execution of an order on level j . It can, however, be challenging to get a handle of execution probability at each orderbook level. Conventional models of orderbook dynamics imply that the probability of execution is an exponentially decaying function of the distance to the midprice. Yet, in practice, it is not only the distance to midprice that matters, but also the fee structure (and possibly other factors). We saw evidence of this in Figure (2), where we noted that a low taker fee seems to imply more liberal “walking of the book” (that is, greater trade impact and hence a larger execution probability on deeper levels compared to markets where the taker fee is large). This makes it difficult in practice to choose such weights, since they must be chosen individually in a bespoke manner for each market. Further complicating things, different markets often have different tick sizes. For instance, the tick size on FTX is 1 USD, while Binance uses a tick size of 0.1 USD.

Another concern with the classical orderbook imbalance (and most of its extensions) is that it does not adequately reflect uncertainty of future returns in times of high volatility when (effective) spreads are large. To illustrate this point, consider an example of an orderbook during non-volatile times 3a and one shortly after a volatility spike 3b. If we compute the classical orderbook imbalance for the first orderbook, we obtain the value -0.98 which reflects well what one would intuit from looking at the orderbook: the ask liquidity is large and overwhelming compared to the bid liquidity, hence we would expect, all else being equal, a price move down to be much more likely than a price move up. The second orderbook gives a classical imbalance of 0.69 which one would interpret as a relatively strong predictive price signal of a price move up. However, a glance at deeper levels quickly reveals that also in this orderbook a price move down should be far more likely than a price move up. However, a glance at deeper levels quickly reveals that the liquidity distribution in this orderbook implies a far greater probability of a price drop than a price rise. What causes the discrepancy between the computed orderbook imbalance value and our intuition are the minuscule amounts posted at the two tops of the book, which ultimately have little bearing on the distribution of future returns due to their small sizes. Such small top-of-the-book amounts are a common occurrence during volatility breakouts and they often persist for several hundred milliseconds. For example, a submission of a mere 5,000 USD to the top ask would result in a vastly different orderbook imbalance of -0.42 . This phenomenon makes the classical orderbook imbalance values extremely noisy, rendering it unsuitable for practical purposes.

We propose an alternative pair of orderbook-based features (one for each of the ask and bid sides) which are more robust during volatility breakouts, and which do not depend on a complicated choice of weights. Let us define the following quantities for each market $i = 1, \dots, 14$ and for fixed quantities $N_1, \dots, N_{14} \in \mathbb{R}_{\geq 1}$

$$IMB_t^{a,i} := \left(\frac{p_{a,t}^i(N_i)}{p_{a,t}^i(1)} - 1 \right) \cdot 10000, \quad (3)$$

$$IMB_t^{b,i} := \left(\frac{p_{b,t}^i(1)}{p_{b,t}^i(N_i)} - 1 \right) \cdot 10000, \quad (4)$$

Orderbook (XBTUSD)		
Price	Size	Total
32829.0	83,000	3,817,200
32828.5	67,500	3,734,200
32828.0	41,800	3,666,700
32827.5	201,900	3,624,900
32827.0	63,700	3,423,000
32826.5	61,000	3,359,300
32826.0	253,000	3,298,300
32825.5	113,900	3,045,300
32825.0	74,000	2,931,400
32824.5	64,400	2,857,400
32823.0	8,600	2,793,000
32822.5	300,000	2,784,400
32822.0	226,200	2,484,400
32821.5	131,800	2,258,200
32821.0	155,300	2,126,400
32820.5	1,971,100	1,971,100

32817.5 ↓

32823.37 / 32819.88 | □□□□□

32820.0	20,900	20,900
32818.5	1,300	22,200
32817.5	500	22,700
32815.5	100	22,800
32814.5	1,000	23,800
32811.5	25,000	48,800
32810.0	10,100	58,900
32806.0	200	59,100
32803.0	5,800	64,900
32802.5	100	65,000
32802.0	400	65,400
32800.5	100	65,500
32800.0	15,700	81,200
32799.0	21,600	102,800
32798.5	9,300	112,100
32797.5	100	112,200

(a) Bitmex orderbook during non-volatile times. Sizes are in USD.

Orderbook (XBTUSD)		
Price	Size	Total
32731.5	25,000	3,387,600
32731.0	66,100	3,362,600
32730.5	94,500	3,296,500
32730.0	118,400	3,202,000
32729.5	110,000	3,083,600
32727.5	102,500	2,973,600
32727.0	72,300	2,871,100
32725.5	133,900	2,798,800
32725.0	331,300	2,664,900
32724.5	50,400	2,333,600
32722.0	1,250,000	2,283,200
32721.5	39,900	1,033,200
32720.5	304,300	993,300
32720.0	593,800	689,000
32719.5	94,800	95,200
32719.0	400	400

32719.5 ▼

32743.85 / 32740.38 | □□□□□

32718.5	2,200	2,200
32717.5	1,000	3,200
32717.0	100	3,300
32716.0	100	3,400
32713.0	100	3,500
32712.5	100	3,600
32712.0	20,600	24,200
32711.5	100	24,300
32711.0	190,300	214,600
32710.5	5,100	219,700
32710.0	7,100	226,800
32709.5	500	227,300
32709.0	200	227,500
32708.5	100	227,600
32708.0	6,300	233,900
32707.5	19,600	253,500

(b) Bitmex orderbook during volatile times. Sizes are in USD

where $p_{a,t}^i(x)$ denotes the average price one would pay at time t for a market *buy* order of size x on market i . Note that $p_{a,t}^i(1)$ is simply the top ask price. Similarly $p_{b,t}^i(x)$ is the average price for a market *sell* order of size $x \in \mathbb{R}_{\geq 1}$, so that $p_{b,t}^i(1)$ is just the top bid price. The quantity $IMB_t^{a,i}$ can be described as the difference in basis points between the top ask price on market i and the average price of a market order of size N_i , and analogously for the bid version $IMB_t^{b,i}$.

What are sensible choices of $N_i \in [1, \infty)$ for $i = 1, \dots, 14$? Note that $N_i = 1$ always yields $IMB_t^{a,i} = 0$, while letting $N_i \rightarrow \infty$ we have $IMB_t^{a,i} \rightarrow \infty$; thus the two extreme ends of the spectrum are void of any signal. We set N_i to be the median liquidity within the top five basis points of the top of the book on market i . More precisely, we compute this median value for the ask side and the bid side separately, and then take the average of the two resulting values. Preliminary experiments indicate this to be a sensible choice, in the sense that the resulting features (4) appear to contain much signal and little noise. We leave it to future work to calibrate the choice of N_i in a more rigorous fashion.

Let us revisit the two orderbooks considered above. The orderbook from Figure 3a has $(IMB_t^{a,i}, IMB_t^{b,i}) \approx (0, 7)$. The small $IMB_t^{a,i}$ value means that the ask side is heavily populated, while the large $IMB_t^{b,i}$ points at the sparsity of the bid side. From the comparison of these two values, we would unequivocally conclude that a price drop is far more likely than an increase. For the orderbook from Figure 3b we have $(IMB_t^{a,i}, IMB_t^{b,i}) \approx (1, 7.4)$. As in the

Table 3: Comparison of R^2 values of classical orderbook imbalance and our proposed imbalance

	$R^2_{\text{classical}}$	R^2_{new}
bybit_BTC/USD	0.1176	0.1646
hbdm_BTC-USD	0.0893	0.13
thbdm_BTC-USDT	0.0957	0.1059
tbybit_BTC/USDT	0.0724	0.0818
bitmex_BTC/USD	0.0547	0.0786
binancecmfut_BTC/USD	0.0228	0.0645
hbdm_BTC_CQ	0.0179	0.0425
okex_BTC-USD-SWAP	0.0315	0.0358
okex_BTC-USDT-SWAP	0.0103	0.0275
ftx_BTC-PERP	0.0157	0.0273
huobipro_BTC/USDT	0.0135	0.0204
okex_BTC-USD-210326	0.0035	0.0159
deribit_BTC-PERPETUAL	0.0115	0.0155
binancefut_BTC/USDT	0.0	0.0061

previous case, the far larger $IMB_t^{b,i}$ value suggests a much greater probability of a price decrease, which agrees with our intuition from observing the liquidity distribution in the orderbook, and which is in contrast to the classical orderbook imbalance value that pointed to a greater likelihood of a price increase.

Moreover, our pair of features retains its high fidelity during highly volatile times when liquidities on both sides of the book are sparse. In such cases, the greater uncertainty in the distribution of futures returns is represented by values $IMB_t^{a,i}$ and $IMB_t^{b,i}$ which are both large.

As further empirical evidence of the suitability of the features defined in (4), we compare their explanatory power with that of the classical orderbook imbalance. Specifically, for each market $i = 1, \dots, 14$, we fit the following pair of univariate linear models

$$fret_t = \alpha + \beta \left(IMB_t^{a,i} - IMB_t^{b,i} \right) + \epsilon_t, \quad (5)$$

$$fret_t = \alpha' + \beta' \left(\frac{b_t^i - a_t^i}{b_t^i + a_t^i} \right) + \epsilon'_t, \quad (6)$$

where a_t^i and b_t^i are the top ask and bid on market i , respectively. The prediction target $fret_t$ used here consists of 500ms future returns on the Bybit BTC perpetual. We fixed this target market for the sake of concreteness; other markets give similar results. The above linear models are fitted using OLS on training data spanning February 22nd until February 27th. See Table 3 for a comparison of the corresponding R^2 values. With the new imbalance measure, we find an average increase in R^2 of 0.019 compared to the classical one. The R^2 value corresponding to the classical imbalance is only approximately 58% that of the new imbalance which, additionally, is higher in every single case.

3.1.2 Trade Imbalances

Let us define the so-called *trade flow imbalance* indicators, for each market $i = 1, \dots, 14$, computed over time horizons $\delta \in \{100\text{ms}, 250\text{ms}, 500\text{ms}, 1000\text{ms}, 2000\text{ms}\}$

$$TFI_t^{i,\delta} := B_{[t-\delta,t]}^i - S_{[t-\delta,t]}^i, \quad (7)$$

where $B_{[t-\delta,t]}^i$ represents the total volume of buy orders on market i in the time interval $[t - \delta, t]$. Similarly, the quantity $S_{[t-\delta,t]}^i$ represents the total volume of sell orders on market i over the same time interval.

This feature can be construed as the “aggressive” component (relating to taker flow rather than submissions or cancellations of passive limit orders) of the order flow imbalance indicator defined by Cont et al. in [20]

$$ofi_t = [\text{buy flow} - \text{sell flow}]_{[t-\delta,t]} \quad (8)$$

$$= \underbrace{[(\text{buy trades} - \text{sell trades})]}_{\text{“aggressive” component}} + (\text{bid submissions} - \text{ask submissions}) + (\text{ask cancels} - \text{bid cancels})_{[t-\delta,t]}. \quad (9)$$

The order flow imbalance indicator was demonstrated in [20] to have significant explanatory power over *contemporaneous* returns. In ongoing work by a subset of the authors, similar results are demonstrated for future returns, when employing a variation of the order flow imbalance from [20], showcasing the usefulness of this family of indicators.

The fragmented nature of the Bitcoin trading universe means that volumes are scattered across several venues. This implies an asynchronous information arrival: trade flow typically arrives first on one venue and then “trickles down” to other venues in the form of cancellations or taker orders. Following this logic, we would expect to also have explanatory power over future returns based on the cross-section of trade flows. This motivates our inclusion of the trade flow imbalance indicators (7) into our set of base features. Additional remarks are in order.

- We believe that the aggressive component in (9) holds a fundamentally greater importance (more signal) than the remaining two “passive” terms. The underlying rationale is that passive flow (submissions or cancellations) are almost exclusively the work of HFT market makers who are largely only reacting to the cross-section of aggressive flow and orderbook imbalances. That is, passive flow tends to be *reactive* in nature rather than contain genuine new information. A rigorous investigation of this topic is left for future work.
- Notice that future returns are a dimensionless quantity measured in basis points, while the unit of trade flow imbalance is USD. Based on this discrepancy of units, one might predict that future returns vary non-linearly on trade flow imbalance. We will see later that this is indeed the case, and will devise a methodology for addressing this concern.
- One might get the idea to consider a relative measure of the form $\frac{B-S}{B+S}$, where B and S denote buy and sell volume (respectively) instead of the trade flow imbalance as we defined it. However, this relative indicator suffers from some of the same drawbacks as does the classical orderbook imbalance, namely a far noisier predictive signal than its non-relative USD denominated analogue. For instance, if the sell volume is zero, a trivial buy volume of size 1 USD gives rise to the same imbalance value as a significant buy volume of size, say, 10 million USD.
- The nature of our orderbook data (snapshots data) makes it impossible to precisely compute the remaining two “passive” terms that make up the order flow quantity (9). A precise computation requires order-by-order feeds, and is left as future work, for the case of equity data, for which such feeds are available¹.

3.1.3 Past returns

Following a similar logic as above for the trade flow imbalance, it is to be expected that a drastic price change on one market should cross-impact short-term future returns on another market. It is therefore natural to include past returns on each market in our feature set. We use the same time horizons here as for the trade flow imbalances. Formally, for each market $i = 1, \dots, 14$ and time horizon $\delta \in \{100\text{ms}, 250\text{ms}, 500\text{ms}, 1000\text{ms}, 2000\text{ms}\}$, we define the following indicator at each time t

$$PRET_t^{i,\delta} := \left(\frac{p_t^i}{p_{t-\delta}^i} - 1 \right) \cdot 10000, \quad (10)$$

where p_t^i denotes the price on market i at time t . This quantity in (10) computes the price change in basis points on market i over the time interval $[t - \delta, t]$. It needs to be clarified precisely what notion of price we are using in p_t^i , since there are many possible choices (mid price, last trade price, and so on). Our initial choice of last trade price gave rise to ambiguities in the case when a number of trades are executed simultaneously at multiple different price levels. This happens for example for “sweeping” orders which consume multiple orderbook levels. To circumvent this concern, we compute the average price in these cases. More precisely, we will use the average price over a short lookback window of 50ms (which is the minimum time step in the preprocessing of our data). That is, we define p_t^i as

$$p_t^i := \frac{1}{A} \sum_{(a,p) \in T_t^i} a \cdot p, \quad (11)$$

where T_t^i is the set of all trades, i.e. pairs (a, p) of amount and price, executed on market i in the time period $[t - 50\text{ms}, t]$ and $A = \sum_{(a,p) \in T_t^i} a$. Note that we include both buy and sell trades in T_t^i .

¹To this end, intraday LOBSTER data could be employed, which provides detailed tick-by-tick limit order book data for NASDAQ stocks. LOBSTER contains all limit order submissions, cancellations and executions on each trading day. All events are time-stamped with millisecond precision [17], and are reconstructed from NASDAQ’s historical TotalView-ITCH data [23].

3.1.4 Mean Divergence


The instruments traded on the markets under consideration in this paper are closely related to one another, although not identical. The similarity stems from the fact that they are all either perpetuals or futures contracts on a Bitcoin underlying (a BTC/USDT or BTC/USDT index), or a BTC/USDT spot market. We would, of course, expect price differences to be mean reverting processes since if this were not the case there would supposedly be arbitrage opportunities. This furnishes the motivation for the inclusion in our set of base features of the price differences between pairs of markets, as indicators for future returns. One would expect, for instance, that if a particular market is cheap relative to some (or all) other markets (e.g. due to a price move up on the other markets) there should be a short-term price increase on the cheap market, as arbitrageurs and market makers rush to profit from the arbitrage. However, there are several assumptions underlying this hypothesized market behaviour which do not hold in practice, and which therefore make the actual dynamics more complicated. We list some of these confounding factors below.

1. The presence of trading fees restricts the set of arbitrages. That is, price differences that would represent arbitrages in a world without trading fees are in fact not arbitrage opportunities. The consequence of this is that the price difference between a pair of markets can fluctuate within a “no arbitrage band” whose size is given by the sum of the taker fees of the two markets. Only when the price difference is greater than the sum of the two taker fees does there exist an arbitrage that can be immediately profited from with a pair of taker orders.
2. The futures premium of, say, quarterly futures contracts is (mostly) nonzero, so that the price difference between a futures contract and another market is, a-priori, not indicative of where the price will move in the near future. Rather, what matters is how far the price difference between the futures contract and the other market strays from its “equilibrium” level. This equilibrium is difficult to get a handle of in practice, since it depends on future expected interest payments whose magnitude is uncertain.
3. The price difference between a BTC/USDT based market (e.g. a perpetual whose underlying is a basket of BTC/USDT spot markets) and a BTC/USD based market alone cannot be expected to be indicative of future returns since it depends on the USD/USDT rate, which is generally only *approximately* 1, but can fluctuate by several basis points, and in rare cases even percentage points. As we discussed before, such fluctuations are usually related to varying investor confidence in the peg between USD and USDT which can decrease the value of USDT (see [21]). On the other hand, the value of one USDT can temporarily exceed that of an actual USD during periods where investors are willing to pay a premium for the conveniences associated with transacting in USDT rather than USD (usually when sentiment is extremely bullish). Deviations of the USD/USDT rate from 1 usually cannot be *immediately* arbitrated since the conversion typically takes several business days and involves a bank transfer.
4. The dominant trader demographics differ from exchange to exchange. For instance, US citizens are completely prohibited from trading on Bitmex and must therefore resort to other exchanges. These types of restrictions can result in persistent price differences, say, between Bitmex and another exchange which is popular with US traders. One reason for this could be that prevailing investor sentiment and expectations of future returns differ between distinct demographics, for example due to divergent macroeconomic backdrops. Other reasons for persistently large price divergences include differences in capital controls, financial regulation, “know-your-customer” (KYC) requirements, etc. across different countries and exchanges. A prominent example is the premium of the Bitcoin price on Korean exchanges relative to US exchanges, dubbed as the *Kimchi premium*, which at various points in the past was notoriously large for extended periods of time. More details can be found in [7].

Altogether, these factors diminish the predictive signal of price differences alone. For visual evidence of this, consider 4a where we plot the price difference over an eight hour period between a pair of perpetuals where one uses BTC/USD as its underlying, while the other uses a BTC/USDT index.

We propose to circumvent the concerns listed above by correcting for average price differences over a certain lookback window. That is, instead of computing “vanilla” price differences, we compute the deviation of the price difference from the mean price difference over a past time period ranging from several seconds to several minutes. The rationale for this is that persistent large price differences arising for some of the reasons mentioned above (nonzero futures premium, USD/USDT rate different from 1, or differing investor outlook between distinct groups of traders) are usually unchanged over such short time horizons. For instance, barring the occurrence of an extreme outlier event, the repricing of a futures premium most often happens over the course of several hours or days as investors gradually adjust expectations of future interest payments. With our proposed new indicator, we are

therefore, in a sense, “purging” the price difference of such confounding longer time scale effects. In Figure 4b we show the same price difference depicted in Figure 4a purged by its five minute rolling mean. Note the seemingly much stronger tendency for mean reversion.



(a) Price difference between the Binance USDT perpetual and the Huobi BTC perpetual (b) Mean divergence feature for the Binance USDT perpetual and the Huobi BTC perpetual

Figure 4: To avoid an excessively large number of samples (and hence to facilitate legibility), the data was resampled to a 2 minute frequency.

This altogether motivates the definition of the following features for pairs of markets $(i, j) \in \{1, \dots, 14\}^2$ and time horizons $\Delta \in \{5s, 9s, 19s, 38s, 75s, 150s, 300s, 600s\}$

$$DIV_t^{i,j,\Delta} = d(p_t^i, p_t^j) - \text{rolling}^\Delta \left(d(p_t^i, p_t^j) \right), \quad (12)$$

where $d(p, q) = \left(\frac{p}{q} - 1 \right) \cdot 10000$ is the difference in basis points between $p, q \in (0, \infty)$, and the $\text{rolling}^\Delta(\cdot)$ function returns the rolling mean of samples from its input over the past Δ seconds. We shall call $DIV_t^{i,j,\Delta}$ the *mean divergence* feature between markets i and j .

As further evidence of the superiority of this indicator compared to the vanilla price difference, we display in Table 4 a comparison of R^2 values corresponding to the univariate regression models

$$fret_t = \alpha + \beta \left(d(p_t^{i_0}, p_t^j) \right) + \epsilon_t, \quad (13)$$

and

$$fret_t = \alpha' + \beta' \left(DIV_t^{i_0,j,150ms} \right) + \epsilon'_t, \quad (14)$$

for each $j = 1, \dots, 14$. For sake of concreteness, we fix a reference market i_0 , namely the Bybit BTC perpetual, whose future returns are used in the above models. That is, we are comparing the explanatory power over Bybit’s future returns of, one the one hand, the “plain” price difference between Bybit and other markets, and, on the other hand, the mean divergence feature of Bybit with other markets. On average, the explanatory power of our new indicator is larger than that of the previous one by 0.074. This represents an average *15-fold* increase. Note that the improvement is particularly drastic for the two quarterly futures contracts `hbdm_BTC_CQ` and `okex_BTC-USD-210326` where we see a more than 70-fold increase.

3.2 Eliminating Nonlinearities

Our overarching objective in this paper is to train powerful *linear* price prediction models using the features defined above. The performance of a linear model is, however, limited by the extent to which the future returns vary *linearly* as a function of the feature variables. That is, if the relationship between variates and covariates is highly nonlinear, we cannot possibly expect to achieve good generalizability of our model.

In this subsection, we will demonstrate the nonlinear nature of the dependence between future returns and some of our indicators. We will then devise a methodology that addresses these concerns. Our approach involves the construction of bespoke nonlinear transformations which will subsequently be applied to our features prior to their usage in linear models.

To illustrate the presence of some of the aforementioned nonlinearities consider Figure 5a. This scatter plot exhibits the relationship between 500ms trade flow imbalance of the Binance USDT perpetual and 500ms future returns of the Huobi BTC perpetual. Each x -value in this plot corresponds to a trade flow imbalance bucket

$$B_n := [10000 \text{ USD} \cdot (n - 1), 10000 \text{ USD} \cdot n]$$

Table 4: Comparison of R^2 values of vanilla price difference (left column) with divergence from mean (right column)

	R^2_{diff}	R^2_{div}
binancecmfut_BTC/USD	0.057	0.1426
binancefut_BTC/USDT	0.0252	0.1331
hbdm_BTC_CQ	0.0017	0.1224
okex_BTC-USD-210326	0.0015	0.119
okex_BTC-USD-SWAP	0.0337	0.1169
okex_BTC-USDT-SWAP	0.0175	0.1113
thbdm_BTC-USDT	0.0099	0.0961
hbdm_BTC-USD	0.0175	0.0778
ftx_BTC-PERP	0.031	0.0754
deribit_BTC-PERPETUAL	0.0149	0.0744
huobipro_BTC/USDT	0.0118	0.0661
tbybit_BTC/USDT	0.0032	0.0393
bitmex_BTC/USD	0.0038	0.016

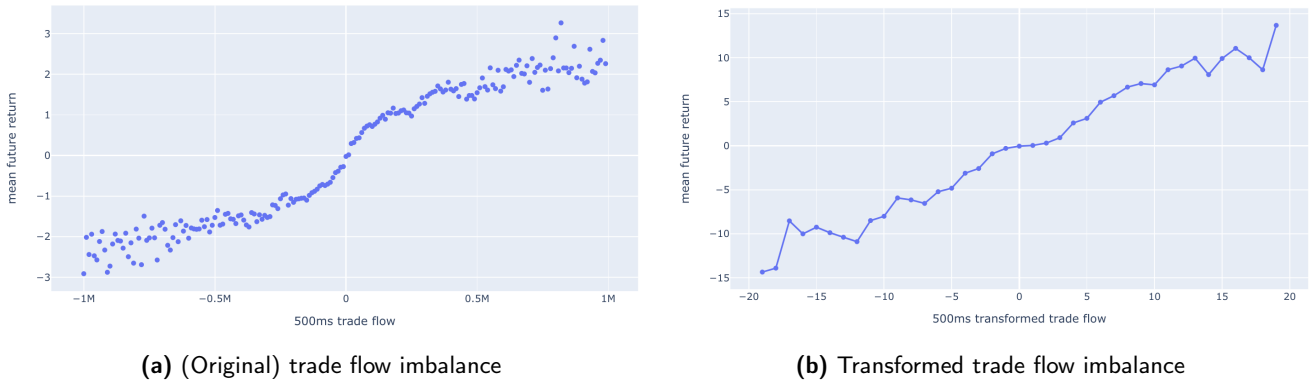


Figure 5: 500ms future returns of the Huobi BTC perpetual as a function of Binance USDT perpetual 500ms original and transformed trade flow imbalance.

for $n \in \{-100, 100\}$. The y -value associated to the bucket B_n was obtained by computing the mean 500ms future returns on the Huobi perpetual over all samples where the Binance trade flow falls into the interval B_n . The functional form we observe in Figure 5a clearly resembles a sigmoid or hyperbolic tangent function more than it does a linear one (although samples for buckets B_n with large $|n|$ are sparser, hence the data noisier). The sigmoid-like relationship between trade flow imbalance and future returns agrees well with prior work, which noted the same empirical observation [27, 28].

