Essays on Asset Management and Crypto Asset Pricing

by

Ainsley To

A dissertation submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

EDHEC Business School

6 July 2023

Dissertation Committee

Laurent Calvet, EDHEC Business School
Wei Xiong, Princeton University
Emmanuel Jurczenko, EDHEC Business School
Enrique Schroth, EDHEC Business School
Raman Uppal, EDHEC Business School

APPROVAL FORM

CANDIDATE'S NAME: Ainsley Nathan TO

TITLE OF DISSERTATION: **"Essays on Asset Management and Crypto Asset Pricing"**

APPROVAL:

Enrique Schroth		*Chair*
Wei Xiong	p/o	*External Examiner*
Laurent Calvet		*Member*
Emmanuel Jurczenko		*Member*
Raman Uppal		*Member*

☐ with minor revisions

☐ with major revisions

DATE: 6 July 2023

To my father, in loving memory

Abstract

The Delegation Gap: The Disposition Effect and Other People's Money: Using a unique dataset, this paper documents a "delegation gap" in investor trades - the disposition effect is reduced when investment decisions are fully delegated to an investment manager compared to when decisions are made by the account owner. The delegation gap is present in both retail and professional accounts, as well as when comparing trades made by investment managers on behalf of clients to those they made in their own personal accounts. The delegation gap varies across investor demographic groups and is positively related to economic news sentiment, suggesting a larger benefit to delegation during periods of negative sentiment. The evidence is consistent with the view that under certain conditions, delegating investment decisions can reduce behavioural mistakes made by investors trading emotionally with their own money.

Magical Internet Money? On-chain Cashflows and the Cross-section of Cryptocurrency Returns: I find that crypto valuation measures derived from on-chain fundamental cashflow characteristics, analogous to valuation metrics used in equity markets, are priced in the cross-section of token returns. A cashflow-based value factor constructed from these measures is not spanned by crypto factor models in the literature. I test different measures of cashflow and find that revenues retained by protocols show the strongest results, whilst token incentives as a cost of revenue measure have little pricing power. I also find evidence that different characteristics are significant for native tokens of a blockchain compared to tokens issued by decentralised applications. Lastly, I test a set of novel crypto native characteristics, unique to public blockchains, that proxy for capital gains overhang, insider ownership, and investor sophistication.

Contents

1 Acknowledgements

The author would like to thank his supervisor Laurent Calvet, and his thesis committee Wei Xiong, Emmanuel Jurczenko, Enrique Schroth, and Raman Uppal for their insightful comments. He would also like to thank Deon Gouws, Harry Houldsworth, Antti Ilmanen, Kasim Khorasanee, Calvin Mclean, and Lasse Pedersen for their suggestions. Additionally, he would like to thank Vikki To for her support and endless patience.

List of Figures

List of Tables

The Delegation Gap:
The Disposition Effect and Other People's Money

Ainsley To

6 July 2023

Abstract

Using a unique dataset, this paper documents a "delegation gap" in investor trades - the disposition effect is reduced when investment decisions are fully delegated to an investment manager compared to when decisions are made by the account owner. The delegation gap is present in both retail and professional accounts, as well as when comparing trades made by investment managers on behalf of clients to those they made in their own personal accounts. The delegation gap varies across investor demographic groups and is positively related to economic news sentiment, suggesting a larger benefit to delegation during periods of negative sentiment. The evidence is consistent with the view that under certain conditions, delegating investment decisions can reduce behavioural mistakes made by investors trading emotionally with their own money.

The author would like to thank the wealth management platform for providing the data and gratefully acknowledges helpful comments from Laurent Calvet.