Why does the price impact decrease as trade flow imbalance grows? One possible economic interpretation is that extreme trade flow values (either very large or very small) could correlate negatively with trader sophistication and correlate positively with the “sloppiness” of execution. That is, an informed trader wishing to enter a large position will rarely do so by using, for example, one large market order or a large “meta order” executed over a time window as small as 500ms. It is perhaps more likely that extreme trade flow values originate from uninformed traders, or, in the extreme case, are the result of forced liquidations where traders involuntarily liquidate their entire positions as a consequence of a margin call. If our hypothesis is correct, the relative lack of sophistication reflected by extreme trade flow imbalances could be correctly recognized by other (more astute) market participants who would then only minimally adjust their estimation of the “fair price”, and hence contribute towards a more limited price impact.

Further below, we describe our methodology for mapping our (original) features to nonlinearly transformed ones. Before we proceed with this description, it is interesting to visually compare Figure 5a, which we already considered earlier, with Figure 5b where we plot the functional dependence of Huobi returns on the *transformed* Binance trade flow imbalance feature. Note the stronger degree of linearity in the dependence of future returns on the transformed feature.

Methodology for feature transformation. The intuition behind our methodology is to define, for each feature, a sequence of transformations that starts with the identity map and ends with a function which yields a significant boost in explanatory power when it is applied to the feature. This sequence is defined through an iterative procedure where each step of the iteration achieves an R^2 improvement.

Before specifying our general algorithm, it is instructive to see what happens in case of the trade flow imbalance feature. To this end, let us preliminarily define, for a finite set of real numbers T , a mapping $f_T : \mathbb{R} \rightarrow \mathbb{R}$ as follows

$$f_T(x) := \text{sgn}(x) \sum_{t \in T} \mathbf{1}\{|x| \geq t\}, \quad (15)$$

where $\text{sgn}(\cdot)$ is the sign function, and $x \mapsto \mathbf{1}\{|x| \geq t\}$ denotes the indicator function which takes the value 1 when $|x| \geq t$ and 0 else. Now suppose $X = \{x_t\}_t \subset \mathbb{Z}$ is the set of observations of the trade flow imbalance on some fixed market. Let us then set $m := \min(X)$, $M := \max(X)$ and $T_0 := \{m, m+1, \dots, M-1, M\}$. Note that we can then write the identity function in a more complicated way as follows

$$\text{id}_X = \left(x \mapsto \text{sgn}(x) \sum_{t \in T_0} \mathbf{1}\{|x| \geq t\} = f_{T_0}(x) \right). \quad (16)$$

This stems from the fact that, in the sum, we have precisely $|x|$ many nonzero indicator functions. The next step of our procedure involves the elimination of one element i^* from the set T_0 , giving rise to an updated set $T_1 = T_0 \setminus \{i^*\}$, which in turn defines a new mapping

$$\left(x \mapsto \text{sgn}(x) \sum_{t \in T_1} \mathbf{1}\{|x| \geq t\} = f_{T_1}(x) \right), \quad (17)$$

representing a slight alteration of the identity mapping. How do we decide which element i^* to prune from the original set T_0 ? We do so by trying all possibilities, and then ultimately choosing the element which provided the largest improvement in average explanatory power of future returns.

Specifically, we fit the following univariate linear models using OLS regression for every market $j = 1, \dots, 14$ and every element $i \in T_0$

$$fret_t^{500\text{ms},j} = \alpha + \beta f_{T_0 \setminus \{i\}}(x_t) + \epsilon_t. \quad (18)$$

We denote the coefficient of determination of the model specified in Eqn. (18) by R_{ij}^2 . That is, R_{ij}^2 is the explanatory power of the linear model that predicts future returns on market j using the trade flow imbalance feature of the fixed market which has been transformed by the function $f_{T_0 \setminus \{i\}}$. We then form an average across market j , and select the transformation that resulted in the greatest average coefficient of determination. That is, we set

$$\overline{R}_i^2 := \frac{1}{14} \sum_{j=1}^{14} R_{ij}^2, \quad (19)$$

and finally define $i^* := \text{argmax}_{i \in T_0} \overline{R}_i^2$. This yields the set $T_1 := T_0 \setminus \{i^*\}$ with corresponding transformation f_{T_1} . We then prune another element in the same manner as above, in order to obtain a next set T_2 which defines a new mapping f_{T_2} , and so on. This process of successive elimination is terminated when none of the eliminations yields an improvement in average explanatory power. That is, we halt the process when the average explanatory power under a transformation $f_{T_{k+1}}$ is not greater than that corresponding to the previous transformation f_{T_k} . At this point, we have obtained our final set $T^* := T_k \subset T$ which defines the (nonlinear) feature transformation f_{T^*} for the trade flow imbalance feature we fixed in the beginning. In practice, it is computationally infeasible to begin with the initial set T_0 defined above, whose cardinality is $|M - m + 1|$, which, for the trade flow imbalance feature, is on the scale of tens of millions. We instead choose an initial set T_0 of cardinality 100 representing an evenly spaced partition of the interval $[m, M]$. Moreover, to mitigate the impact of outliers, we do not use $m = \min(X)$ and $M = \max(X)$, but instead define the two quantities m and M as a very small quantile and a very large quantile, respectively, of the set X of feature observations. The data used to calibrate the feature transformation spans February 21st 2021, the day prior to our one week training and test period.

We now specify the general algorithm for obtaining the final "optimal" feature transformation for a fixed feature with observations $X = \{x_t\}_t$. The initialization is as follows. First, we define \overline{R}_0^2 to be the average of the R^2 values of the linear models

$$fret_t^{500\text{ms},j} = \alpha + \beta x_t + \epsilon_t,$$

for $j = 1, \dots, 14$. Now let m and M denote the q -th and $(1-q)$ -th quantiles of X , respectively, and initialize T_0 as an evenly-spaced partition with 100 elements of the interval $[m, M]$. In these calculations, we employed $q = 0.0001$ which was revealed as a sensible choice in preliminary experiments.

The iterative step for obtaining T_{k+1} from $T_k = \{t_1, \dots, t_n\}$ proceeds in the following manner.

1. Fit the linear model $fret_t^{500ms,j} = \alpha + \beta f_{T_k \setminus \{t_i\}}(x_t) + \epsilon_t$ for each market $j = 1, \dots, 14$ and every $i = 1, \dots, n$ and denote the corresponding coefficient of determination by $R_{j,i,k}^2$.
2. With $\overline{R}_{i,k}^2 := \frac{1}{14} \sum_{j=1}^{14} R_{j,i,k}^2$ we make the definitions $i^* := \operatorname{argmax}_{i=1, \dots, n} \overline{R}_{i,k}^2$ and $\overline{R}_{k+1}^2 := \overline{R}_{i^*,k}^2$
3. If $\overline{R}_{k+1}^2 > \overline{R}_k^2$ we define $S := \{t_{i^*}\}$, else we let $S := \emptyset$.
4. Finally, we set $T_{k+1} := T_k \setminus S$.

Clearly, this iteration terminates eventually since either S becomes \emptyset or the set T_k itself becomes empty. If the process has terminated after K steps, we will further define $T^* := T_K$ and $\overline{R}^2_* := \overline{R}_K^2$.

Results. For each feature, we compute its optimal transformed feature according to the above method. A natural question to ask is for which features does the above transformation procedure result in a significant improvement in explanatory power of future returns by a linear model? By far the biggest improvement was found for the trade flow imbalance feature. In Figure 6 we plot the average R^2 values obtained using the (untransformed) 500ms trade flow imbalance feature on the one hand, and its optimal transformed version on the other hand.

More specifically, we compare the two quantities \overline{R}_0^2 and \overline{R}^2_* (as defined above) which capture the average explanatory power of the original and the transformed feature, respectively. The results are shown in Figure 6. The transformation yielded a significant improvement for every market. The average improvement is by a factor of 2.5. The improvement is largest on the Deribit perpetual contract, where we observe an increase by a factor of 7.3. The smallest improvement occurs on Bybit’s perpetual contract, where we nevertheless note an increase of 20%.




Figure 6: Trade flow imbalance improvement of transformation versus original feature.

The only other feature where we found, on average, a significant improvement using our previously described transformation procedure is the past returns feature. See Figure 7 for a barplot of the results. On average, we note an improvement of 28%. This is far lower than the one attained for the trade flow imbalance feature, but it is still significant. On the FTX perpetual, the increase in average explanatory power was largest, with an increase by a factor of 2.11. The Bybit USDT-margined perpetual actually saw a decrease in explanatory power by around 33% (although it was quite low to begin with). This was the largest decrease of any markets.

The two remaining features (orderbook imbalances and mean divergence) did not show a significant average boost in explanatory power of the transformed feature over the original version. For the orderbook imbalance feature, we found an average increase of around 3%, while the mean divergence feature with lookback window 150s produced an average decrease of around 5% (with similar results holding true for other lookback windows). This lack of R^2 gain points towards an already linear relationship between the aforementioned features and future returns.

In the remainder of this paper, we will use the transformed trade flow imbalance feature instead of the original version. On markets where the transformation yielded an improvement for the past returns feature, we shall use the transformed version. Where a decrease was found, we will continue using the original version. For the orderbook imbalance and mean divergence features, where we did not find a significant improvement, we will simply continue using the untransformed features.




Figure 7: Past returns improvement of transformation versus original feature.

3.3 Optimal Time Horizon Selection

With the exception of the orderbook imbalance features, all of our base features are defined over multiple time horizons. From a standpoint of complexity reduction, however, it is desirable to select only one time horizon per base feature. This is because base features of a single feature category at different time horizons share a high degree of collinearity. For instance, it is clear that 500ms past returns are highly correlated to 1s past returns (the values will often coincide). In this subsection, we set out to devise a methodology by which we select a single time horizon for any feature defined over multiple time horizons. With an eye towards eventually training large composite models that employ the full feature set, this additional feature selection step will be useful in ameliorating concerns of overfitting (it is common knowledge that OLS regression is prone to overfitting when the number of covariates is large and the covariates exhibit high cross-correlations). Additionally, it is an interesting research question in and of itself to investigate which lookback windows for which features are most predictive of, say, 500ms future returns.

Our methodology for selecting the “optimal” time horizon involves a simple process of computing an average predictability score based on R^2 values for each time horizon, and then choosing the one with the highest score.

Methodology for selection of optimal time horizon. Suppose we are trying to determine the optimal time horizon δ_{k^*} corresponding to some set of features $f^{\delta_1}, \dots, f^{\delta_n}$. We proceed in two steps. The first step consists of using OLS regression to fit the following model for every time horizon $k \in \{1, \dots, n\}$ and target market $j = 1, \dots, 14$

$$fret_t^{500ms,j} = \alpha + \beta f_t^{\delta_k} + \epsilon_t, \quad (20)$$

whose coefficient of determination we denote by $R_{k,j}^2$. For the second step, we compute the following averages for each time horizon $k \in \{1, \dots, n\}$

$$\bar{R}_k^2 := \frac{1}{14} \sum_{j=1}^{14} R_{k,j}^2, \quad (21)$$

and then set $k^* := \operatorname{argmax}_k \bar{R}_k^2$. In this case, we shall call δ_{k^*} the *optimal time horizon* for the feature f . In summary, a feature’s optimal time horizon is the one whose average explanatory power over 500ms future returns is largest.

Empirical results. We now examine the empirical results of the procedure described above.

(1) Let us first consider the **trade flow imbalance (TFI) indicators**. Since we found that the *transformed* TFI yielded a considerable improvement over the untransformed one, we will in fact use the transformed version, though we shall, for the sake of brevity, generally omit the qualifier “transformed”. In Figure 8 we illustrate for each market’s (transformed) TFI the quantities $\bar{R}_1^2, \dots, \bar{R}_4^2$ corresponding to the five time horizons.

We observe that the smallest and largest time horizons are not selected for any market, indicating a sensible choice of lower and upper bound for the size of the lookback windows with the optimal value falling somewhere in-between. Next, we note that the most commonly preferred horizon (eight cases) is 500ms, while the 250ms horizon and the 1000ms horizon are preferred in two and four cases, respectively. Finally, it is noteworthy that the five markets whose trade flows have the most predictive power are all on either Binance or Huobi. Specifically, this set of markets is comprised of the two perpetual contracts (one of them being margined with USDT, the other one with




Figure 8: Average explanatory power of TFI features for each time horizon

Bitcoin) on Binance and Huobi, as well as the Huobi quarterly futures contract. This furnishes our first indication of the leading behaviour of Binance and Huobi, which is a theme we that will repeatedly reappear later in this work.

(2) We now examine the **past returns indicators**. For these features, we previously also found an improvement when using the transformed feature instead of the original one, so we shall again use the transformed version in the following analysis. In Figure 9, we show for each market the average explanatory powers $\overline{R^2}_1, \dots, \overline{R^2}_5$ corresponding to the five time horizons. As was the case with the TFI indicators, we note that also for the past returns features,




Figure 9: Average explanatory power of past returns features for each time horizon

neither of the optimal time horizons is attained at the minimal or maximal time horizon, providing further validation for the choice of time windows. In four cases, the 1000ms time window is preferred, while the 250ms horizon is selected in two cases. The most commonly selected optimal horizon is 500ms for six markets. It is noteworthy that the two markets where the smaller time horizon (250ms) is preferred are also among the markets where the smaller time horizon is preferred for the TFI feature. Furthermore, the markets which appeared to be leaders based on their comparatively large R^2 values for the TFI features, also have among the highest explanatory powers in terms of their past returns features. This lends additional confidence to the previously conjectured leading behavior of the two Binance and Huobi perpetual contracts, as well as the Huobi quarterly futures market.

(3) The third and final feature whose definition relied on the choice of a time horizon is the **mean divergence feature**, which we shall examine now. Here, the feature transformation did not yield an improvement, so we will simply use the original feature. The corresponding illustration of each market's average R^2 values can be seen in Figure 10. Our smallest time horizon (5s) appears twice as the optimal choice. The time horizons 9s, 19s, and 38s




Figure 10: Average explanatory power of mean divergence features for each time horizon.

were each selected as the optimal choice for exactly four markets. The previously noted pattern of leading markets persists in this feature class, as well with the Binance USDT-margined perpetual achieving the largest average R^2 value, followed by the Huobi futures contract and then the Binance BTC perpetual.

4 Network Effects in Bitcoin Markets

With our feature set in its final form, we now turn to examining network effects amongst the markets considered in this work. We begin by investigating the presence and strength of *lead-lag relationships* between pairs of markets. After performing this analysis, we will examine how it translates into the realities of trading.

4.1 Leader-Lagger Network

What does it mean for a market j to lead another market i ? Intuitively, one would say that a market j *leads* a (lagging) market i if future returns on the lagging market i can reliably be anticipated based on information from market j . In our case, this information consists of microstructural data encapsulated in the features we previously calculated.

Let us make the previous formulation more precise. A market i is said lag another market j if a large portion of the total variation in future returns of market i can be explained by features of the leading market j . It is therefore natural to fit the following linear models for each pair of markets $(i, j) \in \{1, \dots, 14\}$

$$fret_t^{\delta, i} = \mu_{ij} + \beta_{ij,1}IMB_j^{a,j} + \beta_{ij,2}IMB_t^{b,j} + \beta_{ij,3}TFI_t^j + \beta_{ij,4}PRET_t^j + \beta_{ij,5}DIV_t^{ij} + \epsilon_{ij,t}, \quad (22)$$

and inspect the resulting coefficients of determination. We fit these models using OLS regression using training data spanning the period from February 22-27, 2021. The data was normalized by subtracting the mean and dividing by the standard deviation. For the sake of readability, we have suppressed in our notation the time horizons over which those features are calculated, with the understanding that the optimal time horizon has been selected. Furthermore, for those features where the nonlinear transformation yielded an increase in linear explanatory power, we will let it be implicit in the notation that the transformed version is used. The models will be fitted for both future returns lookahead windows $\delta \in \{500ms, 1000ms\}$, although the results turn out to be similar in both cases, and we will report results primarily based on the 500ms horizon.

Let us denote the coefficient of determination of the above fitted linear model by R_{ij}^2 . This is the total variation in future returns on market i explained with features from market j , and it will be our first indicator of the lead-lag relationship between the two markets.

We visualize the matrix $R := (R_{ij}^2)_{i,j=1,\dots,14}$ in Figure 11 as a pair of heatmaps, where the first one uses an ordering by column sum, and the second uses an ordering by row sum.

The largest R^2 value, of about 25.9%, is achieved when predicting FTX using features from the Binance USDT perpetual. The smallest R^2 value is about 1% for predicting the Binance USDT perpetual using features of the Bybit USDT perpetual. Overall, we achieve an average R^2 of 9.6% and a median of 9%. The results are similar




Figure 11: Entries in the matrix R illustrating the lead-lag relationships between pairs of markets.

when we use $\delta = 1000\text{ms}$ as our future returns time horizon, where we find average and median values of 10.4% and 9.2%, respectively. This is quite remarkable considering the fact that only information of a single market was used in these regression models. In other words, even information from just a single market can be quite effective in explaining future returns. It demonstrates that it is possible and indeed quite feasible to anticipate future price moves in Bitcoin markets on sub-1s time scales, and it is a reassuring demonstration of the utility of our features.

One can compare the R^2 values we obtain here with those reported in [20] where the authors achieve values of around 60% using order flow imbalance to explain *contemporaneous* returns. Our next set of observations pertains to the hierarchy of markets implied by our results. See Figure 12 for a visualization of the column averages and the row averages. A large column average shows that the market is easily predicted by others (hence a laggard), while a large row average shows that a market is particularly useful in predicting price action on other markets. Notably high are the column averages of the Bybit, FTX, and Bitmex perpetuals. Many of the markets that are especially useful in predicting Bybit and Bitmex are from the exchanges Binance and Huobi, as it can be seen in the above heatmaps when one searches, for instance, in the Bybit row, for the cells corresponding to the Binance USDT perpetual or the Huobi perpetual. Indeed, the top 5 markets in terms of row average are all from Binance and Huobi, providing strong evidence for the central role occupied by these exchanges in the price formation process. Note also that this agrees well with the ranking of markets in terms of volume: the Binance USDT perpetual has the largest average daily volume in the month of February by a sizable margin, as seen in Figure 1.




Figure 12: The left bar plot shows row averages of the matrix R , while the right bar plot shows column averages.

4.2 Accounting for Trading Realities

So far we have assessed the goodness-of-fit of our models only by an in-sample measure, namely their R^2 values. This measure suffers from two complications. First, the R^2 of a model is an inherently dimensionless quantity, and second, it is not clear how an R^2 value translates to the realities of trading. Out-of-sample measures like RMSE or misclassification error can give a more interpretable quantity having a dimension, but they too suffer from the latter complication mentioned above. It is unclear how, for example, an accuracy score maps to PnL. One can conceive of cases where a model has an extremely high accuracy score which maps to a large number of slightly profitable trades, but where a single misclassification results in a disastrous loss that wipes out any prior profits and more. To address these concerns, we propose an alternative out-of-sample way of evaluating goodness-of-fit which is more in line with trading realities. Our measure will be defined as the PnL of a natural trading strategy associated with the model from Eqn. (22). This PnL value will be computed from a synthetic walk-forward on an out-of-sample data set, consisting of the two days following our training period. The training period spanned February 22-27 2021, and the test period includes February 28th and March 1st.

Mapping a linear model to a trading strategy. We now describe how we map each model from Eqn. (22) to a trading strategy. The basic idea is very simple: we buy when the prediction is large (and positive), and sell when the prediction is small (and negative). More precisely, the starting point is the strategy which generates predictions on out-of-sample observations and places a hypothetical buy order at the top ask price when the model predicts a value $> T$. When the prediction is $< -T$, the strategy places a hypothetical sell order. Clearly, this strategy involves the choice of a threshold T . However, there is a relatively canonical choice that we can make here: we set T as a quantile value of the set of predictions that the model produced on the training period. Let us introduce some notation to make this more precise. We denote by M_{ij} the model from Eqn. (22), and let us write $\text{preds}_{\text{in-sample}}(M_{ij})$ for the set of in-sample predictions produced by M_{ij} . Then we define

$$T_{ij} := 95\text{-th percentile of the set } \text{preds}_{\text{in-sample}}(M_{ij}). \quad (23)$$

The strategy in its current state still suffers from (at least) two technical issues which we circumvent with a simple additional constraint on the strategy. The first problem is that when we, for example, place a hypothetical buy order anytime the prediction is $> T_{ij}$, we often obtain bursts of many buy orders over a period of time while the prediction remains large. However, of the trades belonging to such a burst of buy orders often only the first one is actionable since it will impact the market and usually change the top ask price. The second issue is that the PnL calculation becomes more complicated and noisy when we have frequent bursts of buy or sell orders (compared with our alternative approach outlined below). For these reasons, we impose the additional constraint on the strategy that the hypothetical trader's position shall never exceed "one unit". That is, in our sequence of hypothetical trades, a buy order may only be followed by a sell order (and a sell order only by a buy order).

We can summarize the strategy associated to M_{ij} by the following set specifications

1. Initialize the trader's position as $p = 0$.
2. If $p \leq 0$ and the prediction of M_{ij} on an unseen observation is $> T_{ij}$, we execute a hypothetical buy order at the top ask price and set $p = 1$.
3. If $p \geq 0$ and the prediction of M_{ij} on an unseen observation is $< -T_{ij}$, we execute a hypothetical sell order at the top bid price and set $p = -1$.

By chronologically scanning over the out-of-sample period (February 28th and March 1st) in this manner, we obtain a sequence of hypothetical trade prices of the following form

$$\mathcal{S}_{ij} = \left(\dots, p_k^{ij, \text{buy}}, p_{k+1}^{ij, \text{sell}}, p_{k+2}^{ij, \text{buy}}, \dots \right). \quad (24)$$

Next, we will compute various measures of PnL from the above sequence.

Results without execution fees. The first PnL we calculate from the sequence (24) is the simplest one: we compute the trader's total PnL over the test period without accounting for the exchange's execution fees yet.

Note that we already implicitly incorporated the spread component of transaction cost since the trader buys at the top ask price and sells at the top bid price. However, it is also important to observe that oftentimes the top quote (ask or bid) contains only a fleetingly small amount of volume. Not infrequently, this volume is a sub 1000 USD amount and occasionally it is just pennies. This implies that the strategy in its current form is often not scalable.

With this caveat stated, we define our first PnL measure by the following equation

$$\text{PnL}_{1,ij} := \sum_{k=1}^{|\mathcal{S}_{ij}|} \left(\frac{p_{k+1}^{ij,sell}}{p_k^{ij,buy}} - 1 \right) \cdot 10000, \quad (25)$$

which can be regarded as the difference in basis points between the average buy price and the average sell price over the sequence \mathcal{S}_{ij} , multiplied by the total number of trades.

The resulting matrix $\text{PnL}_1 = (\text{PnL}_{1,ij})_{i,j=1,\dots,14}$ is visualized as a pair of heatmaps in Figure 13, where the left heatmap is ordered by row sum and the right heatmap by column sum.




Figure 13: Visualization of the matrix PnL_1 with two orderings, one by row sum and the other by column sum.

Row and column averages are provided in Figure 14. Immediately we are struck by the agreement of the results here with those found in the previous subsection, where our analysis was based on R^2 which, reassuringly, provides evidence of the consistency of our approaches. Markets that yield among the largest PnL values are the Bybit at Bitmex perpetuals, which were previously identified as some of the most predictable ones. The most useful markets in producing high PnL values on other markets are the two Binance perpetuals and the Huobi quarterly futures contract, as well as the Huobi USDT perpetual.




Figure 14: Row and column averages of the matrix PnL_1 .

The mean and median value of the matrix PnL_1 are both approximately 11000. Note that this number is the total number of basis points “accumulated” by the strategy. This basically says that, on average, we accumulate 11000 basis points over the course of the two day test period, or, in other words, we achieve roughly a doubling of the initial capital. Of course, this is not realistic, at least not with any sizable amount of capital. How do these results change when we take execution fees into account? This is the question we will turn to next.

Results with default execution fees. As previously noted, exchanges typically have tiered fee structures whereby highly active traders often pay a far lower execution fee than what the default fee is that novice traders are subject to. Another way of obtaining fee discounts on some exchanges can be by holding a large amount of the exchange’s native cryptocurrency. This applies in particular to Binance and FTX where steep discounts can be obtained this way. We refer the reader to Table 2 where, in our previous remarks on execution cost, we display the default and the lowest possible taker fee per market.

For our second PnL measure, we will assume that the default execution fee applies to the trader on each executed trade. Specifically, we define

$$\text{PnL}_{2,ij} := \text{PnL}_{1,ij} - |\mathcal{S}_{ij}| f_i^{\text{default}}, \quad (26)$$

which is the first PnL measure in (25), adjusted for the default execution fee. We are assuming each trade incurs a fee of f_i^{default} , which is our notation for the default taker fee in basis points on market i . The corresponding matrix will be denoted by $\text{PnL}_2 = (\text{PnL}_{2,ij})_{i,j=1,\dots,14}$ and it is illustrated in Figure 15 along with its column and row averages in Figure 16.

The obtained results are more economically sensible. All PnL values are negative, as one would expect from such a simple model which incorporates information of only one market. The smallest loss is 2819 basis points, while the overall average loss is 18899 basis points.