3 Introduction

The disposition effect is referred to as “one of the most robust facts about the trading of individual investors" (Barberis and Xiong (2009)) and has been extensively documented empirically in the literature. Shefrin and Statman (1985) coined the term ‘disposition effect’ when using data from US mutual funds back to 1961 they found that investors are more likely to realise gains than to realise losses. The effect has subsequently been found in investors at a US discount brokerage (Odean (1998)), using logit regressions in a dataset of Finnish investors (Grinblatt and Keloharju (2001)), in institutional and individual Israeli investors (Shapira and Venezia (2001)), for investors on the Taiwan Stock Exchange (Barber et al. (2007)), as well as prior studies on UK brokerage investors (Richards et al. (2015)), amongst many others. Consistent with these studies, I also find strong and pervasive evidence of the disposition effect in individual investor transactions.

Whilst many of the studies on the disposition effect focus on the trading of common stocks across various geographies and investor types, there are also those looking at its presence in investors across other asset classes. Woo (2016) studies aggregate trading books of financial institutions in the South Korean bond market and finds little evidence of the disposition effect. Hincapie-Salazar and Agudelo (2019) analyse the Colombian treasury bond market and find a significant effect in fixed income, albeit smaller than in stocks. Studies have also looked at the disposition effect in derivatives markets - Frino, Johnstone, and H. Zheng (2004) and Choe and Eom (2009) both find the disposition effect in futures markets, with the latter finding a larger effect for individual investors than in institutional and foreign investors in Korean index futures markets. Bergsma, Fodor, and Tedford (2019) find evidence consistent with a disposition effect in U.S. equity options. There is also evidence that real estate investors are reluctant to realise their losses, as Genesove and Mayer (2001) find in the Boston condominium market. Consistent with the literature, I find the disposition effect is also present across equities, bonds, and real estate transactions in the dataset used for this study.

Explanations for the disposition effect include Prospect Theory (Kahneman and Tversky (1979)), which describes a kinked utility function, and mental accounting (Thaler (1985)) that suggests a tendency for investors to segregate gambles in separate mental accounts, implying differences in decision making when investors are faced with losses and gains. Barberis and Xiong (2009) study prospect theory as a driver and instead find that a model with realised gains and losses predicts the disposition effect better than simply gains or losses - namely that investors derive utility from the act of realising gains and losses itself Barberis and Xiong (2012). An et al. (2019) find evidence the disposition effect is driven by portfolio level gains and losses, suggesting hedonic mental accounting and preferences as a potential explanation.

Other studies suggest there might be more nuance to the existence and nature of the disposition effect. Eom (2018) finds an opposite disposition effect when the magnitude of gains and losses is large enough - investors tend to hold onto large winners longer than large losers. Da Silva Rosa, To, and Walter (2005) find that U.K. fund managers are more likely to realise losses in their stock holdings, with propensity to sell more related to stock specific factors than to whether or not they are winners. Hirschleifer and Ben-David (2012) find a nonlinear relationship between the level of profit and differences in probability of selling, with no evidence of a jump in probability between gains and losses.

There are also many studies documenting contrasting evidence for trades in mutual funds: Ivkovich and Weisbenner (2009) find that individual investors are more likely to realise losses in mutual fund investments than gains. Calvet, Campbell, and Sodini (2009a) find that investors in Sweden are slightly more likely to exit mutual fund investments when they have performed badly and more likely to exit direct stock holdings when they have performed well. This is consistent with the evidence that fund manager returns and flows into their funds display a positive relationship (Chevalier and Ellison (1997), Ivkovich and Weisbenner (2009)). This is tested directly in Chang, Solomon, and Westerfield (2016), who attribute it to cognitive dissonance. They explore the theory that investors are reluctant to realise losses in stocks as it would be admitting a mistake, which comes at a psychological cost. However, in the case of mutual fund holdings, the fact the investor delegated the investment decisions enables them to fire an underperforming fund manager by attributing blame without the need to admit a mistake, which provides an alternative (and preferable) way to resolve cognitive dissonance. They provide further evidence in an experimental setting, in which they find subjects are more likely to show a difference in disposition effect between stocks and mutual funds when the salience of intermediaries is increased. They also compare actively managed mutual funds to index funds, in which delegation is minimal, and find that the disposition effect is significantly different. I find that compared to single stocks, the disposition effect is less than half the size in index funds and ETFs, and is insignificant in mutual funds. However, results differ when taking into account whether decision making on the account is by the underlying investor or has been delegated to an investment manager.