A striking observation is that markets which were formerly identified as the most predictive are now among those which yield the highest (although still negative) PnL values. The top five markets by row average are all either on Huobi or Binance, while the lowest row average is attained by the FTX perpetual which we formerly noted as one of the most predictable ones.




Figure 15: Visualization of the matrix PnL_2 with two orderings, one by row sum and the other by column sum.

Results with VIP execution fees. The third and final PnL measure is similar to the second one, with the only difference that we use the lowest possible execution fee instead of the default fee. The lowest fee typically requires on the scale of around 500m USD monthly notional volume, and in some cases significant holdings of the exchange’s native cryptocurrency in the trader’s exchange wallet. For these reasons, there are only a small number of wealthy individual traders or significant trading firms that receive this fee. To this end, we define

$$\text{PnL}_{3,ij} := \text{PnL}_{1,ij} - |\mathcal{S}_{ij}| f_i^{\text{VIP}}, \quad (27)$$

which is the first PnL measure adjusted for the lowest possible execution fee f_i^{VIP} on market i . We denote the corresponding matrix by $\text{PnL}_3 = (\text{PnL}_{3,ij})_{i,j=1,\dots,14}$. A visualization can be found in Figure 17. The row and column averages can be seen in the bar plots in Figure 18.




Figure 16: Row and column averages of the matrix PnL_2 .

The surprising implication of these results is that the most leading markets are the ones where the largest PnL can be achieved. For instance, the Binance USDT perpetual has the second largest row average. Conversely, many of the most lagging markets are the ones where the lowest PnL is achieved. An example of this would be the Bitmex perpetual which has the second lowest row average. The Huobi spot market seems to be an outlier, where even when averaged across markets, a positive PnL can be achieved. It should be noted that this market is one with often particularly small top of the book amounts which implies that the capital that can be deployed towards this strategy is particularly small (perhaps a single digit USD amount) since we are assuming execution at the top of the book price.

Overall, in the PnL_3 matrix, we see an average loss of 9621 basis points. The highest value is observed for trading on Huobi spot using information from Binance BTC perpetual, where the PnL is 11426 basis points. The largest loss (of size -42435 basis points) is incurred by the model trading on the Bybit USDT perpetual using features of the Huobi spot market.

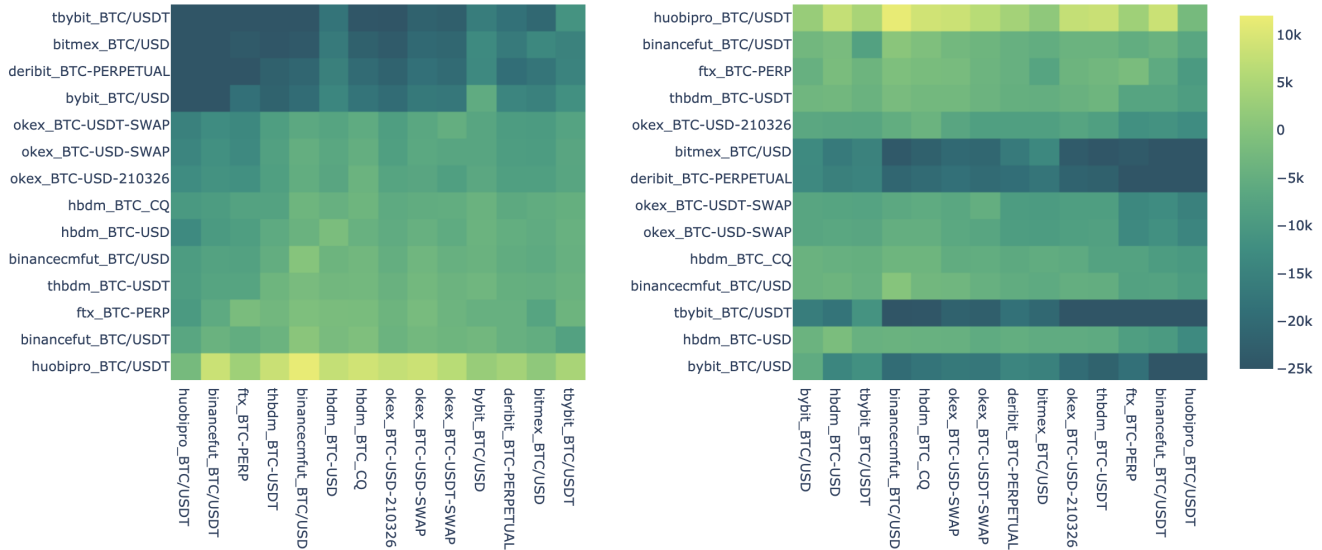


Figure 17: Visualization of the matrix PnL_3 with two orderings, one by row sum and the other by column sum.

Remarks on the Role of Exchange Fees. Exchange fees have been used in studies of price formation, e.g. by [25]. We identified a set of leading markets and a set of lagging markets using an approach based on R^2 values of the models M_{ij} . In our first comparison with the PnL results where we ignored execution fees, we found strong agreement with the R^2 based lead-lag analysis. In line with basic intuition, we found that the most lagging markets




Figure 18: Row and column averages of the matrix PnL_3 .

yielded the largest PnL values. With execution fees accounted for, however, we obtained a very different picture. Some of the most leading markets attained the highest PnL values, while some of the most lagging markets attained among the lowest PnL values. What are some possible explanations of this phenomenon? We will argue that a market's fee regime naturally predisposes it to either being a leader or a lagger (or something in between). This natural tendency in and of itself does not correspond to trading "alpha" (positive PnL). Only when a market exhibits "leadingness" in excess of its natural disposition to being a leader, can we extract alpha to produce positive PnLs on other markets. By the same token, a market's "laggardness" in and of itself only implies alpha (positive PnL) if the lag effect is greater than what is implied by the market's fee regime.

So how does the fee structure on a market correspond to a natural disposition to being a leader or laggard? To this end, let us consider the example of a pair of markets $(1, 2)$ with differing fee structures and reason about how these affect the exchange's lead-lag behaviour in each case. Let us assume that the difference between taker fee and maker rebate is the same on both exchanges. That is, there exists a scalar c such that for $i \in \{1, 2\}$

$$fee_i - rebate_i = c, \quad (28)$$

where fee_i denotes the taker fee on market i and $rebate_i$ denotes the maker rebate on market i . Note that the taker fee is the execution cost in basis points paid for a market order (or a marketable limit order). The maker rebate is the rebate, also in basis points, received by the market maker upon receiving a fill for a passive limit order. The maker rebate can also be negative, in which case the market maker incurs a cost for getting a fill. In practice, one finds a somewhat uniform difference

$$fee - rebate \approx 5, \quad (29)$$

when one compares across exchanges (with VIP traders often enjoying a smaller difference, hence paying less net fees). Eqn. (28) of course does not impose any restrictions on the values of fee_i or $rebate_i$ themselves. Let us suppose that $c = 5$ and that the first market has fee structure $(fee_1, rebate_1) = (10, 5)$ while the second one uses $(fee_2, rebate_2) = (5, 0)$.

An informed market taker (a trader contemplating the placement of a taker order) will rationally place a buy order on market i if and only if they expect an upward price move of at least fee_i basis points over the course of their expected holding period. In other words, the trader sends a taker buy order if and only if the following condition is met

$$\mathbb{E} \left[\left(\frac{p_{t+h,i}}{p_{t,i}} - 1 \right) \cdot 10000 \right] > fee_i, \quad (30)$$

where $p_{t,i}$ is the price on market i at time t , and h is the trader's expected holding time. Since $fee_1 > fee_2$, this implies that the trader requires a stronger signal to send a taker order on market 1 and is therefore disincentivized from sending taker orders. From the perspective of the market maker, the larger rebate $rebate_1$ on market 1 implies a disincentive to cancel passive orders, since a larger amount of adverse selection can be compensated for by the maker rebate on market 1 relative to market 2. The combination of these twin effects implies a natural tendency of market 1 to be a laggard. In the limiting case $fee_1, rebate_1 \rightarrow \infty$, the orderbook would be completely static with

no transactions ever occurring as market makers can tolerate an arbitrary amount of adverse selection, and takers require an infinitely strong price signal to overcome the taker fee. The conclusion therefore is that market 2 is naturally disposed toward lagging behaviour when compared to market 1 (which is consequently a leader in this comparison).

This line of reasoning leads us to conclude that the extent to which a market is a leader or a lagger does not, in and of itself, have any implications regarding trading "alpha" that can be achieved for trading on said market or for utilizing its information to trade on other markets. The reality is more subtle due to potentially different fee regimes (even when the difference between taker fee and maker rebate is uniform across exchanges). Only when a market exhibits lagging behaviour *in excess* of the laggardness to which it is naturally predisposed by its fee structure do we obtain trading "alpha" that can be exploited to generate positive PnL values on the market in question. For instance, counter-intuitively, the Binance USDT perpetual, while generally a strong leading market, is nevertheless sufficiently predictable, owing to its fee structure, to give rise to positive PnL values.

5 Cross-Sectionally Combining Information

Only by taking into account information from the full set of relevant markets, can one hope to reliably make accurate predictions on short term order flow, and hence price action on any given market. Evidence in support of this assertion was provided in the previous subsection, where we observed that a single market's features are insufficient to produce "alpha" in excess of execution cost on another market (with rare exceptions for VIP traders that receive heavy fee discounts). The underlying explanation is that market makers, who are responsible for the bulk of liquidity in the orderbook, continuously adjust their quotes on the basis of *all* available information sources which consistently bear relevance for price moves. For if this were not the case, they would be vulnerable to being picked off by quick arbitrageurs who are "on the prowl" for resting limit orders which are no longer in line with the rest of the market. For instance, suppose the price moved up substantially on all markets except one. The market participant who continues to post a passive sell order on the market where a price move up has not yet occurred is apparently ignorant to the new information dispersed among other markets, and will no doubt shortly incur a loss by having his order consumed by an arbitrageur. Hence the market makers that "survive" (the profitable ones that do not consistently incur such losses), and therefore ultimately dictate the bulk of the liquidity in the orderbook and thus also the price, are precisely the ones that are capable of parsing all available relevant information so as to prevent being adversely selected.

Therefore, if we keep in mind our objective of predicting order flow and price change on sub-second time scales, we ought to search for ways to combine information across markets to achieve greater predictive power. One natural starting point is to simply fit linear regression models (as before with OLS) using the full set of features from all markets instead of just a single market. This will be the starting point and a baseline for comparison of our analysis in this subsection. One concern with this approach is that the model might be overfitted due to a significant number of irrelevant features and/or a strong degree of collinearity between features. We will therefore compare the baseline approach with the following two alternative approaches.

The first alternative approach is to deal with the large number of features by using L^1 -regularization in the regression model. That is, instead of fitting a simple OLS model on the full feature set, we will train a LASSO regression model which ends up selecting a subset of the features.

The second alternative approach involves linearly combining features across the markets to form a set of "meta features" which are subsequently leveraged to generate parsimonious models for price prediction on each market. The linear combinations will be formed using coefficients that are proportional to predictive power, as measured by R^2 .

5.1 Baseline Approach

Model specification. Our baseline model consists of a linear regression trained on the full feature set. Each market has 5 features and there are a total number of 14 markets, resulting in a total of 70 covariates. Formally, we use OLS regression to fit the following linear models for all $i \in \{1, \dots, 14\}$ and $\delta \in \{500\text{ms}, 1000\text{ms}\}$

$$fret_t^{\delta,i} = \mu_i + \sum_{j=1}^{14} \left(\beta_{i,j,1} IMB_t^{a,j} + \beta_{i,j,2} IMB_t^{b,j} + \beta_{i,j,3} TFI_t^j + \beta_{i,j,4} PRET_t^j + \beta_{i,j,5} DIV_t^{ij} \right) + \epsilon_{i,t}. \quad (31)$$

As usual, to keep notation simple, we let it be implicit that each feature is dependent on the choice of a time horizon, using the one that was previously deemed optimal for it. It is also implicit in the notation that we use the

Table 5: In-sample and out-of-sample R^2 values for baseline models

	R^2		R^2_{oos}	
	500ms	1000ms	500ms	1000ms
ftx_BTC-PERP	0.364	0.342	0.273	0.264
bybit_BTC/USD	0.294	0.331	0.289	0.32
thbdm_BTC-USDT	0.273	0.231	0.25	0.215
hbdm_BTC-USD	0.229	0.262	0.214	0.243
huobipro_BTC/USDT	0.219	0.136	0.191	0.126
bitmex_BTC/USD	0.217	0.292	0.193	0.27
tbybit_BTC/USDT	0.214	0.279	0.219	0.28
binancecmfut_BTC/USD	0.209	0.193	0.187	0.177
deribit_BTC-PERPETUAL	0.187	0.212	0.178	0.214
hbdm_BTC_CQ	0.186	0.152	0.171	0.144
okex_BTC-USDT-SWAP	0.173	0.196	0.156	0.184
okex_BTC-USD-210326	0.17	0.164	0.158	0.155
okex_BTC-USD-SWAP	0.164	0.18	0.151	0.167
binancefut_BTC/USDT	0.106	0.071	0.094	0.069

transformed feature rather than its original version, whenever an improvement was found for the feature. Let us denote the model corresponding to Eqn. (31) by $M_{i,\delta}^{\text{baseline}}$.

Results. We now examine the coefficients of determination of the models $M_{i,\delta}^{\text{baseline}}$ for $i = 1, \dots, 14$ and $\delta \in \{500\text{ms}, 1000\text{ms}\}$. We compute this number both in-sample and out-of-sample. In other words, we report the portion of total variation in future returns explained by each model on the training period, as well as on the two-day test period following the training period. The results can be found in Table 5. Significance values and p-values for all models are reported in Tables 19 and 20 in the Appendix.

We observe that the difference between the two time horizons 500ms and 1000ms is generally rather small. As one would expect, the 500ms time horizon is somewhat more easily predictable. On average, the in-sample R^2 values for this horizon are 4.6% larger than their 1000ms counterparts. When we perform the same computation on the out-of-sample values, we find that the R^2 on the 500ms horizon is, on average, greater than that the 1000ms R^2 by 1.4%. There are, however, some noteworthy outliers. Notably larger values on the 500ms horizon are found on the Binance USDT perpetual and the Huobi spot market, while the Bitmex and Bybit USDT perpetuals exhibit considerably larger R^2 values for the 1000ms horizons. What are some possible explanations of these phenomena?

When a market's 1000ms R^2 is much greater than the 500ms R^2 it means that the features at time t contain information that is not yet reflected in the price at time $t + 500\text{ms}$ but is priced in at $t + 1000\text{ms}$. This suggests the possibility that cross-sectional arbitrages can survive for more than 500ms between a strong leading market and a strong lagging market. Note that this agrees well with our previous identification of Bitmex and Bybit as lagging exchanges, and Binance and Huobi as leading exchanges.

Conversely, when the 500ms R^2 is much larger than its 1000ms counterpart on a given market, this indicates that this market is prone to the arrival of new information in the time window $t + 500\text{ms}$ to $t + 1000\text{ms}$ which is not yet encapsulated in the feature values at time t . We would predict this to be the case for strongly leading markets, as indeed appears to be the case when we compare the R^2 values with our previous findings on the leader-lagger network.

Our next observation pertains to the comparison between in-sample and out-of-sample values. On average, we find that on the 500ms time horizon, the out-of-sample R^2 is around 91% of the in-sample one. The largest discrepancy, both in relative and absolute terms, can be seen for the FTX perpetual where the out-of-sample R^2 is 25% lower than the in-sample R^2 . Overall, the small to moderate loss in accuracy as we pass from the in-sample to the out-of-sample values ameliorates concerns of overfitting to an extent.

Next, we note that by using a market's actual "leadingness" or "laggardness", as measured by R^2 , in conjunction with its fee structure, one can form a set of expectations about what PnL values could be obtained. See Figure 19 for a visual overview of each market's lowest possible taker fee (usually accessible only to highly active traders). As we reasoned in the previous section, we would a-priori expect a market with a low taker fee to be a leading market, which should express itself in a low R^2 value. If we find that a market which "should" be a leader based on its fee structure is in fact a lagger, we would predict a high PnL value for this market. When we compare Figure 19 with Table 5 we observe that, remarkably, the FTX perpetual has both the highest R^2 (indicating that its returns can be




Figure 19: Lowest possible (across tiers) taker fee per market (in basis points).

anticipated well) and the lowest taker fee. On this basis, one would predict a large PnL value for those VIP traders with access to this fee.

The Bybit BTC perpetual is among the most lagging markets, although this is in agreement with the fact that it has the highest taker fee, so we would not necessarily foresee a high PnL value for this market, unless its lagging nature is so significant that it outweighs the large fee.

On the other end of the spectrum, the most leading market, the Binance USDT perpetual, also has one of the lowest taker fees, so it is not straightforward to predict what the PnL might be on this market. Whether a trader subject to the lowest fee is capable of producing a positive PnL boils down to the question of whether the general leading nature of this market can be overcome by its low taker fee.

In Table 6 we display the PnL values produced by the trading strategies corresponding to the models $M_{i,500ms}^{baseline}$ for $i = 1, \dots, 14$. The values are computed by the same procedure we introduced in Subsection 4.2. That is, we map a model to a natural trading strategy and calculate the PnL of this strategy in a synthetic walk-forward over the out-of-sample period. As before, we compute the PnL in three different ways: (1) ignoring execution fees, (2) adjusting for the default execution fee, (3) adjusting for the lowest possible ("VIP") execution fee. The PnL formulae are precise analogues of the ones from Equations (25), (26), and (27). Note that, as before, the PnL therefore tells us what the total number of accumulated basis points over the two day test period is.

In line with the expectations we formed on the basis of predictability and fee regime, we find that FTX gives the largest PnL_3 value. Somewhat surprisingly, even the most leading market, the Binance USDT perpetual, exhibits a positive PnL_3 value. This can be interpreted as a suggestion that the extremely low taker fee is sufficient to "overpower" the market's relative lack in predictability. The lowest PnL is found on Deribit which has one of the highest taker fee and middling R^2 values. Similarly, Bitmex has slightly larger R^2 values but also a larger taker fee, so it is not surprising to see it having the second smallest PnL_3 value.

Note that all PnL_2 numbers are negative. This hints to "the rich get richer" phenomena in the crypto markets, whereby traders with access to the lowest taker fee have a much greater ability to turn profits than novice traders. Since a trader's fee regime is determined by trade volume, and in some cases, holdings in the exchange's native cryptocurrency, the privilege of being in the best fee tier clearly correlates highly with the trader's net worth. On the other hand, a more novice trader wishing to ascend to the best fee tier in order to boost profits might have to "burn through" some losses in lower fee tiers until they have accumulated enough trading volume to be granted a better fee.

5.2 LASSO Regression Approach

The models we trained in the preceding subsection used a total of 70 features from 14 different markets. Many of these features exhibit high cross-correlations, as evidenced in Table 7 where we display the cross-correlations amongst the 500ms past returns features across markets.

In this subsection, we will pursue an alternative methodology for fitting the parameters of Eqn. (31) based on *LASSO regression*. The difference between an OLS fit and a LASSO fit resides in an additional regularization term

Table 6: PnLs of baseline models

	PnL ₁	PnL ₂	PnL ₃
ftx_BTC-PERP	19792.4	-27527.6	9652.4
huobipro_BTC/USDT	28844.0	-28251.0	5645.4
binancecmfut_BTC/USD	16519.1	-22010.9	2648.3
binancefut_BTC/USDT	12107.6	-17724.4	696.8
thbdm_BTC-USDT	20784.9	-10703.1	-469.5
hbdm_BTC-USD	17703.1	-8186.9	-1455.5
hbdm_BTC_CQ	16432.0	-23400.0	-3484.0
okex_BTC-USD-SWAP	18041.0	-19319.0	-4375.0
okex_BTC-USD-210326	17436.4	-27273.6	-4918.6
okex_BTC-USDT-SWAP	15012.4	-20097.6	-6053.6
bybit_BTC/USD	19811.2	-11508.8	-11508.8
tbybit_BTC/USDT	17036.4	-15828.6	-15828.6
bitmex_BTC/USD	17804.5	-16140.5	-16140.5
deribit_BTC-PERPETUAL	10401.1	-19738.9	-19738.9

Table 7: Correlation matrix of 500ms past returns. For brevity we used the market acronyms from Table 17.

	m1	m2	m3	m4	m5	m6	m7	m8	m9	m10	m11	m12	m13	m14
m1	1.0	0.5	0.4	0.4	0.5	0.4	0.5	0.5	0.4	0.5	0.5	0.4	0.6	0.6
m2	0.5	1.0	0.3	0.4	0.4	0.6	0.5	0.6	0.4	0.6	0.5	0.3	0.6	0.6
m3	0.4	0.3	1.0	0.3	0.4	0.2	0.3	0.3	0.3	0.3	0.3	0.3	0.4	0.3
m4	0.4	0.4	0.3	1.0	0.4	0.3	0.4	0.4	0.3	0.4	0.3	0.3	0.4	0.4
m5	0.5	0.4	0.4	0.4	1.0	0.3	0.4	0.4	0.4	0.4	0.4	0.4	0.5	0.4
m6	0.4	0.6	0.2	0.3	0.3	1.0	0.4	0.5	0.3	0.7	0.4	0.2	0.6	0.6
m7	0.5	0.5	0.3	0.4	0.4	0.4	1.0	0.5	0.4	0.5	0.5	0.3	0.6	0.6
m8	0.5	0.6	0.3	0.4	0.4	0.5	0.5	1.0	0.4	0.5	0.5	0.3	0.6	0.6
m9	0.4	0.4	0.3	0.3	0.4	0.3	0.4	0.4	1.0	0.4	0.4	0.3	0.5	0.5
m10	0.5	0.6	0.3	0.4	0.4	0.7	0.5	0.5	0.4	1.0	0.4	0.2	0.6	0.6
m11	0.5	0.5	0.3	0.3	0.4	0.4	0.5	0.5	0.4	0.4	1.0	0.3	0.5	0.5
m12	0.4	0.3	0.3	0.3	0.4	0.2	0.3	0.3	0.3	0.2	0.3	1.0	0.4	0.3
m13	0.6	0.6	0.4	0.4	0.5	0.6	0.6	0.6	0.5	0.6	0.5	0.4	1.0	0.7
m14	0.6	0.6	0.3	0.4	0.4	0.6	0.6	0.6	0.5	0.6	0.5	0.3	0.7	1.0

Table 8: Number of surviving coefficients, as we vary the λ regularization coefficient.

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
deribit_BTC-PERPETUAL	67	64	61	53	44	36	25	9	4
binancefut_BTC/USDT	67	65	62	54	43	26	16	6	1
ftx_BTC-PERP	67	64	57	50	45	40	26	14	6
huobipro_BTC/USDT	67	67	61	51	41	28	21	9	2
binancecmfut_BTC/USD	67	62	59	49	35	27	17	9	2
hbdm_BTC-USD	66	64	56	43	34	27	15	6	6
okex_BTC-USD-SWAP	66	62	53	43	38	33	20	10	5
okex_BTC-USDT-SWAP	66	64	58	47	41	35	21	9	4
thbdm_BTC-USDT	66	61	56	51	33	20	14	10	3
hbdm_BTC_CQ	65	63	58	52	32	26	22	6	3
okex_BTC-USD-210326	65	61	58	48	37	30	24	9	2
bitmex_BTC/USD	62	60	54	48	38	28	22	15	4
bybit_BTC/USD	62	60	55	33	28	23	17	9	5
tbybit_BTC/USDT	62	61	51	43	39	26	19	10	5

introduced in the objective function of the latter. Specifically, OLS regression minimizes the standard least-squares objective function, while LASSO regression minimizes the regularized objective function

$$\min_{w \in \mathbb{R}^p} \|y - \mathbf{X}w\|_2^2 + \lambda \|w\|_1, \quad (32)$$

where $\mathbf{X} \in \mathbb{R}^{n \times p}$, $y \in \mathbb{R}^n$ and $\lambda \in \mathbb{R}_{>0}$ is the regularization parameter. The additional ℓ_1 regularization term has the effect of zeroing a subset of the parameters. This subset grows in size as the regularization parameter λ is increased. For a full discussion on LASSO regression, the reader is referred to [22].

We fit the coefficients in Eqn. (31) using a LASSO regression with regularization parameters from the set $\{0.001 \cdot 2^k | k = 0, \dots, 8\}$. The lower bound $\lambda = 0.01$ of this set was chosen since it was found to be the smallest (round) number which, when used as the regularization parameter in the LASSO regression, results in a positive number of zero coefficients for each market’s model. Similarly, the upper bound $\lambda = 0.256 = 0.01 \cdot 2^8$ was selected since it was the largest doubling of the lower bound $\lambda = 0.01$ for which a positive number of nonzero coefficients persist.

For each market and each of the nine regularization parameters, we then examine the number of nonzero coefficients, their in-sample and out-of-sample explanatory power over future returns, and the PnL values of their naturally associated trading strategies. Furthermore, we will also inspect exactly which coefficients tend to survive as λ is increased. Comparisons of the results in this subsection with the results of baseline approach are deferred until we introduce our third alternative methodology for training powerful linear models to explain short-term future returns.