Though not testing the impact of delegation directly, a related approach taken in Bashall, Willows, and West (2018) was to compare transactions between investment accounts of individuals acting in their own capacity to accounts in which the client had input from a professional advisor. Using data from a South African stockbroker, they compare the disposition effect in investors with those receiving assistance from professional advisors. They find advised clients to have a reduced propensity to realise gains ahead of losses. Similarly, I find the disposition effect to be higher for investors acting in their own capacity than for accounts managed by an intermediary. A key distinction in the delegated accounts that this paper focuses on is the intermediaries have full autonomy in decision making as opposed to an end investor which has input from an advisor but still has the final decision on transactions.

Many studies document how the disposition effect varies across demographic characteristics. Dhar and Zhu (2006) find a smaller disposition effect in individuals with more wealth and professional occupations. Calvet, Campbell, and Sodini (2009b) find greater investor education reduces the disposition effect, whilst

Frina, Lepone, and Wright (2015) find female investors and older investors prone to a larger disposition effect, with nationality also playing a role. Using transactions from Brazilian asset management firm, Oreng, Yoshinaga, and Junior (2021) find that risk averse investors and female investors are more prone to the disposition effect. Ahn (2021) uses feature selection techniques from classical machine learning to identify the most relevant factors associated with the disposition effect - concluding that gender, risk attitude, and investor sophistication are the most significant. However, none consider these demographic differences in the context of the delegation gap. Using feature selection I find that nationality, the percentage of wealth invested, and the time horizon of investors are key demographic variables driving differences in the disposition effect in the cross section of investors. Amongst these, the effectiveness of delegation in reducing the disposition effect is not significant for accounts where investors have shorter time horizons.

Though there is a large body of literature on the disposition effect, there have been fewer studies focusing on its time series variation and relationship with exogenous conditions. Barber et al. (2007) aggregate the transactions of investors on the Taiwan stock exchange and find a significant disposition effect which was robust across years, with the propensity to sell winners relative to losers declining following strong market returns. Using transaction histories of traders on the Taiwan Futures Exchange, Cheng, Lee, and Lin (2013) find a stronger disposition effect following negative returns in equity markets, controlling for differences in demographic characteristics. Bernard, Loos, and Weber (2021) analysed transaction data from a German brokerage and found that the disposition effect is more prevalent during bust periods in equity markets, suggesting both preference and belief based explanations (higher risk aversion and/or more pessimism about future returns in bust periods lead to investors locking in gains). Using zip codes for individual investors, Henriksson (2021) found that the disposition effect increased following local natural disasters, with a larger impact after more severe disasters, suggesting a higher marginal utility from realising gains after a negative external shock to utility. I explore the impact on the disposition effect of prevailing news sentiment at the time of transactions utilising an economic sentiment index constructed via textual analysis of economic news articles Shapiro, Sudhof, and Wilson (2020). I find that sentiment is a significant factor explaining delegation's impact on the disposition effect and that delegated accounts show a smaller (larger) delegation gap during periods of positive (negative) sentiment - the impact of delegation in reducing the disposition effect is larger during pessimistic times.

This study contributes to the literature in a number of ways. First, I add to the body of evidence documenting the disposition effect using a unique dataset of investor transactions from a UK wealth management platform. The transactions include a large number of accounts that invest in both individual securities and funds across asset classes. Second, by grouping transactions based on whether investment decisions were either made by an individual investor or delegated to an investment manager, I provide new insights on the effect of delegation on trading behaviour. Third, I explore how the relationship between trading behaviour and delegation differs across investor demographics and news environments. To my knowledge, this is the first study to explore intermediaries and the disposition effect in this way. The aim of this study is to further explore the relationship between delegation and the disposition effect, which was the focus of Chang, Solomon, and Westerfield (2016), through new analysis of a novel dataset.