Model Selection Results. First, we inspect the number of nonzero coefficients of each of the LASSO fits, as the regularization parameter λ is increased. The results can be seen in Table 8. With $\lambda = 0.001$, the majority of the 70 coefficients are found to be nonzero. In line with expectation, as λ is increased the number of surviving coefficients decreases. With our largest choice of $\lambda = 0.256$, we typically do not retain more than a handful of features in the model.

Our next set of observations pertains to the explanatory power of each of the LASSO models. It would be reasonable to hypothesize that perhaps explanatory power (particularly out-of-sample) might increase as we discard superfluous features in a LASSO model, via a suitable regularization parameter. However, this turns out to be the case only very rarely, as evidenced in Tables 9 and 10, where we show in-sample and out-of-sample (respectively) R^2 values for each model. In fact, R^2 seems to generally be monotonically decreasing in the regularization parameter λ . That said, the decrease happens quite slowly, especially for $\lambda \leq 0.032$. For example, the LASSO models with $\lambda = 0.016$ retain on average 97.2% of the out-of-sample R^2 compared to the LASSO models with $\lambda = 0.001$. When we compare the $\lambda = 0.001$ models with the $\lambda = 0.032$ ones, we still find an average decrease in out-of-sample R^2 of only 5.9%. Since the number of features retained in the latter models ranges between 20 and 40, we can conclude therefore that only around half of the original 70 features are sufficient to retain the vast majority of explanatory power. When we increase the regularization parameter to $\lambda = 0.064, 0.128$ and 0.256 , we lose on average 13.3%, 25.97%, and 39% of explanatory power relative to the $\lambda = 0.001$ models, respectively. Both of the aforementioned

Table 9: In-sample R^2 values for LASSO models using 500ms future return horizon, as we vary the λ regularization coefficient.

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
ftx_BTC-PERP	0.364	0.364	0.363	0.362	0.36	0.354	0.335	0.291	0.237
bybit_BTC/USD	0.293	0.293	0.291	0.289	0.286	0.283	0.27	0.237	0.131
thbdm_BTC-USDT	0.273	0.273	0.272	0.269	0.262	0.252	0.234	0.203	0.107
hbdm_BTC-USD	0.229	0.229	0.227	0.224	0.221	0.212	0.202	0.177	0.092
huobipro_BTC/USDT	0.219	0.219	0.218	0.216	0.21	0.201	0.177	0.108	0.064
bitmex_BTC/USD	0.217	0.216	0.215	0.214	0.212	0.206	0.192	0.165	0.088
tbybit_BTC/USDT	0.214	0.214	0.213	0.212	0.21	0.205	0.191	0.168	0.1
binancecmfut_BTC/USD	0.209	0.209	0.208	0.206	0.202	0.193	0.171	0.141	0.076
deribit_BTC-PERPETUAL	0.187	0.187	0.186	0.184	0.182	0.177	0.163	0.121	0.069
hbdm_BTC_CQ	0.185	0.185	0.185	0.182	0.179	0.173	0.154	0.108	0.058
okex_BTC-USDT-SWAP	0.173	0.173	0.172	0.171	0.169	0.164	0.148	0.117	0.07
okex_BTC-USD-210326	0.17	0.17	0.17	0.168	0.165	0.161	0.144	0.1	0.059
okex_BTC-USD-SWAP	0.164	0.164	0.163	0.162	0.16	0.154	0.138	0.108	0.059
binancefut_BTC/USDT	0.106	0.106	0.105	0.103	0.096	0.083	0.058	0.036	0.002

Table 10: Out-of-sample R^2 values for LASSO models using 500ms future return horizon, as we vary the λ regularization coefficient.

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
bybit_BTC/USD	0.289	0.289	0.287	0.283	0.281	0.279	0.272	0.263	0.236
ftx_BTC-PERP	0.273	0.272	0.272	0.271	0.269	0.265	0.25	0.211	0.193
thbdm_BTC-USDT	0.25	0.25	0.249	0.247	0.242	0.234	0.223	0.211	0.149
tbybit_BTC/USDT	0.219	0.219	0.218	0.217	0.216	0.212	0.202	0.194	0.167
hbdm_BTC-USD	0.213	0.213	0.212	0.21	0.206	0.198	0.193	0.188	0.176
bitmex_BTC/USD	0.193	0.193	0.192	0.192	0.191	0.188	0.181	0.171	0.141
huobipro_BTC/USDT	0.191	0.19	0.19	0.188	0.183	0.176	0.159	0.095	0.073
binancecmfut_BTC/USD	0.187	0.187	0.186	0.184	0.181	0.173	0.153	0.136	0.101
deribit_BTC-PERPETUAL	0.178	0.178	0.178	0.177	0.176	0.173	0.167	0.143	0.122
hbdm_BTC_CQ	0.171	0.171	0.17	0.168	0.164	0.16	0.146	0.107	0.082
okex_BTC-USD-210326	0.158	0.158	0.158	0.156	0.154	0.151	0.14	0.103	0.083
okex_BTC-USDT-SWAP	0.156	0.156	0.156	0.155	0.153	0.15	0.138	0.115	0.096
okex_BTC-USD-SWAP	0.151	0.151	0.15	0.149	0.148	0.144	0.131	0.111	0.09
binancefut_BTC/USDT	0.094	0.094	0.093	0.092	0.087	0.076	0.051	0.038	0.026

tables were created using the 500ms future returns horizon. For the 1000ms versions, we refer to Tables 21 and 22 in the Appendix. The results are very similar in both cases.


Next, we examine the profits and losses attained by the LASSO models. In Tables 23 and 24 of the Appendix, we show the PnL_1 and PnL_2 values (respectively) for each model. Most interesting are the PnL_3 values which we display in Table 11. We note that in most cases the maximal PnL is achieved with either $\lambda = 0.001$ or $\lambda = 0.002$, although the difference between the maximum PnL and any other PnL with $\lambda \leq 0.032$ is marginal. This is in agreement with our prior observation that explanatory power is almost identical between all models with $\lambda \leq 0.032$. The only markets where the maximal PnL is not attained with $\lambda \in \{0.001, 0.002\}$ are the Huobi spot market and the Binance USDT perpetual. The former achieves its maximum with $\lambda = 0.128$ and nine features while the latter does so with $\lambda = 0.032$ and 26 features.

We have seen that the PnL and explanatory power of a market's LASSO models is almost unchanged for any regularization parameters $\lambda \leq 0.032$. It is interesting to examine which of the 70 original features are retained in these models. In Figure 20a we provide an illustration of the number of markets whose $\lambda = 0.032$ model uses the features listed on the x -axis. For ease of legibility, we include in this plot only those features which appear in the models of at least six markets.

We are immediately struck in Figure 20b that the features used by any of the $\lambda = 0.256$ models are *exclusively* from a market on Binance or Huobi. This underlines the outsized role that these two exchanges play in price

Table 11: PnL₃ values of LASSO models as regularization parameter varies

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
ftx_BTC-PERP	9576	9439	9371	9253	9444	9407	9370	8847	6987
huobipro_BTC/USDT	5738	5679	5654	5637	5610	6196	7264	8908	3598
binancecmfut_BTC/USD	2576	2497	2229	1802	1991	1855	1511	-179	-2188
binancefut_BTC/USDT	766	824	844	1274	1627	2335	2095	1340	1297
thbdm_BTC-USDT	-502	-516	-732	-1013	-1246	-1756	-2527	-3722	-8110
hbdm_BTC-USD	-1517	-1572	-1648	-1551	-1816	-2104	-2483	-3197	-4461
hbdm_BTC_CQ	-3579	-3669	-3711	-4104	-4232	-4061	-4129	-5479	-9427
okex_BTC-USD-SWAP	-4380	-4376	-4603	-4979	-5314	-5929	-6544	-7829	-12021
okex_BTC-USD-210326	-4951	-5006	-5327	-5784	-5966	-6776	-7439	-8122	-11823
okex_BTC-USDT-SWAP	-6050	-6105	-6418	-6770	-6951	-7605	-7752	-8334	-11518
bybit_BTC/USD	-11631	-11440	-11504	-11808	-11895	-12510	-13771	-15709	-19746
tbybit_BTC/USDT	-15855	-15917	-16264	-16521	-16601	-17025	-17592	-19935	-25232
bitmex_BTC/USD	-16209	-16091	-16341	-16386	-16311	-15813	-16328	-18577	-23969
deribit_BTC-PERPETUAL	-19666	-19621	-19784	-20189	-20368	-20438	-20441	-20387	-25174



(a) Regularization $\lambda = 0.032$. Only features with a count of at least seven are included for ease of legibility.

(b) Regularization $\lambda = 0.256$

Figure 20: Count of nonzero features

discovery. Let us next inspect Figure 20a. Here we observe the same dominance of Binance and Huobi with notably large counts also occurring from Okex markets. We note that the most common type of feature is a (transformed) trade flow imbalance feature. This suggests signed volume is the most powerful of our indicators. Of the orderbook imbalance features, the one corresponding to the Huobi BTC perpetual is the most commonly selected. The most frequent *DIV* feature comes from the Huobi quarterly futures contract, which indicates that the price difference with this market is particularly useful to predict price action.

5.3 Meta Features Approach

In the previous subsection, we found that a much smaller subset retaining approximately 30 of the original 70 features used in the baseline model is sufficient to produce models whose explanatory power and PnL values are only marginally lower than those of the baseline models. In this subsection, we pursue an alternative methodology for drastically reducing dimensionality of our linear models. The procedure involves an additional feature processing step, which cuts down the number of features from 70 to five. The resulting models will turn out to achieve a greater average PnL than the baseline models, even though average explanatory power is diminished.

We observed in Table 7 that past returns features exhibit large cross-correlations across exchanges. This is of course not surprising, since the markets we are considering are either Bitcoin spot markets themselves or derivative contracts whose underlying is an index composed of Bitcoin spot markets. Motivated by this consideration, we devise a methodology by which we form a single “meta” past returns feature from the set of all individual past

returns features. That is, we will define for each market $i \in \{1, \dots, 14\}$ a meta feature $mPRET_i$ using

$$mPRET_t^i := \sum_{j=1}^{14} \alpha_{ij}^{PRET} PRET_t^j, \quad (33)$$

where the coefficient α_{ij}^{PRET} is chosen in proportion to the predictive power that the past returns feature of market j has over the future returns of market i . Other meta features (orderbook imbalances, trade flow imbalance, mean divergence) are computed analogously. More specifically, we define a set of so-called *meta features* for each market $i \in \{1, \dots, 14\}$ by the equations

$$\begin{aligned} mIMB_t^{a,i} &:= \sum_{j=1}^{14} \alpha_{ij}^{IMB^a} IMB_{jt}^{a,j}, & mIMB_t^{b,i} &:= \sum_{j=1}^{14} \alpha_{ij}^{IMB^b} IMB_{jt}^{b,j}, \\ mTFI_t^i &:= \sum_{j=1}^{14} \alpha_{ij}^{TFI} TFI_t^j, & mDIV_t^i &:= \sum_{j=1}^{14} \alpha_{ij}^{DIV} DIV_t^{ij}, \end{aligned} \quad \text{where}$$

the coefficients are determined by the procedure described below. We will demonstrate this process on the example of the past returns feature. Let us fix $i \in \{1, \dots, 14\}$ and $\delta \in \{500\text{ms}, 1000\text{ms}\}$. The steps are then as follows

1. For every $j \in \{1, \dots, 14\}$, we train the univariate models $fret_t^{\delta,i} = \mu_{ij} + \beta_{ij} PRET_t^j + \epsilon_{ij,t}$ using an OLS regression. We denote its coefficient of determination by R_{ij}^2 .
2. We set $\alpha_{ij}^{PRET} := \frac{R_{ij}^2}{\max_j R_{ij}^2}$.

The procedure is completely analogous for other features. Our normalization procedure involves division by a maximum value. Note that there are other possible weight normalizations, such as division by the $\sum_j R_{ij}^2$; we leave it for future work to compare these approaches.

Now that we have defined the set of meta features, we can define the following linear models for every market $i \in \{1, \dots, 14\}$ and the two future returns time horizons $\delta \in \{500\text{ms}, 1000\text{ms}\}$

$$fret_t^{\delta,i} = \mu_i + \beta_{i,1} mIMB_t^{a,j} + \beta_{i,2} mIMB_t^{b,j} + \beta_{i,3} mTFI_t^i + \beta_{i,4} mPRET_t^i + \beta_{i,5} mDIV_t^i + \epsilon_{i,t} \quad (34)$$

We fit these models using an OLS regression and analyze their performance along the same axes we have previously considered (explanatory power and PnL).

Results. One main benefit of the meta feature models is their simplicity. Compared to the baseline models, we have cut down the number of covariates from 70 to five. In Table 12, we report t-statistics and p-values for each of the coefficients in all of the 14 models we trained. We find strong statistical significance for all meta features in each of the models. The smallest absolute t-statistic corresponds to the meta past returns feature for the model predicting future returns on the Okex quarterly futures contract. It is perhaps not too surprising that we would see a comparably small t-statistic in this case, since the futures premium can fluctuate somewhat freely in the sense that it is not linked by any strong arbitrage bounds to the price of its underlying. This argument applies more strongly the more distant the expiration of the futures contract is (note that in our case the expiration date lays approximately one month in the future). The coefficients of determination for the meta models defined by Eqn. (34) are given in Table 13, where we show the in-sample and out-of-sample R^2 values for the two time horizons under consideration. Here we find that a similar set of observations holds true as in the baseline and LASSO models.

First, the average difference between the 500ms and 1000ms R^2 values tends to be quite small. For example, in case of the in-sample R^2 values we observe that the 500ms version is on average 7% larger than the 1000ms one, in line with the expectation that shorter time horizons would be more easily predicted. Some outliers are the lagging markets Bybit USDT perpetual and Bitmex perpetual, where passing from the 500ms to the 1000ms future returns horizon yields an increased R^2 of 32.6% and 37%, respectively. On the other hand, markets that see a notably high decrease in accuracy in passing from the 500ms to the 1000ms horizon are leading markets such as the Binance USDT perpetual or the Huobi spot market, where we see decreases of 39.7% and 40.5%, respectively. A similar phenomenon was noted for the baseline and LASSO models.

Secondly, the in-sample R^2 is not significantly larger than its out-of-sample counterpart, which provides evidence of the absence of significant overfitting. We find average decreases of 10.7% for the 500ms horizon and 4.9% for the 1000ms horizon.

In Table 14, we report PnL values for the meta models (34). As before, we compute three separate values: PnL₁ which is the total number of basis points accumulated over the test period by the strategy associated with the respective meta model, PnL₂ where we adjust PnL₁ by the default fee, and PnL₃ where we adjust PnL₁ by the VIP

Table 12: T-statistics and p-values for meta models

	t-stat					p-value				
	mDIV	mIMBa	mIMBb	mTFI	mPRET	mDIV	mIMBa	mIMBb	mTFI	mPRET
hbdm_BTC-USD	28.85	53.29	-24.7	64.47	47.59	0.0	0.0	0.0	0.0	0.0
hbdm_BTC_CQ	-9.37	44.81	-12.37	98.57	-16.33	0.0	0.0	0.0	0.0	0.0
bitmex_BTC/USD	44.19	37.24	-16.03	37.36	73.89	0.0	0.0	0.0	0.0	0.0
deribit_BTC-PERPETUAL	28.31	41.18	-12.16	50.14	51.19	0.0	0.0	0.0	0.0	0.0
bybit_BTC/USD	54.91	34.91	-31.58	50.55	74.97	0.0	0.0	0.0	0.0	0.0
binancefut_BTC/USDT	31.76	33.36	-6.21	104.06	-35.93	0.0	0.0	0.0	0.0	0.0
okex_BTC-USD-SWAP	-6.72	48.1	-24.71	62.38	20.09	0.0	0.0	0.0	0.0	0.0
okex_BTC-USD-210326	-13.5	38.55	-13.5	79.51	2.11	0.0	0.0	0.0	0.0	0.04
ftx_BTC-PERP	53.45	40.31	-10.22	77.34	95.26	0.0	0.0	0.0	0.0	0.0
huobipro_BTC/USDT	47.64	27.6	-5.86	110.8	-30.58	0.0	0.0	0.0	0.0	0.0
okex_BTC-USDT-SWAP	-12.02	46.86	-21.47	57.19	29.97	0.0	0.0	0.0	0.0	0.0
tbybit_BTC/USDT	28.17	34.46	-21.09	26.64	88.36	0.0	0.0	0.0	0.0	0.0
thbdm_BTC-USDT	29.99	57.25	-20.41	108.79	18.96	0.0	0.0	0.0	0.0	0.0
binancecmfut_BTC/USD	19.76	60.37	-22.4	89.62	15.52	0.0	0.0	0.0	0.0	0.0

Table 13: In-sample and out-of-sample R^2 values for the meta models.

	R^2		R^2_{oos}	
	500ms	1000ms	500ms	1000ms
ftx_BTC-PERP	0.306	0.288	0.213	0.212
bybit_BTC/USD	0.262	0.3	0.261	0.294
thbdm_BTC-USDT	0.215	0.18	0.202	0.173
hbdm_BTC-USD	0.201	0.229	0.19	0.217
bitmex_BTC/USD	0.192	0.263	0.174	0.249
tbybit_BTC/USDT	0.19	0.252	0.198	0.26
binancecmfut_BTC/USD	0.156	0.139	0.137	0.13
deribit_BTC-PERPETUAL	0.141	0.166	0.152	0.186
okex_BTC-USDT-SWAP	0.13	0.144	0.117	0.139
huobipro_BTC/USDT	0.126	0.075	0.106	0.069
okex_BTC-USD-SWAP	0.122	0.131	0.112	0.124
hbdm_BTC_CQ	0.116	0.091	0.104	0.09
okex_BTC-USD-210326	0.112	0.105	0.105	0.102
binancefut_BTC/USDT	0.068	0.041	0.055	0.041

Table 14: PnLs of meta models

	PnL ₁	PnL ₂	PnL ₃
huobipro_BTC/USDT	21220.4	-6377.1	10007.1
ftx_BTC-PERP	17140.6	-19357.4	9319.6
binancefut_BTC/USDT	10790.5	-12097.5	2035.8
binancecmfut_BTC/USD	10085.4	-13284.6	1672.2
thbdm_BTC-USDT	12393.6	-6910.4	-636.6
hbdm_BTC_CQ	8051.1	-10916.9	-1432.9
hbdm_BTC-USD	15057.7	-8112.3	-2088.1
okex_BTC-USD-210326	8739.4	-15220.6	-3240.6
okex_BTC-USD-SWAP	10398.4	-13091.6	-3695.6
okex_BTC-USDT-SWAP	9803.9	-13946.1	-4446.1
bybit_BTC/USD	19613.6	-14451.4	-14451.4
deribit_BTC-PERPETUAL	9116.7	-14623.3	-14623.3
bitmex_BTC/USD	17561.6	-17433.4	-17433.4
tbybit_BTC/USDT	17659.5	-18580.5	-18580.5

fee. As before with the other models, we find particularly large PnL_3 values for FTX where the fee structure implies that it “should be” a leader, but its large R^2 value points at lagging behaviour. We also note that the Binance USDT perpetual yields a positive PnL_3 value despite the fact that it has (by quite a margin) the lowest predictability score (as measured by R^2). An interpretation of this result could be that the extremely low VIP fee on this market is sufficiently small to overcome its general leadingness behaviour.

The Huobi spot market has one of the lowest VIP taker fees and its corresponding meta model has middling R^2 value. Based on these observations, we would predict a relatively large PnL_3 value although we would not anticipate it having the highest PnL_3 , as it turns out to be the case. As mentioned before, this might be partially explained by the fact that this market tends to have extremely low top of the book liquidity, implying that the strategy whose PnL we are computing is not immediately scalable to the deployment of substantial amounts of trading capital. The additional cost of slippage could become an important consideration at that point.

Surprisingly, the Binance BTC-margined perpetual, which has a similarly low taker fee as the USDT-margined one (see Figure 19), exhibits a lower PnL_3 when we compare the two (a relative difference of 18%), despite its considerably higher R^2 values. At the other end of the spectrum, the four markets whose PnL_3 is lowest are the precisely ones with the largest taker fees. This is a similar finding to what we noted for other models: the greater predictability is not sufficient to overcome the larger fee.

5.4 Comparisons of Models

We pursued three alternative methodologies for fitting linear models that anticipate price change on each market. The starting point were the baseline models that used all 70 features. Next, we introduced ℓ_1 regularization to the baseline model with nine different regularization parameters to obtain more parsimonious models with fewer non-zero coefficients. Finally, we devised an alternative methodology for reducing the number of features down to only five using an additional feature processing step. In this subsection, we set out to draw comparisons between the respective explanatory powers and PnL s of each of the trained models.

We begin by investigating how the R^2 values of our models compare to one another. Figure 21 shows the out-of-sample R^2 values of each model. The in-sample and 1000ms versions of the same illustration were deferred to the Appendix since all of the R^2 -based plots strongly resemble one another. The interested reader can find these in Figures 32, 31, and 33. In examining Figure 21, we first note that the maximal R^2 is attained by the baseline




Figure 21: 500ms out-of-sample R^2 values.

model in each case. LASSO models with small regularization parameter λ perform similarly. For example, the LASSO models with $\lambda = 0.016$ experience an average drop of only 2.8% relative to the baseline model and the $\lambda = 0.032$ models experience an average decrease of 5.9%. The average number of nonzero coefficients retained is 38 for $\lambda = 0.016$, and 29 for $\lambda = 0.032$. As noted earlier, this suggests that of the 70 original features used by the baseline model, around 30-40 contribute very little to the model’s performance. When λ is increased beyond 0.032, and as more and more coefficients are thus set to zero, we find progressively larger drops in explanatory power.

The meta model uses only five features and is by far the most parsimonious in terms of explanatory power per feature. To illustrate this point, we can compare the results of the meta model with those of the $\lambda = 0.256$ LASSO models which have a similar number of features (namely an average of four). Here we find a lower average R^2 of 21.2% for the LASSO models. The R^2 values of the meta models are actually quite similar to those of the $\lambda = 0.128$ models which retain an average of nine features. The average difference in R^2 values between the two is 3.5%. When we compare the meta models to the baseline models, we note a substantial reduction in explanatory power. The average reduction is 23.7%.

Let us now compare the PnL_3 values of each of the models shown in Figure 22. Remarkably, the meta models exhibit the largest average PnL_3 of -4114 bpts compared to -4666 bpts for the baseline models, which have the second largest average. These findings are unexpected in the sense that they do not agree with the earlier findings on explanatory power, which potentially hints at the power of the meta models and their simplicity for actual trading applications. With the exception of the meta models, the average PnL_3 is in fact monotonically decreasing in the number of coefficients. That is, as the regularization parameter is increased, the average PnL_3 decreases. See Figure 23 for a visual overview of these average values across markets.




Figure 22: PnL_3 values for each model.

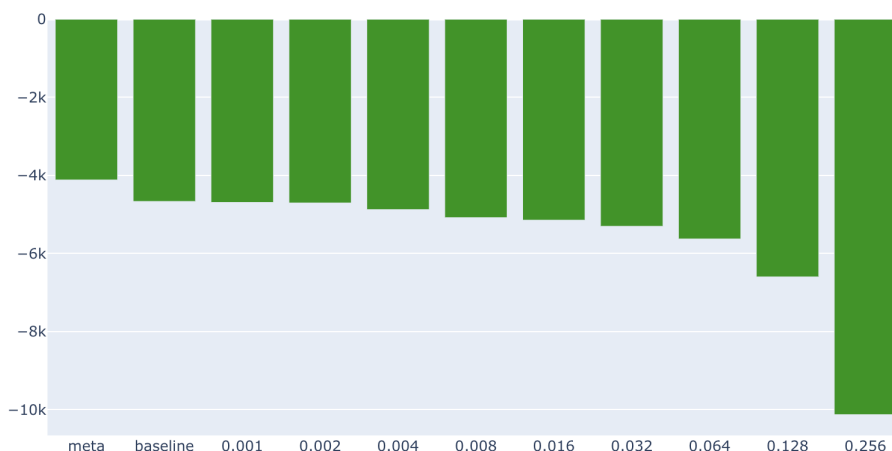


Figure 23: PnL_3 values averaged across markets for each type of model. The numeric values denote the LASSO regularization parameters.

It is interesting to contrast these findings with the analogous results for the the PnL_1 values of all models where execution cost is ignored, as shown in Figure 24. One is immediately struck by the fact that the meta models perform considerably worse than, say, the baseline model according to PnL_1 . As a matter of fact, the meta models

now have the *smallest* average value of 13402 whereas the baseline model has an average of 17695. For the LASSO models, the average PnL_1 decreases monotonically in the regularization parameter λ , starting at 17602 for $\lambda = 0.001$ and ending with 15128 for $\lambda = 0.256$.

How can we reconcile the fact that the meta models yielded the best results in terms of PnL_3 , but yet delivered the worst results if we consider PnL_1 as the metric? Since $\text{PnL}_3^i = \text{PnL}_1^i - \text{fee}_i \cdot n_i$ where fee_i is the taker fee on market i and n_i is the number of trades on market i , we must have a smaller number of trades for the meta models compared with the baseline models. That is, the strategy associated with the baseline model trades more often than that associated with the meta model, hence accumulating basis points in profit when fees are ignored, but accumulating losses when fees are accounted for.




Figure 24: PnL_1 values for each model.

The difference in the number of trades per model is illustrated in Figure 25. The strategy we associate to the




Figure 25: Number of trades for each model

meta models trade an average of 25.3% less often relative to its baseline model counterpart. The discrepancy is especially large on the Huobi spot market and quarterly futures contracts, where the number of trades is less than half that of the baseline models.

6 Market Making Experiments

The strategies we associated to each of the previously trained linear models have been *taker* strategies. That is, they rely on limit orders that would lead to immediate execution: buy orders use the top ask as their limit price, while sell orders use the top bid. We found positive PnL values even after taking execution cost into account, in a number of cases when the taker fee was particularly low. Some markets (such as the Bybit BTC perpetual) which have among the highest taker fees gave low PnL values, despite their strong tendency of being laggards.

However, as we noted before, a large taker fee usually also implies a large maker rebate. This is because the net fee (the difference between taker fee and maker rebate) is largely uniform across markets, as one would expect in a competitive exchange environment, since a market with comparatively very large net fees would naturally find it difficult to attract traders. On the other hand, a market which undercuts competing exchanges with exceptionally low net fees would likely struggle to be profitable. Let us consider more closely those markets which are laggards, have a high taker fee and consequently also a high maker rebate. The maker rebate on the Bybit BTC perpetual, for instance, is 2.5 bpts, which is the maximum among our set of 14 markets. It is natural to wonder whether we can leverage the large maker rebate on Bybit in conjunction with its lagging nature to instead devise a *maker* strategy which achieves positive PnL values.

In contrast with taker strategies, maker strategies employ *passive* limit orders (also called maker orders). That is, for a sell order, the limit price is strictly greater than the top bid price, while for a buy order, the limit price is strictly smaller than the top ask price. The order thus submitted rests “passively” in the orderbook, can be observed by others, and consumed by others with a taker order (an exception is the iceberg order type supported by some exchanges which allows one to submit passive orders that are not visible to other market participants). Execution of the maker order hinges on the arrival of a taker order which consumes it. It is therefore uncertain at what point in time a passive order will be executed, if at all. The compensation for this uncertainty is that a maker order submitted at time t necessarily achieves a better price (if it is executed) than a taker order submitted at time t would. In addition, maker orders incur less (in some cases negative) execution fees relative to taker orders. The trade-off between passive and aggressive orders can be concisely summarized as one of immediacy on one hand, versus “goodness” of execution price on the other hand.

What is required of a successful maker strategy? Since the execution of a maker order relies on the arrival of a taker order in the opposite direction, the maker order is naturally subject to a certain adverse price move. If we use the notion of microprice defined below, then it is true by definition that at the precise time of the fill, there is an adverse price move against the maker order. We define the microprice as

$$p_m := \frac{p \cdot a + p' \cdot a'}{a + a'}$$

where p and p' are the top bid and ask prices, respectively, while a and a' are the top bid and ask amounts, respectively. The goal of a maker strategy therefore must be to minimize the magnitude of the adverse price move following a fill event. In order to achieve this goal, the trader must constantly monitor market conditions for the potential future arrival of a significant adverse price move. When such a price move is imminent, the trader should cancel their passive order (and possibly resubmit at a different price level).

In this section, we set out to devise a maker strategy that builds on the previously trained meta models. These strategies will then be rigorously tested which, due to the difficulties in backtesting maker strategies, will involve a series of real world trading experiments on Bybit where the maker rebate is largest. Finally, we draw comparisons between the results of our maker strategy with those of a naive benchmark strategy, in order to showcase the efficacy of the trading alpha provided by the meta models.

6.1 Strategy Specification

This subsection is devoted to constructing a maker strategy that builds on the meta models. The reason we use the meta models as the foundation is two-fold. First, the meta models are far simpler than for instance the baseline models if we gauge model complexity by the number of covariates. Second, while the meta models had lower explanatory power than some of the other models, their efficiency in the sense of R^2 per feature was by far the largest, and they achieved the largest PnL values when fees were taken into account.

What is the simplest, most natural maker strategy one can define? Let us first consider what data is required in order to fully specify a maker strategy. For the sake of concreteness, let us consider in the following the case of a single passive sell order. It is straightforward to dualize this discussion to the case of a passive buy order. The specification of a strategy requires

1. a criterion for when to post a new order,
2. a limit price $p > \text{top bid}$ and an amount $a \in \mathbb{N}$ for the posted order, and
3. a criterion for when to cancel the ask order.

In order to reduce the problem to its most plain form, we shall utilize the following simplifications

- $p := \text{top bid} + \text{tick size}$,
- $a := 2000 \text{ USD}$, and
- Post a new order : \iff cancel criterion returns “*do not cancel*” and we currently do not have an ask order posted.

Here the *tick size* refers to the minimum price movement of the market in question. That is, we use a fixed price, namely the lowest possible ask level and a fixed (somewhat arbitrarily chosen) small order amount for the sake of testing. It is left for future work to investigate the scalability to larger order amounts. With these simplifications in mind, we need only to specify a condition for when to cancel an order posted at the top ask. From this, we will then obtain a full maker strategy. The cancel criterion for a passive order at the top ask level should be capable of anticipating when an adverse price move in excess of the maker rebate is imminent, and in that case return a “*cancel*” decision in time to successfully pull the order. As mentioned earlier, adverse selection is unavoidable - the question is just how much of it our fills are subject to it. If this is less than the maker rebate received per fill, we achieve profitability.

How can we leverage the meta models to define a natural cancel criterion? Let us fix a market $i \in \{1, \dots, 14\}$ on which to trade. In fact, as previously noted, we will conduct our experiments on the Bybit BTC perpetual. The meta model for market i is the fitted linear model

$$fret_t^{\delta,i} = \mu_i + \beta_{i,1}mIMB_t^{a,j} + \beta_{i,2}mIMB_t^{b,j} + \beta_{i,3}mTFI_t^i + \beta_{i,4}mPRET_t^i + \beta_{i,5}mDIV_t^i + \epsilon_{i,t}, \quad (35)$$

and we shall fix a future returns horizon $\delta = 1000 \text{ ms}$ for the remainder of this section. Let

$$F_i : \mathbb{R}^5 \rightarrow \mathbb{R}, \quad (36)$$

denote the market-specific function that maps an observation of meta features to the prediction of the fitted meta model according to Eqn. (35). Suppose we observe a new sample $x_t \in \mathbb{R}^5$ at a time t , where we have a sell order posted on the top ask level. We can then define the following cancel criterion parametrized by a constant T_0

$$\text{cancel top ask order} : \iff F_i(x_t) > T_0. \quad (37)$$

That is, when the meta model predicts a rise price of a certain magnitude (parametrized by T) based on the most recent feature observation, we shall cancel an order posted at the top ask if indeed we have one posted there at the time.

We are now left with the choice of the threshold T . If we select a value which is too large, we will end up cancelling very rarely and hence experiencing a lot of adverse selection. In the limit case $T = \infty$, we in fact never cancel (which will serve as our benchmark for comparison, as we will describe later). On the other hand, if we select a value T which is too close to zero, we will cancel extremely often and thus receive very little fills (and consequently collect very few maker rebates). As a matter of fact, if we were to choose $T = -\infty$ the cancel condition would always return “cancel” and we would receive no fills whatsoever, since we would never post a new order to begin with. Given the trade-offs highlighted by the limiting cases, how can we arrive at a sensible parameter value where we maximize the number of fills we receive, while simultaneously minimizing the amount of adverse selection our fills are subject to?

Our approach is to come up with an approximate answer to this question by learning from cancellation behaviour of other market participants. That is, we will attempt to characterize, by means of meta model predictions, times at which a significant amount of cancellations from the top ask level is imminent. More specifically, we search for meta feature observations in the union of training and test set (spanning February 22 – March 1) where, at times t the following conditions holds

- $a_t > M$,
- $p_{t+\delta} = p_t$ and $a_{t+\delta} < m$ for some $\delta \in (0, 500\text{ms})$.

Here a_t and p_t denote the top ask liquidity and price at time t , respectively. The constant $M > 0$ represents an at least moderately large amount of liquidity, while the scalar $m > 0$ represents a small amount of liquidity. Experimentation has revealed definitions of m and M as the first and second quartile (respectively) to be sensible ones. The quartiles are computed over all observations of the top ask liquidity in the training and test sets combined. Let us denote the subset of meta feature samples defined by the above two condition by \mathcal{F} .

The intuition behind our approach is to scan our data set for examples where we pass from a moderate-high liquidity regime (on the top ask level) to a low liquidity regime. These are cases representing the cutoff point at which the market (i.e. most market participants) agree that it is no longer profitable to post an order at the top ask level. We want to understand such cases by means of our meta models. To this end, consider Figure 26a where we display a histogram of values from the set

$$\mathcal{P} := \{F_i(x) : x \in \mathcal{F}\}. \quad (38)$$

That is, we show the distribution of predictions made by the meta model on samples from the set \mathcal{F} defined by the two conditions above. More precisely, we uniformly select a random subset of size 20,000 from these predictions in order to avoid many successive samples and to keep the histogram more legible. The average and median predictions are 2.04 and 1.83 respectively, while the standard deviation is 5.14. We compare this with the histogram shown in Figure 26b, where we display a uniformly random size 20,000 subset of $\{F_i(x) : x \text{ is any training or test sample}\}$. Here we find an average prediction value of 0.47, a median of 0.44 and a standard deviation of 7.88.

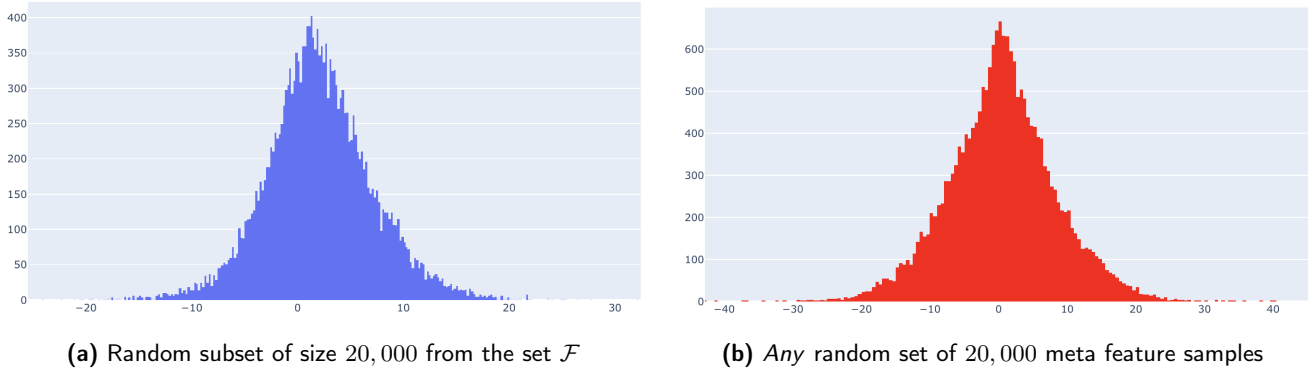


Figure 26: Distribution of meta model predictions.

From the preceding analysis, we might extract $T := 1.83$, the median prediction on the set \mathcal{F} , as a reasonable candidate for the definition of our cancel criterion and hence the full specification of the maker strategy. In the following, we will examine this choice from another perspective and draw comparisons with other possible values for T .

Next, we would like to measure the expected adverse selected when the prediction exceeds T . To this end, let us consider Figure 27 where we computed the average price (top ask) increase (in basis points) after δ seconds, for $\delta \leq 500$ s, conditioned on a meta model prediction greater than T_q for $q \in \{0.3, 0.4, 0.5, 0.6, 0.7\}$. The quantity T_q is defined as the q -th quantile of the set \mathcal{P} from Equation 38. Note that $T_{0.5} = 1.83$ is our candidate threshold. We

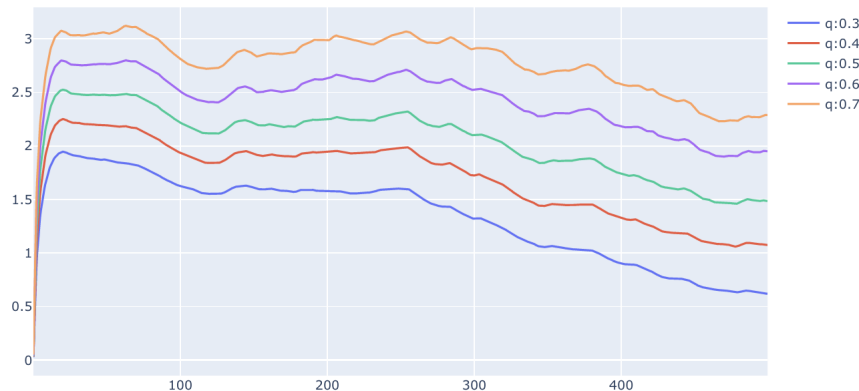


Figure 27: Average top ask price returns after δ second conditioned on meta model prediction greater than T_q

observe in Figure 27 that our candidate threshold (corresponding to $q = 0.5$) exactly attains the 2.5 bpts mark after

a few seconds, which is precisely the breakeven point of tolerable adverse selection, since the maker rebate is 2.5 bpts. This provides us with further evidence of the suitability of the $T = 1.83$ threshold.

Impossibility of backtesting. At this point, we have a candidate maker strategy (or a set of candidate strategies) that we would like to test, in order to determine whether it is capable of achieving positive PnL. But how exactly can we test performance of a maker strategy? It turns out that there are a number of difficulties associated with this, making it essentially impossible to do so with any reasonable degree of precision. Let us contrast this with the case of backtesting a taker strategy on historical data.

The situation is straightforward in the taker case due to the immediacy of execution for taker orders: when a candidate strategy provides a buy signal, we assume execution at the top ask price p_a or, in more elaborate versions, at a price $p > p_a$ that takes (i.e. *sweeps*) liquidity from a number of ask levels into account. Similarly, when a sell signal occurs, we assume execution at a price less or equal to the top bid price. As we scan over a historical test set in this manner (computing buy and sell signals at each point), we obtain a sequence of hypothetical buy and sell orders, which can be used as the foundation for subsequent PnL calculations. While this analysis neglects some market micro-structural effects such as market impact and latency concerns, it nevertheless provides a good approximation of the actual PnL of the strategy. Moreover, there are also fairly natural ways of mitigating some of the aforementioned concerns. For instance, latency concerns largely relate to uncertainty of execution, and can be addressed by assuming that hypothetical orders fail with some nonzero probability, which could, for example, be calibrated from historical execution data. Market impact can be accounted for by using an impact model from the large body of literature on this generally well understood topic.

How does this compare to the case when we would like to backtest a maker strategy? The underlying difficulty is the fact that the execution of a posted maker order is uncertain. It is not known when it will occur, or if it will occur at all. Execution hinges on the arrival of a taker order in the opposite direction of the posted maker order. This means that if a candidate strategy which is being tested posts a hypothetical maker order at some orderbook level, we must henceforth keep track of the queue position of our hypothetical maker order. While this hypothetical order remains posted in the orderbook (i.e. it has not been cancelled in the meantime), we must monitor all opposite-side taker orders and evaluate, based on the size of the taker order and our inferred queue position, whether it would have taken (i.e. filled) our maker order. There are a number of factors that present challenges to doing this accurately, whose severity ranges from mild to critical. On the mild side of the spectrum, we are facing a similar pair of problems as in the taker case:

- The submission of a passive order has an impact on the market. One would expect that the presence of a large maker order discourages aggressive orders that take from the orderbook level at which the maker order is posted. Hence, when we keep track of historical taker orders and attempt to infer whether they filled a hypothetical maker order of ours, we must account for this by assuming a nonzero probability that the taker order in fact does not arrive or that its size would have been less than the one observed historically. This sort of market impact of passive orders is far less well-understood and understudied relative to the impact of taker orders. Nevertheless, we rate this concern as only mild in severity.
- There are latency issues similar to those that one must contend with when backtesting a taker strategy. For instance, at time t our maker strategy outputs a decision to cancel a posted order and at time $t' \approx t$ a large taker order arrives that would certainly have consumed our passive order. It can be difficult to know with certainty whether the cancellation would have been processed first by the exchange's matching engine, or whether the taker order would in fact have consumed our passive order. To address this concern, we must associate a probability of failure to our cancellation that we could, for instance, calibrate from historical cancellation data and which depends on the time difference $t' - t$. We rate this concern as mild.

A much greater challenge, which is essentially detrimental to the ability to backtest maker strategies, is related to the fact that we must keep track of our *queue position* in the orderbook, in order to gauge whether any particular opposite-side taker order consumed our passive order. This turns out to be virtually impossible with the data we have at our disposal in this study, and more broadly, in the crypto space, given the lack of Lever 3 (order-by-order) data in this ecosystem. This is best understood by an example. Consider Figure 28 and suppose that the orderbooks displayed there are subsequent snapshots received from the exchange. Let us denote the arrival times of the respective orderbook snapshots by t_1, t_2 and t_3 . Consider an example where our maker strategy behaves in the following manner.

1. At time t_1 , the maker strategy posts a new hypothetical ask order at the top ask. Hence there is $386194 \approx 368k$ USD in liquidity in front of this order, in the queue on this level.

- At time t_2 , the maker strategy keeps its order on the top ask level. In light of the large amount of new submissions to the top ask, there is now presumably still $\sim 368k$ USD of liquidity in front of our hypothetical maker order and additionally $\sim 2122k$ USD behind our order in the top ask queue.
- Suppose at time t_3 our maker strategy decides to remain posted on the top ask level. There was a significant amount of cancellations from the top ask, with roughly 2165k USD removed. We do not know whether these cancellations correspond to the orders that were originally *in front* of our hypothetical order or *behind* it in the top ask queue. Hence, at this point our queue position is completely unknown. Our hypothetical maker order could now well be the first liquidity on the top ask level or the last.
- Suppose at a time $t_3 + \epsilon$ the exchange informs us about the arrival of a taker buy order of size 100k USD, while the most recent orderbook we received remains the one from time t_3 . We now have no way of assessing whether the taker order consumed our posted ask order or not. This depends entirely on our queue position which, as noted above, is not known. Modeling the queue position is a very interesting problem in itself, and a high impact topic of independent interest, which merits its own study. While the topic is relatively unexplored in the literature, the interested reader can find more information on this subject in, for instance, [26].

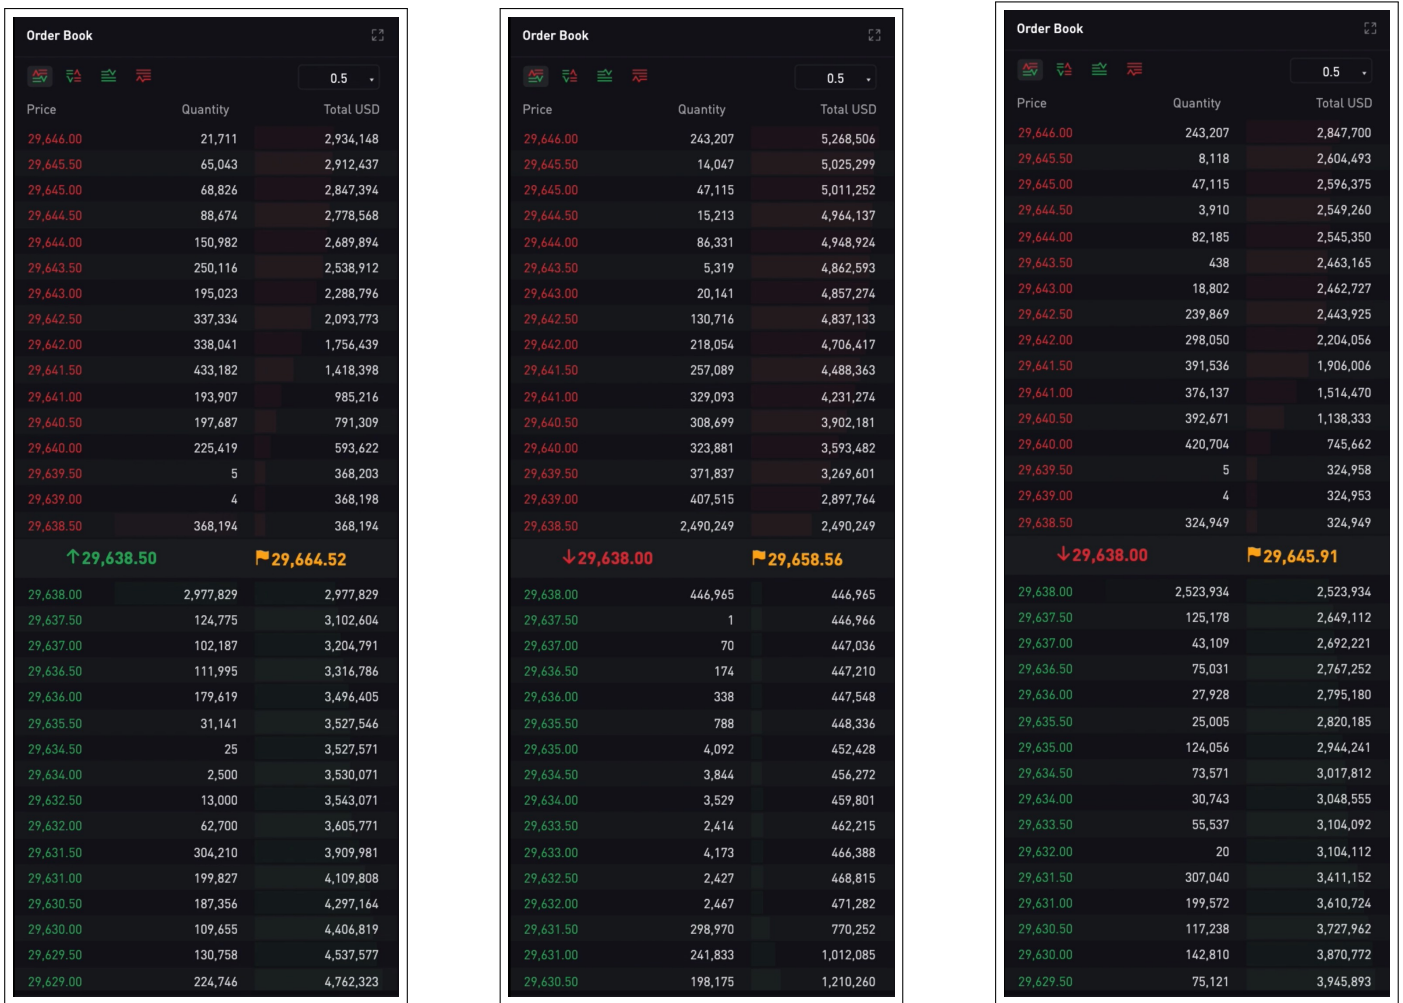


Figure 28: Consecutive orderbook snapshots, highlighting a scenario when it is virtually impossible to backtest a maker strategy, mainly to issues arising from the lack of knowledge of the queue position.

This example sheds light on the difficulty in keeping track of queue position of a hypothetical maker order. Similar examples can be easily constructed. They are the norm, not the exception. Generally speaking, the longer a hypothetical order remains posted, the more difficult it becomes to know its queue position exactly. The core issue which makes it impossible to precisely track queue position is related to the nature of the data we receive from the exchange, namely orderbook snapshot data as opposed to, for instance, order-by-order data. If we instead

received order-by-order data that included unique identifiers for each order, we would know exactly which orders were cancelled and what their queue positions were. From this data, we could then infer our exact queue position and hence whether or not a hypothetical passive order of ours would have filled. Alas, this data is not available in crypto markets, so we must accept the impossibility of precise backtesting of candidate maker strategies. We will circumvent this issue by conducting a series of real world trading experiments, to test performance and compute PnL of maker strategies.

6.2 Methodology and Experimental Setup

While backtesting of maker strategies with any reasonable degree of precision is not possible in crypto markets according to our remarks in the previous subsection, it is relatively simple for individuals to trade in an automated fashion on crypto markets due to the ease of creating exchange accounts, and accessing market data and trade infrastructure through public API endpoints. This is in contrast to traditional financial markets with a high barrier to entry, where one must usually interface with brokers, undergo a number of elaborate verification procedures, set up a prohibitively elaborate trading infrastructure to interface with relevant markets, and so on. We leverage the comparatively easy access to crypto markets in order to test, in the real world, the performance of our previously defined candidate maker strategy. In this subsection, we describe the framework and scope of our trading experiments, along with details of how they are conducted.

Some practical matters. Let us begin by describing a number of practical circumstances involving the real-time data collection and feature calculation process. The data feeds we use here are the same ones used to gather historical data, and thus the descriptions given in Section 2 apply equally here. That is, we use websocket API feeds for each of the 14 markets whose data we require. This is generally the quickest way of obtaining real-time data in crypto markets, since the exchange automatically pushes orderbook or trades updates to the subscriber without them having to explicitly prompt the exchange for the most recent market data (which is the case for the alternative type of data feed, the so-called REST API requests). Orderbook and trades data are received via different websocket feeds.