I find that the disposition effect is positive and significant across subsets, with the exception of active funds. In contrast, there exists a delegation gap between direct accounts and delegated accounts - the impact of delegation on the disposition effect is negative and significant, driven by transactions in funds and index investments, where delegating to an intermediary reduced the disposition effect by half. The delegation gap is insignificant for single stocks, suggesting nuance in the effect of delegation across security types. Delegation also has a different effect on the way winning and losing securities are treated, with delegated accounts more likely to redeem from losing fund managers but no difference in the way they sell losing stocks. I also find that demographic features play a role in the effect of delegation - identifying the key demographic features related to the disposition effect (Nationality, Percentage of Wealth Invested, and Time Horizon), evidence suggests delegation does not reduce the disposition effect for accounts with shorter time horizons. Finally, I also find that both the disposition effect and the delegation gap have a positive relationship with economic news sentiment - the difference in trading behaviour of delegated accounts is more pronounced in periods of negative sentiment, when the impact of delegation in reducing the disposition effect increases.

The rest of the paper is laid out as follows. Section 12 outlines the data and summary statistics. Section 5 looks at the relationship between delegation and the disposition effect across all transactions and in subsets of various asset types. Section 6 links the disposition effect to investor characteristics and investigates the effect of delegation on the subsets of accounts based on key demographic features. Section 7 analyses the relationship between news sentiment and the delegation gap. Section 14 concludes.

4 Data

One of the main challenges facing researchers in household finance is the lack of centralisation of investor transaction data. Whilst much of the research utilises surveys, there have been a number of studies analysing unique datasets: US brokerage clients (Schlarbaum, Lease, and Lewellen (1978),Barber and Odean (2002)), Finnish investors (DØskeland and Hvide (2011)), German brokerage clients (Dorn and Huberman (2005)), Chinese investors (Feng and Seasholes (2008)), investors on the Taiwan Stock Exchange (Barber et al. (2009)), and the entire population of Swedish households (Calvet, Campbell, and Sodini (2007)), amongst others. This study intends to add to the body of evidence around the investment decisions of households with a unique dataset from a UK based wealth management platform, that includes a heterogeneous set of securities and investors.

To understand the dataset, it is helpful to briefly describe the structure of UK platforms in the context of the wealth management industry. Retail consumers and financial advisers use UK investment platforms to access retail investment products from a number of different providers. Through platforms, they can either execute their own investment decisions or delegate to intermediaries who then implement their portfolios on behalf of the underlying clients. As of 2019, UK platforms held over £500 billion in assets under administration with two thirds of assets held on platforms invested in funds [1].

The firm who provided this data has been in existence since 1998, with transactions across $\sim 16,000$ unique client accounts included since its inception to the end of 2020, with an average account size of £590,000. The dataset spans daily transactions across a number of security types on the platform, including stocks, bonds, direct real estate transactions, active mutual funds, hedge funds, index funds, and exchange traded funds. A number of investor characteristics have been documented, including age, gender, nationality, risk tolerance. Given the richness of the data, there are a wide number of potential hypotheses for a researcher to test. A key variable available is the type of mandate of each account, enabling a separation into accounts which are purely for execution purposes, with the final investment decisions made by the account holder (Direct Accounts), and accounts where investment decisions are delegated to an investment manager (Delegated Accounts). There are a variety of intermediary firms within the Delegated Accounts using the platform, spanning different investment approaches and utilising different security types across asset classes. The initial focus of this study is on differences in trading behaviour between these two account types and the extent to which the differences vary across asset types, investor groups, and through time.