We are mining this real-time data in a server cloud-colocated to the server of the exchange on which we are trading, which is Bybit. To explain the term “cloud-colocated”, we first remark that Bybit, as most other crypto exchanges, are hosted in a cloud computing facility. In the case of Bybit this is AWS, and specifically the Singapore region. Colocation can then be easily achieved by renting server space in the same AWS region where the exchange is located. In doing so, we achieve a sub-1ms latency with Bybit.

The market data we receive is processed in real time to first compute values of the features (or transformed features, when applicable), then the meta features, and finally the output of the meta model and the corresponding output decision. This entire computation pipeline takes between zero and two milliseconds. The main bottlenecks are the computations of the average prices $p_{a,t}^i(N_i)$ and $p_{b,t}^i(N_i)$ that appear in the definition of the orderbook imbalance (4). This is due to the fact that their computation can involve iterations over many orderbook levels. We trigger a new feature calculation (on the latest market data) and decision output roughly every 5-10ms, depending on market volatility.

Experimental setup. We now describe the setup of our experiment, whose goal is to evaluate the PnL of the strategy devised in the previous subsection. To this end, let us first recall the specification of the strategy. For a sell order, it is defined by the following set of rules:

- (a) The sell post price at time t is $p_{b,t} + \Delta$, where $p_{b,t}$ is the top bid price on the Bybit BTC perpetual at time t and Δ represents the tick size.
- (b) The order amount is fixed at 2000 USD.
- (c) Cancel a top ask order : $\iff F_i(x_t) > T := 1.83$ where $F_i(x_t)$ is the prediction of Bybit’s meta model on the latest meta feature observation x_t
- (d) Post a new top ask order : $\iff F_i(x_t) \leq T$.

However, due to rate limit constraints imposed by the exchange, we have found it necessary in practice to further restrict condition (d). Specifically, the exchange allows only for a limited number of API requests (order placements or cancellations) per minute. When the limit is exceeded, any subsequent requests fail for some period of time. Luckily, the exchange provides information about the number of remaining API requests. We incorporate this information by additionally requiring that there be at least two remaining API requests. That is, if we let rrl_t denote the remaining rate limit at time t , we restrict (d) by additionally imposing the condition $\text{rrl}_t > 1$ (we need at least two requests so we can cancel after posting an order).

Even with this additional restriction, however, we ran into some issues. Whenever $F_i(x_t) \approx T$, there are often many up- and down-crossings of T in a short time window $[t, t + \delta]$. This results in rapid depletion of the number of permitted API requests. The consequence is that we must subsequently pause for roughly one minute until a new order can be submitted. The effect is a significant portion of inactivity periods. We have addressed this problem by further restricting condition (d). Specifically, we additionally require certain trade flow imbalance in favor of the potential sell order. More explicitly, this condition takes the form

$$mTFI_t^i < T',$$

for some threshold T' . This type of condition was chosen not for any strong theoretical reason, but because it works well in practice. The threshold T' was selected on the basis of intuition, but also trial and error. We leave further calibration and a more rigorous approach to future work. In the meantime, the net result of our restriction measures is that we have reduced the number of new submissions and mitigated extended inactivity periods due to quickly depleted rate limits. The updated condition (d) now takes the form

$$\text{Post a new top ask order} : \iff F_i(x_t) \leq T \quad \wedge \quad mTFI_t^i < T' \quad \wedge \quad \text{rrl}_t > 1.$$

Certainly, we want to not only post sell orders but also buy orders. We have decided to avoid subsequent same-side fills to simplify post-analysis of the results. To elaborate on why this simplifies matters, consider a case when we do not impose this restriction and we receive many subsequent sell fills with no buy fills in the intervening period. The result will be a significant short inventory. If there then happens to be a drastic price jump up, the net results will amount to an extremely high PnL loss. On the other hand, a significant price decrease will result in potentially large profits that are not indicative of long term performance. This is assuming that we use the comparison between average buy and average sell price over the sample of all fills as an evaluation metric, which indeed we do, as we shall discuss further below when we examine results.

This issue is not difficult to address, however. We do so by additionally tracking a variable P_t which tells us the net position (or inventory) at time t . If the inventory is negative (that is, we are short), we rule out submissions of new sell orders. By the same token, if we have long inventory, we do not permit the posting of a new buy order. With this new restriction, the final form of condition (d) from above is as follows

$$\text{Post a new top ask order} : \iff F_i(x_t) \leq T \quad \wedge \quad mTFI_t^i < T' \quad \wedge \quad \text{rrl}_t > 1 \quad \wedge \quad P_t > 0.$$

We define a maker *buy* strategy completely analogously by the following set of rules

- (a) The buy post price at time t is $p_{a,t} - \Delta$ where $p_{a,t}$ is the top ask price on the Bybit BTC perpetual at time t and Δ represents the tick size.
- (b) The order amount is fixed at 2000 USD.
- (c) Cancel a top bid order : $\iff F_i(x_t) < -T$
- (d) Post a new top bid order : $\iff F_i(x_t) \leq T \quad \wedge \quad mTFI_t^i < T' \quad \wedge \quad \text{rrl}_t > 1 \quad \wedge \quad P_t > 0.$

Now we have fully specified our maker buy and sell strategies in their final form. For each of these, we launch an independent "bot" that operates according to the above specified rules. That is, we launch a *sell bot* with access to real time data (including the remaining rate limit rrl_t and net position P_t) which acts according to the set of rules for the sell strategy. Likewise, we launch a *buy bot* which uses the same data feeds, trades on the same exchange, and acts according to the rules of the buy strategy. Note that the data feeds rrl_t and P_t are shared. This means, for instance, that if the buy bot is barred from posting a new buy order due to $P_t > 0$ (see condition (d) above) and the sell bot receives a fill at time $t + \delta$ resulting in $P_{t+\delta} < 0$, then the buy bot will receive the updated value $P_{t+\delta}$ and become active again. Both bots operate in perpetuity according to these guidelines. We initialize the position by $P_0 := 1$.

In order to be able to better interpret results and draw comparisons, we launch another pair of buy and sell bots operating the same way as we outlined above, with the exception that they use $T = \infty$ and fixed order amount $a := 100$ (we reduced the amount since this strategy is expected to be extremely unprofitable). The modification of the threshold T has large ramifications: the resulting pair of bots never cancel after posting an order, hence have no capability of avoiding adverse selection. We shall call this the *naive benchmark* or just *benchmark* strategy. It will also be useful to have a name for the other pair of bots (the non-benchmark ones), which we will refer to as the *metaMM* bots or strategy.

6.3 Results and Comparison with Benchmark

In the preceding two subsections we defined a maker strategy associated to the meta model on the Bybit BTC perpetual, and we outlined our methodology for conducting a real-world market making experiment on the aforementioned market. This subsection is devoted to the examination of the results.

The total filled amount for the metaMM strategy is 1.537 million USD. The bots were active from 2021-07-08 13:44:02 until 2021-07-13 23:54:02+00:00. This implies an hourly fill volume of 11.810k USD, which corresponds to approximately six “units” since the order size is fixed at 2,000 USD. Over the entire sample, the difference between the average sell and buy price is -4.4749 bpts. More specifically, let us define the average sell price by

$$p_s := \frac{1}{A} \sum_{(p,a) \in \mathcal{G}^{(s)}} p \cdot a,$$

where $\mathcal{G}^{(s)}$ is the set of all sell fills over the sample period, and $A := \sum_{(p,a) \in \mathcal{G}^{(s)}} a$. In the same manner, we define the average buy price

$$p_b := \frac{1}{A} \sum_{(p,a) \in \mathcal{G}^{(b)}} p \cdot a.$$

Then it holds that

$$\text{bpts}_{\text{metaMM}} := \left(\frac{p_s}{p_b} - 1 \right) \cdot 10000 = -4.4749,$$

over the set $\mathcal{G}^{(b)}$ of buy fills. In other words, a roundtrip trade loses on average -4.4749 bpts. However, recall that each leg (buy or sell) of the trade receives a 2.5 bpts maker rebate so that a roundtrip trade nets a total of 5 bpts in rebates. This implies an average profit per roundtrip trade of $5 - 4.4749 = 0.5251$ bpts. The median interarrival time between opposite-side fills is 328 seconds.

Note that the trading period of the metaMM strategies is more than four months after the training and test periods that were used to calibrate the models underlying metaMM. We expect that better results could be achieved if we repeat the experiment using more recent data for the model calibration. Indeed, we performed similar real-world trading experiments at the end of February and March which exhibited slightly better performance than the one shown in this paper. These experiments were conducted over a longer period of time and used larger order sizes (up to 20k USD). In total, the notional USD turnover in this initial live trading experiment was more than 100 million USD. The methodological approach was slightly different, however, and does not match well with the narrative of the present work, thus we decided to redesign the experiment to a more fitting approach.

For the benchmark strategy we employed the 20-fold smaller order amount $a = 100$. The total filled amount was 153.26k USD representing 1532 units. The activity period of the benchmark bots was from 2021-07-17 12:11:27 until 2021-07-18 06:59:26, hence considerably shorter than for the metaMM strategy. This amounts to 8.152k USD in fill volume per hour. The benchmark strategy achieved the average value (defined similarly to the metaMM case)

$$\text{bpts}_{\text{benchmark}} = -6.8184,$$

on a roundtrip trade which, in line with our expectation, did not yield a profit despite the 5 bpts maker rebate per roundtrip trade. However, the median interarrival time between opposite-side fills is far lower for the benchmark bots, namely only 27 seconds, which is less than 10% of the corresponding metaMM value. This is sensible, since the metaMM strategy cancels much more often after it posts, and therefore accumulates fills much more slowly.

In Figures 29a and 29b we show histograms, one for the metaMM strategy and one for the benchmark bots, of the differences in basis points between pairs of sell and subsequent buy orders. It is interesting to note here that the benchmark strategy appears to have a much smaller proportion of roundtrip trades with positive basis point difference.

In Figure 30a, we show PnL over time for the metaMM bots, as well as an overlaid Bitcoin price chart, which we included to give an indication of price volatility over the sample period. We suspect that the strategy’s expected performance is worse during periods of high volatility - a potentially interesting subject of investigation for future work. Furthermore, we hypothesize that the variation in the strategy’s PnL is far greater during highly volatile times. This is made plausible by the observation that high volatility implies more frequent price jumps and hence larger PnL jumps due to inventory held during a large price rise or fall (even if the inventory is small). The plot 30a includes both realized and unrealized PnL. To elaborate on what exactly this means, let us introduce some notation. If at time t the set of (buy and sell) fills is \mathcal{G}_t , the inventory is $x \neq 0$ and the last fill price is p , we define the set $\mathcal{G}'_t := \mathcal{G}_t \cup \{(-x, p)\}$. The realized and unrealized PnL computed is then the PnL computed over the set \mathcal{G}'_t , while




Figure 29: Distribution of the basis point difference between buy and subsequent call fill




Figure 30: Cumulative PnL in USD, across time.

the PnL computed over the set \mathcal{G}_t is just the realized PnL. If we denote the output of this calculation by PnL_t , then the pair (t, PnL_t) represents a point on the “rpnl + upnl” graph shown in Figure 30a. That is, the quantity PnL_t takes into account the “PnL potential” stemming from the fact that the inventory can be nonzero. Note that by design of our experiment, we have $|x| \leq 2000$ so the “unrealized” component is generally quite small (it never exceeds 30 USD).

For comparison, we show a similar plot for the benchmark strategy in Figure 30b. It is striking how steadily the PnL decreases in this case. The contrast between the two PnL graphs provides clear visual evidence of the additional “alpha” that the meta model provides for this maker strategy.

For our last performance evaluation metric, we sought a measure for the adverse selection that our fills were subject to. Intuitively, this means that we want to compare the fill price with another price some time into the future. In order to rigorously define the notion of “adverse selection”, one needs to define a time horizon and a notion of price. For our comparisons, we used the last traded price on the Bybit BTC perpetual as the reference price and employed a number of different time horizons, namely $\delta \in \{0.5s, 5s, 10s, 30s, 60s, 150s, 300s, 600s, 1200s, 2400s, 4800s\}$. For each of these time windows, we measure how much the price moved against our fill price on average, what the median adverse price move is, the 25th and 75th percentile, and the minimum and maximum. We additionally show the standard deviation in adverse price move. The results are shown in Table 15 for the metaMM bots, and in Table 16 for the naive benchmark.

When comparing the adverse selection that the benchmark fills are subject to with that of the metaMM fills, we

Table 15: Adverse selection over different time horizons for metaMM strategy

	0.5s	5s	10s	30s	60s	150s	300s	600s	1200s	2400s	4800s
avg	-0.51	-0.97	-1.24	-1.69	-1.99	-2.27	-2.44	-2.19	-2.88	-2.91	-1.73
std	1.61	2.74	3.57	5.41	7.54	12.52	16.21	23.12	30.04	40.94	57.06
min	-12.66	-24.18	-31.75	-23.75	-37.99	-53.18	-69.18	-139.03	-129.19	-215.28	-217.02
25%	-0.15	-1.07	-1.94	-3.74	-5.71	-9.48	-12.45	-15.26	-20.32	-25.09	-34.52
50%	0.00	0.00	-0.15	-0.74	-1.22	-2.32	-2.56	-3.14	-3.69	-1.64	-1.80
75%	0.00	0.00	0.00	0.15	0.30	4.73	6.82	11.52	14.24	18.27	33.22
max	3.87	13.72	17.41	28.14	30.46	42.77	61.28	110.43	108.80	192.34	199.02

Table 16: Adverse selection over different time horizons for naive benchmark strategy.

	0.5s	5s	10s	30s	60s	150s	300s	600s	1200s	2400s	4800s
avg	-0.49	-1.65	-2.28	-2.78	-2.94	-3.32	-3.26	-3.00	-3.16	-3.24	-3.06
std	1.25	2.96	4.12	7.10	10.40	16.33	23.63	30.26	38.49	51.81	89.72
min	-12.05	-19.97	-26.90	-38.16	-61.74	-91.39	-122.95	-164.41	-164.73	-184.83	-368.74
25%	-0.34	-2.78	-4.15	-6.31	-8.58	-13.23	-17.64	-22.91	-28.00	-36.19	-45.80
50%	-0.00	-0.43	-1.19	-2.16	-2.41	-3.47	-3.08	-2.99	-3.33	-3.31	-3.64
75%	0.08	0.00	0.00	0.09	2.04	6.51	11.70	16.21	21.17	30.10	40.48
max	3.53	16.77	31.85	54.40	50.37	73.54	119.77	132.37	159.60	182.54	365.70

note that after a “settling period” or around 10s, the difference between the two stabilizes at around one basis point. That is, the benchmark fills experience an average adverse price move which is one basis point worse than that experienced by the metaMM fills. While a mere basis point may not sound particularly substantial, we see in the PnL graphs what ramifications this has. The 500ms time horizon is anomalous in the sense that the benchmark sees less adverse selection over this time period. The standard deviation in adverse selection is smaller for the metaMM strategy, with exception of the 500ms time window. The difference becomes more pronounced the larger the time horizon.

7 Conclusion

This study used a granular data set comprised of market data from the most liquid Bitcoin markets, in order to closely examine the Bitcoin price formation process on a micro scale. To this end, we defined a set of features that encapsulate microstructural information on each market, which we subsequently leveraged to generate a leader-lagger network illuminating the cross-impact relationships inherent in our set of markets. Perpetual swaps and quarterly futures on Binance and Huobi were found to be particularly strong leading markets, while the Bybit and FTX perpetuals were found to be strong laggards.

The generated features were subsequently utilized further for the development of linear price prediction models. We compared three methodologies of training such linear models. The first was based on a simple OLS regression using the full set of 70 features, the second employed a LASSO regression instead with a number of different regularization parameters, and the third inserted an additional dimensionality reduction step before applying an OLS regression. The largest explanatory powers over future returns were achieved with the first approach, though LASSO models with only around half of the features surviving (that is, having nonzero coefficient) yielded similar R^2 values. On the most lagging markets, we were able to achieve R^2 values as high as 36% for the task of predicting 500ms future returns. Even the most leading market still yielded an R^2 value larger than 10%.

In a next step, we defined a natural trading strategy associated to each model, in order to assess how well predictability (as measured by R^2) translates to trading profit. These strategies employ taker orders, in contrast with another strategy we define later based on maker orders. One would expect large R^2 to correspond to high PnL values, but this turned out not to be the case. We noted that a market’s propensity to either lead or lag is intimately tied to its fee structure. A high fee regime consisting of a large taker fee and a large maker rebate predispose a market to lagging behaviour. Conversely, a low fee regime, where both taker fee and maker rebate are small, implies a predisposition to leading behaviour. Only when a market’s actual laggingness exceeded the one implied by its fee structure, did our strategy yield large (and positive) PnL values therein. A notable example was the FTX perpetual, which was among the most predictable markets despite having an extremely low taker fee.

Finally, we constructed a simple maker strategy associated to our linear models in order to address the question of whether the large maker rebates can be converted into profit on exchanges where we found highly negative PnLs with the previously tested taker strategy. Due to the difficulties involved in testing maker strategies on historical data, we devised a live trading experiment on Bybit, where we implemented and deployed real capital to the strategy in the real world. It was found to yield approximately one basis point of “alpha” compared to a naive benchmark strategy, which was sufficient to produce positive PnL over a sample period of roughly five days.

Future work. Our framework opens up a variety of avenues for future research. For one, a natural extension would be to incorporate alternative cryptocurrencies, such as Ethereum, into our analysis. Nowadays, daily Ethereum trade volumes routinely represent around half of the Bitcoin ones (and occasionally even exceeding them). Given the high correlation between these two assets, we would certainly expect significant cross-impact to take place, even on small time scales and potentially as a useful predictive signal. Similarly, one could also explore the interplay

between equities and Bitcoin markets, particularly around times of market open and close where equities trade volumes tend to be largest.

A second avenue for potential future work is to consider extensions of the taker and maker strategies we presented in this paper. One could, for instance, allow for greater inventory risk in the maker strategy, and possibly vary the post (i.e. submission) price depending on the size of the inventory, thus trading off the strength of the alpha with the inventory risk. For the taker strategy, one could explore how profitability is affected by different execution styles, which go beyond the simple top-of-the-book execution that we assume.

Further related to the market making strategy, it would be interesting to investigate its scalability to larger order amounts, along with the interplay of its profitability with volatility conditions. Building models for estimating queue position and incorporating this into the trading behaviour would likely lead to more realistic backtesting, and potential additional profitability when deployed in a live market experiment.

Lastly, at a more foundational level, our feature set can certainly be refined and extended in a number of ways, or simply processed differently. For instance, instead of applying our nonlinear feature transform to the base features, one might experiment with utilizing nonlinear methods, such as tree ensembles or neural networks, applied directly on the set of base features.

References

- [1] Bitmex perpetual guide. Available at <https://www.bitmex.com/app/perpetualContractsGuide>.
- [2] Bitwise volume report. Available at <https://static.bitwiseinvestments.com/Research/Bitwise-Asset-Management-Analysis-of-Real-Bitcoin-Trade-Volume.pdf>.
- [3] Bybit websocket api documentation. Available at <https://bybit-exchange.github.io/docs/linear/#t-websocket>.
- [4] Coin metrics market manipulation report. Available at <https://coinmetrics.substack.com/p/coin-metrics-state-of-the-network>.
- [5] Cryptocompare spot vs derivatives volumes. Available at https://www.cryptocompare.com/media/37746011/cryptocompare_exchange_review_2021_01.pdf.
- [6] Deribit index documentation. Available at <https://test.deribit.com/pages/docs/general>.
- [7] Kimchi premium investopedia. Available at <https://www.investopedia.com/terms/k/kimchi-premium.asp>.
- [8] Pandas resample function. Available at <https://pandas.pydata.org/pandas-docs/stable/reference/api/pandas.core.resample.Resampler.last.html>.
- [9] Spot vs leveraged trading. Available at <https://www.binance.com/en/support/faq/360033162052>.
- [10] Volume monitor. Available at <https://ftx.com/volume-monitor>.
- [11] Volume report. Available at <https://ftx.com/volume-report-paper.pdf>.
- [12] Websocket wikipedia article. Available at <https://en.wikipedia.org/wiki/WebSocket>.
- [13] Saketh Aleti and Bruce Mizrach. Bitcoin spot and futures market microstructure. *Journal of Futures Markets*, 41(2):194–225, Oct 2020.
- [14] Carol Alexander, Jaehyuk Choi, Heungju Park, and Sungbin Sohn. Bitmex bitcoin derivatives: Price discovery, informational efficiency, and hedging effectiveness. *Journal of Futures Markets*, 40(1):23–43, Sep 2019.
- [15] Carol Alexander and Daniel F. Heck. Price discovery in bitcoin: The impact of unregulated markets. *Journal of Financial Stability*, 50:100776, 2020.
- [16] Aurélien Alfonsi, Antje Fruth, and Alexander Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157, Jun 2009.
- [17] Markus Bibinger, Nikolaus Hautsch, Peter Malec, and Markus Reiss. Estimating the Spot Covariation of Asset Prices - Statistical Theory and Empirical Evidence. *Journal of Business & Economic Statistics*, 37(3):419–435, 2019.

- [18] Jean-Philippe Bouchaud, J. Farmer, and F. Lillo. How markets slowly digest changes in supply and demand. *Capital Markets: Market Microstructure*, 2008.
- [19] Jean-Philippe Bouchaud, Marc Mézard, and Marc Potters. Statistical properties of stock order books: empirical results and models. *Quantitative Finance*, 2(4):251–256, Aug 2002.
- [20] Rama Cont, Arseniy Kukanov, and Sasha Stoikov. The price impact of order book events. *Journal of financial econometrics*, 12(1):47–88, 2014.
- [21] John M. Griffin and Amin Shams. Is Bitcoin Really Untethered? *The Journal of Finance*, 75(4):1913–1964, Jun 2020.
- [22] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA, 2001.
- [23] R. Huang and T. Polak. LOBSTER: Limit order book reconstruction system. *Tech. rep.*, 2011.
- [24] Jui-Cheng Hung, Hung-Chun Liu, and J. Jimmy Yang. Trading activity and price discovery in bitcoin futures markets. *Journal of Empirical Finance*, 62:107–120, 2021.
- [25] Katya Malinova and Andreas Park. Subsidizing liquidity: The impact of make/take fees on market quality. *The Journal of Finance*, 70(2):509–536, Mar 2015.
- [26] C. Moallemi and Kai Yuan. A model for queue position valuation in a limit order book. *Columbia Business School Research Paper No. 17-70*, 2017. Available at SSRN https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2996221.
- [27] Vasiliki Plerou, Parameswaran Gopikrishnan, Xavier Gabaix, and H. Eugene Stanley. Quantifying stock-price response to demand fluctuations. *Physical Review E*, 66(2), Aug 2002.
- [28] Marc Potters and Jean-Philippe Bouchaud. More statistical properties of order books and price impact. *Physica A: Statistical Mechanics and its Applications*, 324(1):133–140, 2003. Proceedings of the International Econophysics Conference.
- [29] Kyle Soska, Jin-Dong Dong, Alex Khodaverdian, Ariel Zetlin-Jones, Bryan Routledge, and Nicolas Christin. Towards understanding cryptocurrency derivatives: a case study of bitmex. *Proceedings of the Web Conference 2021*, Apr 2021.
- [30] Daniel Tkacik. Crypto derivatives markets are booming. Available at <https://engineering.cmu.edu/news-events/news/2021/04/19-bitmex.html>.
- [31] Ke Xu, Martin D. Gould, and Sam D. Howison. Multi-level order-flow imbalance in a limit order book. *Market Microstructure and Liquidity*, 04(03n04):1950011, 2018.

Appendix

In this appendix, we will occasionally abbreviate the names of markets according to the acronyms in Table 17, when it is necessary to save space.

Table 17: Market acronyms

	Acronym
hbdm_BTC-USD	m1
hbdm_BTC_CQ	m2
bitmex_BTC/USD	m3
deribit_BTC-PERPETUAL	m4
bybit_BTC/USD	m5
binancefut_BTC/USDT	m6
okex_BTC-USD-SWAP	m7
okex_BTC-USD-210326	m8
ftx_BTC-PERP	m9
huobipro_BTC/USDT	m10
okex_BTC-USDT-SWAP	m11
tbybit_BTC/USDT	m12
thbdm_BTC-USDT	m13
binancecmfut_BTC/USD	m14

Table 18: Summary statistics for all models M_{ij} from Eqn. (22).