The characteristics of the firm are broadly in line with the UK investment platform industry and it is representative in a number of ways. Total assets under administration of £4 billion were evenly split between funds and single stocks (39% in open ended funds, 37% in individual equities, and also 10% in exchange traded funds). Stock holdings (which include depositary receipts) are diverse across sector, geography, and market capitalisation. Within fixed income, the dataset also contains a number of corporate bond transactions, including some of lower credit quality and also hybrid instruments such as contingent convertible bonds, which may provide additional insights to previous studies of the disposition effect in fixed income which are predominantly focused on government securities. Whilst transactions in exchange-traded funds (ETFs) are diverse across listing geographies, a large proportion of the fund transactions are in European domiciled funds, given the regulatory requirements of the platform at the time and that the majority of account holders are based in Europe. Though the underlying strategies of the funds are diverse across asset class, geography, and style. Virtually all ETFs in the data track an index, thus I interpret them as involving less delegation from an end investor compared to actively managed funds, similar to Chang, Solomon, and Westerfield (2016). The data includes transactions in

[1] FCA 2019

hedge funds, which are included in the aggregate analyses and in the subset of active funds with traditional mutual funds.

I begin by excluding cash related transactions such as account transfers, foreign exchange, as well as money market fund trades - a total of 970,000 transactions are excluded. All non investor related accounts (such as non trust corporate entities) are also excluded. Index funds are then identified within the fund transactions and combined with ETF transactions into a separate group (Index Funds and ETFs). To control for the effect of rebalancing, only transactions which involved full sales of a security were used (Odean (1998)) - since the expectation is that an investor who is transacting for the purposes of rebalancing their portfolio will sell proportions of their holdings in winning securities without fully exiting the position. By excluding transactions with partial sales, it reduces the likelihood the trades, which involve selling their entire holding in a security, were made with the intention of rebalancing at the portfolio level as opposed to an active trading decision. The excluded partial sales made up 17% of the total transactions (184,000) in the original dataset.

Following the exclusions, Table 1 shows summary statistics for the transactions included in the analysis (total of 899,729 observations) across security types, for the two types of accounts. Delegated accounts were larger (averaging £820,000 vs. £480,000 in Direct accounts) and held more securities per account on average, though there is wide dispersion across investors and asset classes. Both Direct and Delegated accounts tended to hold a broader set of individual securities than they did active funds or index funds, with exception of direct real estate deals in which there were fewer observations. Delegated accounts on average have a larger number of transactions though this is in part be due to Delegated accounts having more diversified portfolios and a slightly longer average lifespan than Direct accounts.

5 Delegation and the Disposition Effect

The original method used by Odean (1998) to measure the disposition effect, which was followed by many subsequent studies, was to count the proportion of gains which an investor realised compared to the proportion of losses realised (effectively comparing the difference in turnover between winning and losing positions).

Grinblatt and Keloharju (2001) instead utilised a logit regression based method to measure the disposition effect, which has since been used in a number of studies (Kaustia (2010), Grinblatt, Keloharju, and Linnainmaa (2012), Birru (2015), Chang, Solomon, and Westerfield (2016)). The advantage of this over the method in Odean (1998) is the ease of adding additional controls and also that the standard errors can be clustered to avoid the need to assume every sale choice is independent. It also enables us to disaggregate trading behaviours which differ in the gain and loss domains - not only can we can analyse the disposition effect (greater turnover in wining securities versus losing securities) but also how delegation is related to trading behaviour in winning and losing securities separately.

For this section, I follow the regression approach using the specification in Chang, Solomon, and Westerfield (2016) [2]. However, where they focus on end of month holdings because they were using monthly holdings files, I perform the analysis across individual transactions on any day that a transaction took place. The initial test in Equation 1 is to confirm the presence of the disposition effect across all transactions and in asset type subsets - For any transaction which was a sale, the dummy variable $Sale$ is equal to 1. And for transactions where the executed price was higher than the average historic purchase price of the security in that account, the dummy variable $Gain$ is equal to 1. α is the probability of selling a position that is a loss and β measures the disposition effect (the increase in probability of selling a winning position over a losing position)

$$Sale_{i,j,t} = \alpha + \beta Gain_{i,j,t} + \epsilon \tag{1}$$

Next, the model in Equation 2 is used to measure the impact of delegation on the disposition effect. Similarly, for any transaction which was a sale, the dummy variable $Sale$ is equal to 1. The dummy variable $Gain$ is equal to 1 for transactions where the executed price was higher than the average historic purchase price of the security in that account. α is the probability of selling a position that is a loss and β measures the disposition effect (the increase in probability of selling a winning position over a losing position) - a negative coefficient indicates a reverse-disposition effect. In addition, Del is a dummy variable equal to 1 for all transactions made in accounts where investment decisions were fully delegated to an investment manager. γ is the difference in the disposition effect between these accounts and those in accounts where the account holder made the decisions and $\beta + \gamma$ measures the overall disposition effect for the accounts with delegation. δ is the differential propensity for delegated accounts to sell securities at a loss compared to other securities at a loss.

$$Sale_{i,j,t} = \alpha + \beta Gain_{i,j,t} + \gamma Gain_{i,j,t} \times Del_j + \delta Del_j + \epsilon \tag{2}$$

Analysis is run on all applicable transactions, then on subsets of the applicable transactions which involved individual stocks, actively managed funds, and index funds and ETFs.

As an additional test to control for financial sophistication, I look at transactions made by a subset of personal accounts owned by financial professionals employed by the firm - these are direct accounts in which investment decisions are made by the owner, who separately also makes decisions on behalf of delegated accounts in a professional capacity. However, a controlled comparison of each professional's personal account transactions to those in the delegated accounts that same professional was managing is not possible since there are only records of each account's current account manager (historical data on changes in manager of each account was not available for this analysis). Thus, these personal

[2] The PGR-PLR method from Odean (1998) was also run as a robustness check for the broad disposition effect and produced similar results

account transactions are aggregated with all delegated account transactions and the model in Equation 2 is used to measure the delegation gap in the subset.

Consistent with previous studies, standard errors are two-way clustered by date and by account, to adjust for serial correlation and cross-sectional correlation in investor trading decisions - this has a large impact on the inference across all results. All statistical significance reported going forward are using these robust standard errors.

5.1 Delegation and Disposition - Results

Results of the initial test for the disposition effect in Equation 1 are shown in Table 3. The positive and significant coefficient on the *Gain* variable confirm the presence of the disposition effect across all transactions in the dataset. When run on subsets by asset type, the disposition effect is also positive and significant in single stocks - investors were more likely to sell winning stocks than losing stocks. There is also a positive and significant disposition effect in index funds and ETFs, though the effect is half the size as in single stocks. In active funds, the disposition effect is smaller and insignificant. These results are consistent with Ivkovich and Weisbenner (2009) and Calvet, Campbell, and Sodini (2009a) who don't find evidence of the disposition effect in investors of mutual funds. The presence of the disposition effect in index funds and ETFs also supports the results in Chang, Solomon, and Westerfield (2016), who find a significant difference in the disposition effect between mutual funds and index funds (which they attribute to the increased role of delegation when allocating to an active fund).

Table 4 shows the regression results using Equation 2 across the various subsets, with significant coefficients in **bold**. The disposition effect is positive and significant across security types on a standalone basis, as can be seen from the large, positive, and significant coefficients of $Gain + Gain \times Del$. However, the $Gain \times Del$ coefficient is negative and significant across the full sample, indicating that transactions in delegated accounts show a reduced disposition effect. Trades made in accounts where investment decisions have been delegated to a professional investment manager showed a lower propensity to sell winning securities over losing securities. These results are directionally consistent with Bashall, Willows, and West (2018), though their study focused on investors acting under guidance as opposed to accounts which fully delegated investment decision making.