	IMB_a t-stat	IMB_b t-stat	TFI t-stat	PRET t-stat	DIV t-stat	IMB_a p-value	IMB_b p-value	TFI p-value	PRET p-value	DIV p-value
i=m1, j=m1	88.26	-79.91	82.37	14.41	0.00	0.00	0.00	0.0	0.00	0.00
i=m1, j=m2	52.70	-41.83	42.66	85.17	75.42	0.00	0.00	0.0	0.00	0.00
i=m1, j=m3	56.80	-56.04	51.03	59.07	2.78	0.00	0.00	0.0	0.00	0.01
i=m1, j=m4	39.28	-24.55	60.34	64.63	48.47	0.00	0.00	0.0	0.00	0.00
i=m1, j=m5	73.35	-74.25	18.61	80.13	-8.36	0.00	0.00	0.0	0.00	0.00
i=m1, j=m6	7.95	-0.96	36.20	100.94	79.07	0.00	0.33	0.0	0.00	0.00
i=m1, j=m7	48.24	-40.49	57.14	86.57	81.40	0.00	0.00	0.0	0.00	0.00
i=m1, j=m8	21.00	-16.93	51.09	91.29	87.88	0.00	0.00	0.0	0.00	0.00
i=m1, j=m9	41.75	-35.50	53.47	67.65	59.71	0.00	0.00	0.0	0.00	0.00
i=m1, j=m10	-17.19	25.78	34.24	127.55	50.09	0.00	0.00	0.0	0.00	0.00
i=m1, j=m11	21.28	-17.90	58.17	90.34	80.62	0.00	0.00	0.0	0.00	0.00
i=m1, j=m12	76.07	-78.25	24.34	62.43	37.05	0.00	0.00	0.0	0.00	0.00
i=m1, j=m13	77.26	-75.75	25.86	90.92	41.14	0.00	0.00	0.0	0.00	0.00
i=m1, j=m14	42.49	-38.15	47.26	100.55	77.12	0.00	0.00	0.0	0.00	0.00
i=m2, j=m1	37.61	-35.15	50.91	1.87	-19.89	0.00	0.00	0.0	0.06	0.00
i=m2, j=m2	85.37	-69.81	104.80	-26.87	0.00	0.00	0.00	0.0	0.00	0.00
i=m2, j=m3	10.76	-11.84	33.19	13.24	-48.58	0.00	0.00	0.0	0.00	0.00
i=m2, j=m4	13.99	-18.62	44.35	20.65	-40.22	0.00	0.00	0.0	0.00	0.00
i=m2, j=m5	21.97	-25.45	18.45	14.33	-46.30	0.00	0.00	0.0	0.00	0.00
i=m2, j=m6	9.32	-1.37	32.47	111.78	3.19	0.00	0.17	0.0	0.00	0.00
i=m2, j=m7	44.84	-41.26	70.91	28.80	0.62	0.00	0.00	0.0	0.00	0.54
i=m2, j=m8	28.03	-23.57	69.01	49.49	16.17	0.00	0.00	0.0	0.00	0.00
i=m2, j=m9	29.55	-22.24	51.49	25.40	-28.85	0.00	0.00	0.0	0.00	0.00
i=m2, j=m10	-4.30	8.60	40.00	76.10	-27.34	0.00	0.00	0.0	0.00	0.00
i=m2, j=m11	15.16	-10.83	72.36	35.82	-20.08	0.00	0.00	0.0	0.00	0.00
i=m2, j=m12	27.13	-24.14	14.10	12.56	-36.38	0.00	0.00	0.0	0.00	0.00

Continued on next page

Table 18: Summary statistics for all models M_{ij}

	IMB_a	IMB_b	TFI	PRET	DIV	IMB_a	IMB_b	TFI	PRET	DIV
	t-stat	t-stat	t-stat	t-stat	t-stat	p-value	p-value	p-value	p-value	p-value
i=m2, j=m13	69.57	-71.47	38.85	13.90	-8.92	0.00	0.00	0.0	0.00	0.00
i=m2, j=m14	56.42	-53.35	74.29	38.65	-3.34	0.00	0.00	0.0	0.00	0.00
i=m3, j=m1	64.45	-54.84	39.90	82.26	55.99	0.00	0.00	0.0	0.00	0.00
i=m3, j=m2	46.27	-35.89	31.48	82.22	87.82	0.00	0.00	0.0	0.00	0.00
i=m3, j=m3	73.99	-67.93	84.21	2.07	0.00	0.00	0.00	0.0	0.04	0.00
i=m3, j=m4	34.72	-17.89	55.36	67.19	76.50	0.00	0.00	0.0	0.00	0.00
i=m3, j=m5	73.68	-72.74	16.22	97.16	29.60	0.00	0.00	0.0	0.00	0.00
i=m3, j=m6	11.46	-0.57	30.37	83.37	99.94	0.00	0.57	0.0	0.00	0.00
i=m3, j=m7	40.49	-34.86	42.77	89.66	96.25	0.00	0.00	0.0	0.00	0.00
i=m3, j=m8	18.42	-15.41	40.26	86.91	102.18	0.00	0.00	0.0	0.00	0.00
i=m3, j=m9	41.86	-32.18	39.38	71.57	80.84	0.00	0.00	0.0	0.00	0.00
i=m3, j=m10	-12.00	22.15	33.21	115.94	73.95	0.00	0.00	0.0	0.00	0.00
i=m3, j=m11	27.98	-20.74	43.89	89.08	99.09	0.00	0.00	0.0	0.00	0.00
i=m3, j=m12	72.20	-71.00	23.56	71.70	67.39	0.00	0.00	0.0	0.00	0.00
i=m3, j=m13	60.33	-59.21	14.22	98.58	59.94	0.00	0.00	0.0	0.00	0.00
i=m3, j=m14	36.98	-29.29	29.47	99.31	89.66	0.00	0.00	0.0	0.00	0.00
i=m4, j=m1	57.53	-56.50	46.79	56.16	23.50	0.00	0.00	0.0	0.00	0.00
i=m4, j=m2	49.07	-43.12	37.93	77.41	53.73	0.00	0.00	0.0	0.00	0.00
i=m4, j=m3	36.77	-37.23	38.85	46.20	-16.28	0.00	0.00	0.0	0.00	0.00
i=m4, j=m4	37.33	-33.19	77.87	-28.07	0.00	0.00	0.00	0.0	0.00	0.00
i=m4, j=m5	50.75	-51.59	13.31	69.64	-13.72	0.00	0.00	0.0	0.00	0.00
i=m4, j=m6	12.16	-2.35	32.06	92.58	50.72	0.00	0.02	0.0	0.00	0.00
i=m4, j=m7	42.41	-39.77	49.56	77.27	54.55	0.00	0.00	0.0	0.00	0.00
i=m4, j=m8	21.77	-18.18	48.24	82.04	60.00	0.00	0.00	0.0	0.00	0.00
i=m4, j=m9	34.83	-31.26	44.75	68.55	28.09	0.00	0.00	0.0	0.00	0.00
i=m4, j=m10	-13.34	21.91	32.00	111.09	20.94	0.00	0.00	0.0	0.00	0.00
i=m4, j=m11	23.71	-21.14	54.33	81.29	48.08	0.00	0.00	0.0	0.00	0.00
i=m4, j=m12	62.06	-61.22	16.97	54.76	23.62	0.00	0.00	0.0	0.00	0.00
i=m4, j=m13	58.27	-67.46	26.12	78.89	29.11	0.00	0.00	0.0	0.00	0.00
i=m4, j=m14	43.19	-37.29	42.94	91.44	47.29	0.00	0.00	0.0	0.00	0.00
i=m5, j=m1	66.52	-60.26	51.99	91.85	87.10	0.00	0.00	0.0	0.00	0.00
i=m5, j=m2	48.96	-40.24	39.77	89.47	118.13	0.00	0.00	0.0	0.00	0.00
i=m5, j=m3	63.94	-60.74	61.04	65.10	64.23	0.00	0.00	0.0	0.00	0.00
i=m5, j=m4	42.40	-21.37	59.36	72.19	108.71	0.00	0.00	0.0	0.00	0.00
i=m5, j=m5	117.57	-125.39	30.34	61.82	0.00	0.00	0.00	0.0	0.00	0.00
i=m5, j=m6	9.56	2.85	43.67	99.09	129.11	0.00	0.00	0.0	0.00	0.00
i=m5, j=m7	44.98	-39.81	61.14	96.36	125.83	0.00	0.00	0.0	0.00	0.00
i=m5, j=m8	20.17	-19.93	55.49	93.43	130.25	0.00	0.00	0.0	0.00	0.00
i=m5, j=m9	45.77	-38.75	50.48	65.06	111.28	0.00	0.00	0.0	0.00	0.00
i=m5, j=m10	-12.11	23.18	39.89	134.52	101.42	0.00	0.00	0.0	0.00	0.00
i=m5, j=m11	31.18	-27.80	61.12	95.21	130.99	0.00	0.00	0.0	0.00	0.00
i=m5, j=m12	82.44	-83.85	32.84	63.61	100.33	0.00	0.00	0.0	0.00	0.00
i=m5, j=m13	66.47	-66.21	20.14	106.64	86.22	0.00	0.00	0.0	0.00	0.00
i=m5, j=m14	33.58	-30.17	43.67	108.06	123.22	0.00	0.00	0.0	0.00	0.00
i=m6, j=m1	32.04	-30.35	41.30	5.14	22.69	0.00	0.00	0.0	0.00	0.00
i=m6, j=m2	47.86	-40.95	54.45	13.57	45.86	0.00	0.00	0.0	0.00	0.00
i=m6, j=m3	6.30	-9.68	29.56	14.21	-12.39	0.00	0.00	0.0	0.00	0.00
i=m6, j=m4	14.70	-12.60	36.17	19.96	3.16	0.00	0.00	0.0	0.00	0.00
i=m6, j=m5	19.93	-23.56	18.17	16.29	-6.27	0.00	0.00	0.0	0.00	0.00
i=m6, j=m6	20.81	-5.67	48.00	6.62	0.00	0.00	0.00	0.0	0.00	0.00

Continued on next page

Table 18: Summary statistics for all models M_{ij}

	IMB_a	IMB_b	TFI	PRET	DIV	IMB_a	IMB_b	TFI	PRET	DIV
	t-stat	t-stat	t-stat	t-stat	t-stat	p-value	p-value	p-value	p-value	p-value
i=m6, j=m7	38.02	-35.32	56.83	16.33	36.78	0.00	0.00	0.0	0.00	0.00
i=m6, j=m8	25.72	-23.81	56.15	25.47	38.47	0.00	0.00	0.0	0.00	0.00
i=m6, j=m9	23.96	-20.31	41.55	15.91	0.90	0.00	0.00	0.0	0.00	0.37
i=m6, j=m10	-0.49	2.77	42.22	40.29	-1.55	0.62	0.01	0.0	0.00	0.12
i=m6, j=m11	14.43	-14.94	57.31	21.64	14.51	0.00	0.00	0.0	0.00	0.00
i=m6, j=m12	-13.67	7.58	14.43	12.08	-34.29	0.00	0.00	0.0	0.00	0.00
i=m6, j=m13	48.42	-48.42	32.91	9.01	30.49	0.00	0.00	0.0	0.00	0.00
i=m6, j=m14	49.72	-49.17	60.20	17.76	47.31	0.00	0.00	0.0	0.00	0.00
i=m7, j=m1	52.08	-48.19	47.58	24.20	-8.51	0.00	0.00	0.0	0.00	0.00
i=m7, j=m2	60.38	-48.03	45.19	70.30	43.21	0.00	0.00	0.0	0.00	0.00
i=m7, j=m3	22.74	-23.94	32.89	30.21	-44.53	0.00	0.00	0.0	0.00	0.00
i=m7, j=m4	18.83	-21.50	48.79	36.45	-32.33	0.00	0.00	0.0	0.00	0.00
i=m7, j=m5	27.52	-29.92	14.59	33.53	-49.59	0.00	0.00	0.0	0.00	0.00
i=m7, j=m6	9.98	0.60	23.81	113.68	22.22	0.00	0.55	0.0	0.00	0.00
i=m7, j=m7	54.59	-48.99	80.34	0.53	0.00	0.00	0.00	0.0	0.59	0.00
i=m7, j=m8	25.29	-20.75	52.16	74.04	41.80	0.00	0.00	0.0	0.00	0.00
i=m7, j=m9	33.87	-27.28	49.10	51.39	-8.33	0.00	0.00	0.0	0.00	0.00
i=m7, j=m10	-12.94	17.70	28.54	99.38	-10.86	0.00	0.00	0.0	0.00	0.00
i=m7, j=m11	15.87	-9.37	62.20	64.75	8.91	0.00	0.00	0.0	0.00	0.00
i=m7, j=m12	42.65	-40.13	13.86	34.22	-20.61	0.00	0.00	0.0	0.00	0.00
i=m7, j=m13	70.78	-74.36	31.11	52.82	6.15	0.00	0.00	0.0	0.00	0.00
i=m7, j=m14	52.30	-49.30	55.02	77.63	17.58	0.00	0.00	0.0	0.00	0.00
i=m8, j=m1	39.71	-36.55	47.16	12.50	-20.02	0.00	0.00	0.0	0.00	0.00
i=m8, j=m2	65.15	-52.37	60.34	54.44	29.54	0.00	0.00	0.0	0.00	0.00
i=m8, j=m3	11.25	-12.92	31.91	18.63	-53.56	0.00	0.00	0.0	0.00	0.00
i=m8, j=m4	12.71	-15.82	43.95	26.93	-49.11	0.00	0.00	0.0	0.00	0.00
i=m8, j=m5	23.09	-24.49	13.60	24.85	-50.42	0.00	0.00	0.0	0.00	0.00
i=m8, j=m6	9.55	-0.49	26.58	115.65	-1.39	0.00	0.62	0.0	0.00	0.17
i=m8, j=m7	44.71	-40.97	64.67	38.31	-0.61	0.00	0.00	0.0	0.00	0.54
i=m8, j=m8	30.30	-25.96	92.09	-6.80	0.00	0.00	0.00	0.0	0.00	0.00
i=m8, j=m9	28.51	-21.38	49.36	35.22	-32.42	0.00	0.00	0.0	0.00	0.00
i=m8, j=m10	-10.21	13.18	32.03	85.19	-30.78	0.00	0.00	0.0	0.00	0.00
i=m8, j=m11	11.23	-6.83	66.04	49.58	-26.65	0.00	0.00	0.0	0.00	0.00
i=m8, j=m12	28.02	-24.95	13.92	20.73	-37.71	0.00	0.00	0.0	0.00	0.00
i=m8, j=m13	67.32	-69.01	34.42	32.14	-7.83	0.00	0.00	0.0	0.00	0.00
i=m8, j=m14	51.53	-49.42	64.52	56.45	-4.87	0.00	0.00	0.0	0.00	0.00
i=m9, j=m1	72.29	-69.82	75.31	78.52	46.49	0.00	0.00	0.0	0.00	0.00
i=m9, j=m2	73.68	-59.21	68.32	136.46	72.25	0.00	0.00	0.0	0.00	0.00
i=m9, j=m3	30.03	-33.36	62.32	59.26	-13.39	0.00	0.00	0.0	0.00	0.00
i=m9, j=m4	27.69	-21.18	85.65	74.60	13.96	0.00	0.00	0.0	0.00	0.00
i=m9, j=m5	63.10	-63.29	4.44	99.21	0.24	0.00	0.00	0.0	0.00	0.81
i=m9, j=m6	15.90	-2.66	36.14	172.97	66.79	0.00	0.01	0.0	0.00	0.00
i=m9, j=m7	66.16	-61.36	86.23	112.67	83.41	0.00	0.00	0.0	0.00	0.00
i=m9, j=m8	32.39	-32.48	84.31	131.57	83.55	0.00	0.00	0.0	0.00	0.00
i=m9, j=m9	89.97	-81.26	85.67	-2.31	0.00	0.00	0.00	0.0	0.02	0.00
i=m9, j=m10	-6.30	10.48	49.30	202.94	16.60	0.00	0.00	0.0	0.00	0.00
i=m9, j=m11	39.89	-33.77	93.05	113.62	70.25	0.00	0.00	0.0	0.00	0.00
i=m9, j=m12	100.15	-99.86	20.85	66.37	70.86	0.00	0.00	0.0	0.00	0.00
i=m9, j=m13	95.04	-98.92	42.32	124.99	45.21	0.00	0.00	0.0	0.00	0.00
i=m9, j=m14	77.05	-71.58	73.04	144.50	65.13	0.00	0.00	0.0	0.00	0.00

Continued on next page

Table 18: Summary statistics for all models M_{ij}

	IMB_a	IMB_b	TFI	PRET	DIV	IMB_a	IMB_b	TFI	PRET	DIV
	t-stat	t-stat	t-stat	t-stat	t-stat	p-value	p-value	p-value	p-value	p-value
i=m10, j=m1	46.17	-44.72	55.00	-4.99	52.76	0.00	0.00	0.0	0.00	0.00
i=m10, j=m2	64.70	-51.88	70.44	10.13	66.93	0.00	0.00	0.0	0.00	0.00
i=m10, j=m3	21.53	-22.87	39.75	9.27	20.75	0.00	0.00	0.0	0.00	0.00
i=m10, j=m4	23.76	-19.52	46.56	16.56	41.29	0.00	0.00	0.0	0.00	0.00
i=m10, j=m5	38.57	-41.35	15.39	15.61	24.68	0.00	0.00	0.0	0.00	0.00
i=m10, j=m6	13.03	-3.49	45.38	78.00	47.11	0.00	0.00	0.0	0.00	0.00
i=m10, j=m7	52.39	-49.02	63.57	7.02	68.23	0.00	0.00	0.0	0.00	0.00
i=m10, j=m8	32.73	-30.26	69.92	19.55	70.47	0.00	0.00	0.0	0.00	0.00
i=m10, j=m9	43.96	-39.45	51.72	10.32	44.46	0.00	0.00	0.0	0.00	0.00
i=m10, j=m10	20.60	-17.15	90.06	-1.45	0.00	0.00	0.00	0.0	0.15	0.00
i=m10, j=m11	27.22	-28.10	69.45	13.35	54.65	0.00	0.00	0.0	0.00	0.00
i=m10, j=m12	48.98	-49.03	12.64	5.91	46.48	0.00	0.00	0.0	0.00	0.00
i=m10, j=m13	74.58	-74.60	42.30	0.17	59.66	0.00	0.00	0.0	0.87	0.00
i=m10, j=m14	63.04	-60.89	74.73	12.85	67.88	0.00	0.00	0.0	0.00	0.00
i=m11, j=m1	52.19	-47.47	43.83	32.46	-15.01	0.00	0.00	0.0	0.00	0.00
i=m11, j=m2	60.49	-47.34	47.22	78.16	35.05	0.00	0.00	0.0	0.00	0.00
i=m11, j=m3	21.46	-22.08	35.35	37.61	-51.44	0.00	0.00	0.0	0.00	0.00
i=m11, j=m4	18.53	-20.75	50.71	41.82	-39.12	0.00	0.00	0.0	0.00	0.00
i=m11, j=m5	33.78	-33.68	11.71	46.28	-50.16	0.00	0.00	0.0	0.00	0.00
i=m11, j=m6	8.29	0.31	20.05	122.20	16.30	0.00	0.75	0.0	0.00	0.00
i=m11, j=m7	53.38	-47.53	57.53	63.11	27.61	0.00	0.00	0.0	0.00	0.00
i=m11, j=m8	25.31	-22.04	52.63	81.15	36.05	0.00	0.00	0.0	0.00	0.00
i=m11, j=m9	36.19	-27.84	48.25	59.43	-15.48	0.00	0.00	0.0	0.00	0.00
i=m11, j=m10	-15.43	22.08	25.77	106.89	-17.18	0.00	0.00	0.0	0.00	0.00
i=m11, j=m11	19.58	-12.26	85.45	4.52	0.00	0.00	0.00	0.0	0.00	0.00
i=m11, j=m12	41.94	-39.37	15.12	39.24	-24.49	0.00	0.00	0.0	0.00	0.00
i=m11, j=m13	72.43	-73.27	28.31	61.26	-1.96	0.00	0.00	0.0	0.00	0.05
i=m11, j=m14	52.14	-46.91	53.09	85.66	9.06	0.00	0.00	0.0	0.00	0.00
i=m12, j=m1	64.16	-57.76	36.18	94.82	31.26	0.00	0.00	0.0	0.00	0.00
i=m12, j=m2	46.02	-38.32	27.83	92.17	69.56	0.00	0.00	0.0	0.00	0.00
i=m12, j=m3	48.79	-47.94	47.14	72.47	2.02	0.00	0.00	0.0	0.00	0.04
i=m12, j=m4	35.84	-21.42	59.82	72.28	47.65	0.00	0.00	0.0	0.00	0.00
i=m12, j=m5	65.40	-67.54	13.70	103.30	-2.43	0.00	0.00	0.0	0.00	0.02
i=m12, j=m6	7.99	1.27	31.81	91.95	84.47	0.00	0.21	0.0	0.00	0.00
i=m12, j=m7	40.31	-35.24	46.01	98.68	76.23	0.00	0.00	0.0	0.00	0.00
i=m12, j=m8	18.81	-16.99	41.83	94.65	81.72	0.00	0.00	0.0	0.00	0.00
i=m12, j=m9	41.13	-33.24	41.07	80.64	59.69	0.00	0.00	0.0	0.00	0.00
i=m12, j=m10	-18.82	25.04	27.53	128.04	51.71	0.00	0.00	0.0	0.00	0.00
i=m12, j=m11	28.67	-25.59	45.72	97.68	80.03	0.00	0.00	0.0	0.00	0.00
i=m12, j=m12	77.74	-78.01	39.34	21.21	0.00	0.00	0.00	0.0	0.00	0.00
i=m12, j=m13	60.61	-62.81	15.22	109.50	39.05	0.00	0.00	0.0	0.00	0.00
i=m12, j=m14	32.74	-28.48	28.44	112.38	73.16	0.00	0.00	0.0	0.00	0.00
i=m13, j=m1	69.25	-66.57	71.54	32.25	14.26	0.00	0.00	0.0	0.00	0.00
i=m13, j=m2	67.44	-54.84	71.21	88.19	59.52	0.00	0.00	0.0	0.00	0.00
i=m13, j=m3	34.97	-36.87	51.63	38.93	-24.66	0.00	0.00	0.0	0.00	0.00
i=m13, j=m4	31.24	-23.90	69.65	48.90	6.05	0.00	0.00	0.0	0.00	0.00
i=m13, j=m5	53.36	-56.24	17.99	55.10	-24.92	0.00	0.00	0.0	0.00	0.00
i=m13, j=m6	11.72	-1.77	38.10	140.88	55.37	0.00	0.08	0.0	0.00	0.00
i=m13, j=m7	57.00	-50.90	83.18	71.25	57.67	0.00	0.00	0.0	0.00	0.00
i=m13, j=m8	28.26	-24.43	75.86	92.48	68.06	0.00	0.00	0.0	0.00	0.00

Continued on next page

Table 18: Summary statistics for all models M_{ij}

	IMB_a	IMB_b	TFI	PRET	DIV	IMB_a	IMB_b	TFI	PRET	DIV
	t-stat	t-stat	t-stat	t-stat	t-stat	p-value	p-value	p-value	p-value	p-value
i=m13, j=m9	51.58	-43.95	62.17	57.57	21.25	0.00	0.00	0.0	0.00	0.00
i=m13, j=m10	-7.24	14.07	47.38	141.30	17.98	0.00	0.00	0.0	0.00	0.00
i=m13, j=m11	27.34	-20.04	86.77	75.04	51.32	0.00	0.00	0.0	0.00	0.00
i=m13, j=m12	65.38	-64.74	19.48	42.64	9.33	0.00	0.00	0.0	0.00	0.00
i=m13, j=m13	133.64	-138.28	71.83	6.18	0.00	0.00	0.00	0.0	0.00	0.00
i=m13, j=m14	70.25	-65.71	80.37	91.94	45.25	0.00	0.00	0.0	0.00	0.00
i=m14, j=m1	59.81	-57.90	53.67	22.00	4.81	0.00	0.00	0.0	0.00	0.00
i=m14, j=m2	66.62	-53.92	63.20	71.24	48.42	0.00	0.00	0.0	0.00	0.00
i=m14, j=m3	22.79	-25.24	39.63	33.64	-39.43	0.00	0.00	0.0	0.00	0.00
i=m14, j=m4	24.74	-22.06	59.63	41.53	-13.87	0.00	0.00	0.0	0.00	0.00
i=m14, j=m5	38.04	-40.29	15.62	42.00	-36.49	0.00	0.00	0.0	0.00	0.00
i=m14, j=m6	10.51	-0.97	26.96	123.55	49.13	0.00	0.33	0.0	0.00	0.00
i=m14, j=m7	59.51	-50.94	69.46	56.91	48.20	0.00	0.00	0.0	0.00	0.00
i=m14, j=m8	30.23	-24.23	65.08	77.02	55.03	0.00	0.00	0.0	0.00	0.00
i=m14, j=m9	44.10	-37.08	60.74	55.99	5.04	0.00	0.00	0.0	0.00	0.00
i=m14, j=m10	-10.04	13.88	38.59	118.39	3.25	0.00	0.00	0.0	0.00	0.00
i=m14, j=m11	20.05	-14.62	74.17	65.55	30.41	0.00	0.00	0.0	0.00	0.00
i=m14, j=m12	49.01	-45.76	17.84	36.11	-8.45	0.00	0.00	0.0	0.00	0.00
i=m14, j=m13	88.00	-90.83	39.63	48.04	15.42	0.00	0.00	0.0	0.00	0.00
i=m14, j=m14	84.18	-77.87	97.41	20.98	0.00	0.00	0.00	0.0	0.00	0.00

Table 19: T-statistics for all baseline models

		m1	m2	m3	m4	m5	m6	m7	m8	m9	m10	m11	m12	m13	m14
IMBa	m1	32.3	10.7	7.7	15.2	8.1	13.1	15.2	12.2	14.1	11.0	14.2	13.7	15.9	15.2
	m2	3.3	32.8	2.8	4.5	5.0	17.9	9.9	17.2	10.4	24.3	9.4	2.1	10.1	9.4
	m3	-3.6	-6.5	21.0	-1.2	-8.1	-11.5	-1.5	-4.2	-10.7	-8.1	-2.6	-3.7	-9.3	-7.1
	m4	-4.8	-5.2	-7.5	3.8	-4.6	-1.6	-3.3	-3.5	-5.5	-5.3	-3.4	-0.7	-6.3	-3.1
	m5	9.7	0.2	6.8	6.2	49.8	2.4	1.3	2.6	3.6	2.7	4.5	8.5	2.9	4.9
	m6	1.4	3.8	5.6	7.0	2.6	12.0	4.8	4.8	8.4	7.9	3.1	3.8	3.5	4.0
	m7	1.9	7.7	2.7	1.3	4.3	11.4	9.8	8.1	3.1	10.8	8.4	1.2	1.6	7.3
	m8	0.2	0.8	-0.0	2.2	-0.9	6.5	-1.1	5.0	0.2	4.9	-0.7	0.0	-1.6	0.8
	m9	5.0	8.2	6.4	1.8	3.7	14.3	7.9	9.0	23.6	10.9	9.4	4.2	8.1	9.8
	m10	-4.2	1.7	-3.8	-0.9	-5.5	0.6	0.3	-0.2	-2.0	13.6	-1.0	-8.2	-1.7	-2.1
	m11	-7.3	-5.6	-0.7	-2.5	-0.8	-4.1	-6.8	-8.0	-0.4	-4.3	-7.8	1.3	-3.8	-9.7
	m12	-11.0	-17.4	-7.5	-9.3	-6.9	-32.4	-9.3	-14.8	-8.8	-27.0	-9.3	5.7	-19.7	-16.6
	m13	14.1	20.4	4.3	4.7	5.6	15.6	16.2	20.3	15.5	20.2	17.7	6.3	58.9	23.4
	m14	7.0	18.8	7.5	9.6	6.3	20.4	11.6	15.7	18.7	21.3	12.4	3.7	19.1	29.2
IMBb	m1	-29.5	-10.3	-5.4	-13.4	-4.1	-10.8	-13.6	-11.9	-13.2	-11.4	-13.9	-9.9	-18.5	-18.4
	m2	-6.9	-31.2	-6.6	-8.5	-8.4	-22.0	-11.6	-18.4	-10.7	-24.6	-10.7	-5.0	-14.7	-14.2
	m3	-3.5	0.7	-22.4	0.5	4.9	5.5	-2.3	-0.5	5.9	3.4	-2.2	-0.4	3.8	0.9
	m4	6.1	2.1	8.2	-8.3	6.9	4.1	1.4	2.7	4.8	2.9	1.2	2.2	8.3	6.2
	m5	-12.5	-6.9	-8.1	-8.1	-58.4	-6.6	-4.1	-6.1	-5.8	-6.8	-6.7	-12.2	-7.3	-9.2
	m6	0.3	0.0	-0.2	-1.5	3.7	-3.9	1.0	0.5	-1.4	-1.1	0.7	1.0	0.1	-0.6
	m7	2.0	-3.9	-2.5	-0.2	-3.0	-7.0	-6.9	-5.1	-2.1	-7.7	-6.0	0.9	2.0	-0.2
	m8	1.2	-1.8	-0.1	0.1	-1.1	-7.4	2.1	-5.1	-3.0	-6.0	-0.7	0.5	2.8	2.4
	m9	-4.2	-6.1	-2.5	-4.4	0.0	-11.5	-6.2	-6.1	-24.3	-10.4	-7.3	-0.1	-5.6	-9.0
	m10	6.2	-0.2	7.0	6.2	8.4	-2.5	2.5	2.2	1.8	-17.9	5.5	8.1	4.4	0.9
	m11	6.6	6.5	3.3	3.4	1.5	1.1	11.1	8.3	5.9	2.4	10.2	-0.9	11.6	11.4

Continued on next page

Table 19: T-statistics for all baseline models

		m1	m2	m3	m4	m5	m6	m7	m8	m9	m10	m11	m12	m13	m14
	m12	11.2	19.3	10.2	11.5	10.5	32.2	11.3	16.9	7.1	27.4	11.0	-5.2	21.0	18.9
	m13	-6.7	-16.5	-5.5	-11.6	-5.3	-10.2	-15.9	-16.5	-15.6	-13.8	-15.6	-7.1	-59.3	-20.6
	m14	-9.5	-19.8	-6.3	-5.1	-7.5	-23.4	-14.4	-17.7	-15.0	-20.6	-12.1	-3.4	-15.1	-27.4
TFI	m1	40.9	15.5	8.0	14.0	11.0	14.6	13.8	11.5	11.5	16.9	8.5	2.3	21.6	14.7
	m2	5.9	45.8	4.7	5.8	2.2	28.8	8.1	19.3	9.6	34.0	8.9	-2.3	15.2	16.7
	m3	8.6	5.6	41.9	5.6	10.4	8.6	3.2	3.8	5.7	9.0	3.6	3.4	8.4	5.7
	m4	-3.3	0.8	1.7	9.3	-2.8	4.0	-2.5	-0.2	4.1	4.4	-2.0	2.2	-0.8	-0.4
	m5	7.4	6.8	6.8	4.6	22.5	9.5	4.5	3.4	-2.3	7.0	2.2	5.1	6.8	6.4
	m6	14.4	10.9	12.6	13.9	19.4	24.0	7.7	7.4	11.6	23.4	4.6	15.9	11.7	6.2
	m7	13.3	20.6	6.3	10.0	15.3	23.5	22.3	17.7	14.7	22.9	13.4	9.3	18.6	17.8
	m8	8.0	18.8	5.0	9.8	9.4	24.8	9.0	28.7	15.5	26.2	7.5	4.3	14.2	13.1
	m9	5.2	11.2	-0.9	3.1	2.2	17.1	5.8	9.9	11.1	16.9	5.5	-1.1	4.1	12.0
	m10	4.3	12.4	7.4	6.2	5.3	21.5	5.0	7.8	6.3	41.6	2.7	2.1	8.6	8.0
	m11	8.5	18.1	4.6	8.8	11.2	21.9	11.8	14.2	12.8	23.2	21.2	4.6	16.4	15.1
	m12	-0.6	-3.1	1.2	-2.2	4.4	-2.1	-2.9	-3.0	-2.6	-0.8	-3.5	11.8	-3.6	-2.6
	m13	1.5	9.9	0.1	4.9	-0.3	13.4	5.8	7.5	7.6	13.9	4.2	0.4	36.7	12.0
	m14	11.1	30.4	2.6	11.0	7.0	34.4	18.4	24.5	22.3	36.5	16.4	2.1	26.0	42.9
PRET	m1	-49.8	-5.8	22.5	16.3	30.6	-3.2	1.0	-4.9	8.5	-6.0	1.3	28.5	3.1	-1.3
	m2	-4.7	-97.5	-2.7	-1.8	-6.8	-9.1	4.1	5.9	8.1	-3.8	5.5	-5.8	2.3	3.0
	m3	11.7	-3.4	-55.2	10.4	14.9	0.0	1.5	-2.6	12.2	-3.6	6.2	20.0	1.8	3.9
	m4	9.4	-0.7	13.4	-99.3	13.7	3.4	-0.3	1.0	6.4	2.4	-0.5	9.6	1.7	3.9
	m5	17.8	-7.9	26.4	12.8	-22.7	-3.0	-2.5	-3.3	14.5	-7.3	2.8	28.7	1.8	-1.0
	m6	24.9	41.6	14.6	19.0	26.7	-63.2	28.3	34.2	50.0	66.0	28.6	22.0	43.4	44.6
	m7	12.1	-1.5	13.3	11.7	17.3	-4.4	-78.5	-1.1	10.7	-4.4	7.0	14.1	4.5	3.7
	m8	7.7	5.6	6.4	7.6	6.4	-3.2	8.1	-97.2	13.0	-2.5	6.9	3.9	9.4	7.1
	m9	3.0	2.4	10.9	15.6	0.7	2.4	6.5	4.9	-119.6	3.7	11.5	12.4	7.6	13.7
	m10	-5.9	6.1	-4.9	-5.3	-4.4	6.2	-0.9	6.5	15.9	-134.0	1.7	-3.7	1.7	6.1
	m11	15.1	2.3	14.4	12.1	18.4	-3.1	10.7	8.8	12.2	-5.4	-84.1	15.7	6.8	6.6
	m12	9.6	-6.2	20.5	8.7	10.7	-5.6	0.8	-3.5	14.1	-7.2	1.1	-59.5	1.5	-2.1
	m13	4.2	-11.8	7.7	5.3	4.8	-7.9	2.5	-4.2	20.5	-8.2	4.7	11.6	-79.6	3.8
	m14	18.1	5.5	10.8	14.1	16.5	-1.4	14.2	10.2	24.2	4.7	12.3	12.6	14.8	-77.5
DIV	m1	-21.8	-0.6	9.3	8.1	18.2	6.6	-0.6	-0.1	3.6	5.1	-1.8	6.7	1.9	-0.2
	m2	17.8	20.5	6.7	15.3	17.7	18.3	21.0	23.5	15.4	20.6	18.4	11.7	26.4	21.8
	m3	-3.2	-0.4	-4.3	-13.2	-5.5	-5.1	-1.7	-0.8	-10.9	-0.4	-4.3	-12.5	-4.5	-4.6
	m4	-0.6	0.4	-5.0	-8.7	-0.5	8.6	-1.6	-3.0	-5.5	5.2	-3.2	-0.6	0.4	4.4
	m5	-11.9	-1.0	-8.1	-8.7	-0.4	0.7	-11.8	-2.4	-3.8	3.9	-6.2	-13.4	-2.1	-3.2
	m6	-3.5	10.3	4.9	8.4	2.0	-12.4	7.8	11.6	27.0	17.3	9.7	2.5	6.8	11.3
	m7	-9.7	-3.1	-2.9	-6.5	-11.4	4.0	-4.1	-2.5	-0.5	-0.6	-0.5	-7.3	-8.7	-4.7
	m8	-6.1	-6.4	-1.1	-3.3	-9.6	-1.5	-6.8	-12.1	-1.5	-4.8	-4.9	-5.2	-9.5	-7.1
	m9	6.5	5.4	-0.8	-3.0	0.9	6.0	4.7	4.8	-26.7	2.1	2.9	-2.9	2.0	5.4
	m10	-7.1	-9.7	-9.2	-6.6	-10.6	-10.9	-3.7	-7.0	-6.6	-7.7	-4.9	-6.3	-12.1	-8.1
	m11	-8.0	-9.8	-5.3	-6.1	-5.5	-10.2	-9.5	-11.4	-7.4	-13.7	-0.4	-3.7	-9.6	-8.7
	m12	-16.7	-22.0	-13.5	-15.8	-9.6	-34.6	-16.4	-20.5	-15.6	-29.2	-16.3	-0.4	-24.6	-21.8
	m13	12.4	2.9	6.1	9.5	8.1	11.9	4.5	1.6	-1.3	7.6	2.7	2.6	-7.3	5.4
	m14	21.6	9.7	18.4	17.7	24.3	24.4	13.8	8.3	12.7	17.4	12.0	20.4	19.0	0.0

Table 20: p-values for all baseline models

		m1	m2	m3	m4	m5	m6	m7	m8	m9	m10	m11	m12	m13	m14
IMBa	m1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Continued on next page

Table 20: p-values for all baseline models

		m1	m2	m3	m4	m5	m6	m7	m8	m9	m10	m11	m12	m13	m14
	m2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.04	0.0	0.0
	m3	0.0	0.0	0.0	0.23	0.0	0.0	0.12	0.0	0.0	0.0	0.01	0.0	0.0	0.0
	m4	0.0	0.0	0.0	0.0	0.0	0.12	0.0	0.0	0.0	0.0	0.0	0.49	0.0	0.0
	m5	0.0	0.87	0.0	0.0	0.0	0.02	0.19	0.01	0.0	0.01	0.0	0.0	0.0	0.0
	m6	0.15	0.0	0.0	0.0	0.01	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m7	0.06	0.0	0.01	0.2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.23	0.1	0.0
	m8	0.86	0.42	0.99	0.03	0.36	0.0	0.26	0.0	0.84	0.0	0.5	0.99	0.12	0.42
	m9	0.0	0.0	0.0	0.07	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m10	0.0	0.08	0.0	0.39	0.0	0.52	0.79	0.85	0.05	0.0	0.32	0.0	0.09	0.03
	m11	0.0	0.0	0.47	0.01	0.43	0.0	0.0	0.0	0.66	0.0	0.0	0.21	0.0	0.0
	m12	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m13	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m14	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
IMBb	m1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m3	0.0	0.48	0.0	0.61	0.0	0.0	0.02	0.63	0.0	0.0	0.03	0.66	0.0	0.35
	m4	0.0	0.04	0.0	0.0	0.0	0.0	0.17	0.01	0.0	0.0	0.23	0.03	0.0	0.0
	m5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m6	0.74	0.99	0.83	0.13	0.0	0.0	0.33	0.65	0.15	0.28	0.51	0.33	0.95	0.56
	m7	0.05	0.0	0.01	0.85	0.0	0.0	0.0	0.0	0.03	0.0	0.0	0.37	0.05	0.86
	m8	0.22	0.07	0.92	0.92	0.27	0.0	0.03	0.0	0.0	0.0	0.46	0.62	0.0	0.02
	m9	0.0	0.0	0.01	0.0	0.96	0.0	0.0	0.0	0.0	0.0	0.0	0.92	0.0	0.0
	m10	0.0	0.86	0.0	0.0	0.0	0.01	0.01	0.02	0.07	0.0	0.0	0.0	0.0	0.38
	m11	0.0	0.0	0.0	0.0	0.13	0.28	0.0	0.0	0.0	0.02	0.0	0.39	0.0	0.0
	m12	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m13	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m14	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
TFI	m1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.02	0.0	0.0
	m2	0.0	0.0	0.0	0.0	0.03	0.0	0.0	0.0	0.0	0.0	0.0	0.02	0.0	0.0
	m3	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m4	0.0	0.41	0.09	0.0	0.01	0.0	0.01	0.81	0.0	0.0	0.05	0.03	0.41	0.65
	m5	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.02	0.0	0.03	0.0	0.0	0.0
	m6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m9	0.0	0.0	0.39	0.0	0.03	0.0	0.0	0.0	0.0	0.0	0.0	0.26	0.0	0.0
	m10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.01	0.03	0.0	0.0
	m11	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m12	0.54	0.0	0.24	0.03	0.0	0.04	0.0	0.0	0.01	0.4	0.0	0.0	0.0	0.01
	m13	0.14	0.0	0.96	0.0	0.75	0.0	0.0	0.0	0.0	0.0	0.0	0.71	0.0	0.0
	m14	0.0	0.0	0.01	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.04	0.0	0.0
PRET	m1	0.0	0.0	0.0	0.0	0.0	0.0	0.34	0.0	0.0	0.0	0.2	0.0	0.0	0.18
	m2	0.0	0.0	0.01	0.07	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.02	0.0
	m3	0.0	0.0	0.0	0.0	0.0	0.99	0.12	0.01	0.0	0.0	0.0	0.0	0.07	0.0
	m4	0.0	0.47	0.0	0.0	0.0	0.0	0.8	0.3	0.0	0.02	0.59	0.0	0.09	0.0
	m5	0.0	0.0	0.0	0.0	0.0	0.0	0.01	0.0	0.0	0.0	0.01	0.0	0.08	0.34
	m6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m7	0.0	0.14	0.0	0.0	0.0	0.0	0.0	0.28	0.0	0.0	0.0	0.0	0.0	0.0
	m8	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.01	0.0	0.0	0.0	0.0
	m9	0.0	0.02	0.0	0.0	0.51	0.02	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
	m10	0.0	0.0	0.0	0.0	0.0	0.0	0.34	0.0	0.0	0.0	0.09	0.0	0.09	0.0

Continued on next page

Table 21: In-sample R^2 values for LASSO models, for the 1000ms future return horizon, as the λ regularization parameter varies.

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
ftx_BTC-PERP	0.342	0.341	0.341	0.34	0.338	0.335	0.324	0.293	0.254
bybit_BTC/USD	0.331	0.331	0.33	0.328	0.325	0.324	0.317	0.298	0.25
bitmex_BTC/USD	0.292	0.292	0.291	0.29	0.288	0.285	0.275	0.259	0.212
tbybit_BTC/USDT	0.279	0.279	0.279	0.278	0.276	0.273	0.263	0.246	0.207
hbdm_BTC-USD	0.262	0.262	0.261	0.259	0.256	0.251	0.243	0.228	0.19
thbdm_BTC-USDT	0.231	0.231	0.23	0.229	0.225	0.22	0.207	0.19	0.144
deribit_BTC-PERPETUAL	0.212	0.212	0.212	0.211	0.21	0.206	0.197	0.173	0.131
okex_BTC-USDT-SWAP	0.195	0.195	0.195	0.194	0.193	0.19	0.182	0.153	0.117
binancecmfut_BTC/USD	0.193	0.193	0.193	0.191	0.189	0.185	0.171	0.152	0.106
okex_BTC-USD-SWAP	0.18	0.18	0.18	0.179	0.177	0.174	0.166	0.137	0.103
okex_BTC-USD-210326	0.164	0.164	0.163	0.163	0.16	0.157	0.148	0.121	0.078
hbdm_BTC_CQ	0.152	0.152	0.151	0.15	0.147	0.144	0.135	0.103	0.065
huobipro_BTC/USDT	0.136	0.136	0.136	0.135	0.131	0.125	0.113	0.074	0.04
binancefut_BTC/USDT	0.071	0.071	0.07	0.069	0.066	0.058	0.045	0.026	0.006

Table 20: p-values for all baseline models

	m1	m2	m3	m4	m5	m6	m7	m8	m9	m10	m11	m12	m13	m14
m11	0.0	0.02	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
m12	0.0	0.0	0.0	0.0	0.0	0.0	0.4	0.0	0.0	0.0	0.28	0.0	0.14	0.04
m13	0.0	0.0	0.0	0.0	0.0	0.0	0.01	0.0	0.0	0.0	0.0	0.0	0.0	0.0
m14	0.0	0.0	0.0	0.0	0.0	0.17	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
DIV														
m1	0.0	0.56	0.0	0.0	0.0	0.0	0.58	0.92	0.0	0.0	0.07	0.0	0.05	0.84
m2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
m3	0.0	0.69	0.0	0.0	0.0	0.0	0.1	0.45	0.0	0.66	0.0	0.0	0.0	0.0
m4	0.58	0.67	0.0	0.0	0.63	0.0	0.11	0.0	0.0	0.0	0.0	0.52	0.68	0.0
m5	0.0	0.32	0.0	0.0	0.68	0.51	0.0	0.02	0.0	0.0	0.0	0.0	0.04	0.0
m6	0.0	0.0	0.0	0.0	0.04	0.0	0.0	0.0	0.0	0.0	0.0	0.01	0.0	0.0
m7	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.01	0.63	0.58	0.61	0.0	0.0	0.0
m8	0.0	0.0	0.26	0.0	0.0	0.14	0.0	0.0	0.13	0.0	0.0	0.0	0.0	0.0
m9	0.0	0.0	0.45	0.0	0.36	0.0	0.0	0.0	0.0	0.04	0.0	0.0	0.05	0.0
m10	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
m11	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.72	0.0	0.0	0.0
m12	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.66	0.0	0.0
m13	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.11	0.18	0.0	0.01	0.01	0.0	0.0
m14	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0

Table 22: Out-of-sample R^2 values for LASSO models, for the 1000ms future return horizon, as the λ regularization parameter varies.

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
bybit_BTC/USD	0.319	0.319	0.319	0.316	0.313	0.312	0.308	0.299	0.292
tbybit_BTC/USDT	0.28	0.28	0.28	0.279	0.278	0.276	0.27	0.259	0.248
bitmex_BTC/USD	0.27	0.27	0.27	0.269	0.268	0.266	0.26	0.253	0.24
ftx_BTC-PERP	0.264	0.264	0.263	0.263	0.261	0.259	0.25	0.219	0.2
hbdm_BTC-USD	0.243	0.243	0.243	0.241	0.239	0.234	0.228	0.222	0.216
deribit_BTC-PERPETUAL	0.215	0.215	0.215	0.214	0.214	0.213	0.207	0.195	0.175
thbdm_BTC-USDT	0.215	0.215	0.215	0.214	0.212	0.208	0.198	0.189	0.175
okex_BTC-USDT-SWAP	0.184	0.184	0.184	0.183	0.183	0.181	0.175	0.153	0.131
binancecmfut_BTC/USD	0.177	0.177	0.177	0.176	0.173	0.17	0.158	0.146	0.121
okex_BTC-USD-SWAP	0.167	0.167	0.167	0.167	0.166	0.164	0.158	0.135	0.119
okex_BTC-USD-210326	0.155	0.155	0.154	0.154	0.151	0.15	0.144	0.125	0.094
hbdm_BTC_CQ	0.144	0.144	0.144	0.143	0.14	0.137	0.132	0.106	0.087
huobipro_BTC/USDT	0.126	0.126	0.126	0.125	0.123	0.119	0.112	0.077	0.051
binancefut_BTC/USDT	0.07	0.07	0.07	0.069	0.068	0.062	0.051	0.032	0.019

Table 23: PnL_1 values of LASSO models, as the λ regularization parameter varies.

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
huobipro_BTC/USDT	28867	28774	28551	28276	27874	27723	27336	24912	22400
thbdm_BTC-USDT	20621	20522	20084	19501	19073	17759	16663	17153	18144
bybit_BTC/USD	19718	19654	19590	19736	19604	19934	20098	20410	20468
ftx_BTC-PERP	19689	19519	19406	19306	19533	19556	19627	19245	18516
okex_BTC-USD-SWAP	17891	17757	17446	17046	16627	16036	14545	13236	15272
bitmex_BTC/USD	17720	17613	17603	17663	17663	17726	17676	17692	17490
hbdm_BTC-USD	17663	17601	17451	17340	17098	16810	17800	18314	18316
okex_BTC-USD-210326	17268	17073	16602	15645	15248	14603	13585	11602	12871
tbybit_BTC/USDT	17069	17037	17050	17078	17073	17099	17222	17669	17277
binancecmfut_BTC/USD	16328	16188	15841	15201	15301	15067	14331	13712	15573
hbdm_BTC_CQ	16232	16038	15712	14931	14183	13818	12942	10924	9728
okex_BTC-USDT-SWAP	14961	14894	14527	14241	14222	14120	13463	12647	14095
binancefut_BTC/USDT	12085	12057	11876	11898	11559	11408	10672	10325	2793
deribit_BTC-PERPETUAL	10313	10218	9995	9730	9691	9641	9638	8762	8845

Table 24: PnL_2 values of LASSO models, as the λ regularization parameter varies.

	0.001	0.002	0.004	0.008	0.016	0.032	0.064	0.128	0.256
hbdm_BTC-USD	-8256	-8308	-8358	-8189	-8461	-8749	-9609	-10755	-12463
thbdm_BTC-USDT	-10674	-10645	-10755	-10890	-11030	-11152	-11768	-13774	-20751
bybit_BTC/USD	-11631	-11440	-11504	-11808	-11895	-12510	-13771	-15709	-19746
tbybit_BTC/USDT	-15855	-15917	-16264	-16521	-16601	-17025	-17592	-19935	-25232
bitmex_BTC/USD	-16209	-16091	-16341	-16386	-16311	-15813	-16328	-18577	-23969
binancefut_BTC/USDT	-17506	-17310	-16963	-15877	-14408	-12311	-11751	-13162	-1118
okex_BTC-USD-SWAP	-19228	-19132	-19303	-19663	-19942	-20573	-20604	-21873	-30217
deribit_BTC-PERPETUAL	-19666	-19621	-19784	-20189	-20368	-20438	-20441	-20387	-25174
okex_BTC-USDT-SWAP	-20058	-20105	-20382	-20778	-21067	-22089	-21896	-22322	-28594
binancecmfut_BTC/USD	-21871	-21841	-21968	-22018	-21668	-21632	-21278	-24877	-33766
hbdm_BTC_CQ	-23391	-23377	-23135	-23140	-22648	-21941	-21201	-21883	-28583
okex_BTC-USD-210326	-27171	-27086	-27257	-27214	-27181	-28156	-28464	-27847	-36518
ftx_BTC-PERP	-27504	-27520	-27423	-27607	-27548	-27805	-28238	-29278	-35285
huobipro_BTC/USDT	-28056	-28064	-27802	-27441	-26921	-25258	-22063	-14474	-23873




Figure 31: 500ms in-sample R^2 values




Figure 32: 1000ms out-of-sample R^2 values.




Figure 33: 1000ms in-sample R^2 values.




Figure 34: PnL_{L_2} values for each model.