Separating transactions into subsets based on asset type indicates that there are nuances in the impact of delegation depending on the types of securities being traded. Firstly, the effect of delegation on the propensity to sell winners is insignificant for trades in single stocks ($Gain \times Del$ is not statistically significant). Shefrin and Statman (1985) documented a disposition effect in trades of equity mutual fund managers, who in a related fashion are investing in single stocks on behalf of their unitholders. This result links the research on the disposition effect in individual investors and in mutual fund managers by providing evidence that intermediaries do not show a significantly different disposition effect compared to investors making their own trading decisions in individual stocks. One distinction is that the delegated accounts in this paper may have a personal relationship with the underlying client, compared with equity fund managers who are more often investing for a multitude of underlying fund holders at once. The extent to which the salience of the relationship between the delegator and the delegatee can affect trading behaviour is a potential area of further research.

In contrast, $Gain \times Del$ is significant and negative for transactions on funds and index investments - in both cases, delegation reduces the overall disposition effect by more than half. The propensity to sell winning funds more than losing funds is significantly reduced when investment decisions are delegated. Comparing Table 4 to the univariate results in Table 3 suggests that the smaller disposition effect in active funds and index funds is driven by the delegation gap (the difference in trading behaviour of delegated accounts compared to direct accounts).

When looking at losing trades, in the case of active funds, the positive and significant coefficient on Del indicates that delegated accounts are also more likely to redeem from losing fund managers, which is consistent with the literature on performance chasing in mutual fund investors (Chevalier and Ellison (1997), Ivkovich and Weisbenner (2009)). This was not the case with index funds and ETFs, where there was no statistically significant difference between turnover of losing index investments in delegated accounts.

Though Chang, Solomon, and Westerfield (2016) attributes the difference in disposition effect between active funds and index funds to larger role of delegation in allocating to the former (cognitive dissonance plays less of a role when the responsibility for a loss is due to the decisions of a fund manager and not the investor themself), the negative sign of the delegation gaps in Table 4 suggests delegation reduces the disposition effect for both types of funds. The distinction between the insignificant $Gain \times Del$ coefficient for single stocks and the significant and negative delegation gap for both active funds % index funds raises an interesting question yet to be addressed in the literature - does the inherent diversification of funds relative to single stocks play a role in reducing the disposition effect? Whilst beyond the scope of this study, it may be a fruitful area for further research.

One interpretation of the delegation gap is that delegated accounts are managed by more sophisticated investors, who are less prone to the disposition effect (Dhar and Zhu (2006), Calvet, Campbell, and Sodini (2009b), Bashall, Willows, and West (2018) Ahn (2021)). Client type labels, as defined by the UK Financial Conduct Authority, allow us to assess the effect of investor sophistication on the delegation gap. Table 5 show the results of running the model in Equation 2 when partitioning the data by Retail clients and Non Retail accounts. Though the size of the delegation gap for professional investors is half that of retail investors, the delegation gap is significant in both subsets. Additionally, when splitting by asset type, the delegation gap shows up in index investments for professionals, which is not the case for retail investors (albeit this may be due to the small sample size of retail ETF transactions).

Similarly, the subset of trades made by investment professionals in their own personal accounts also showed a delegation gap. Since the delegation gap is exhibited by both retail and professional investors, as well as investment managers in their personal accounts, it suggests the higher disposition effect in accounts without any delegated decision making is not driven by differences in financial sophistication of account owners and investment managers they have delegated to, in contrast to the conclusion of similar analysis in Bashall, Willows, and West (2018). Some caveats to this are that professional accounts display a reduced delegation gap and that the subset of personal account transactions is small (14,271) and does not cover all the professional investors who are making investment decisions for the delegated accounts. Also, the account type definitions (e.g. retail, professional) are self selected by account owners upon account opening - other measures of financial sophistication would enable a better exploration of the relationship between financial sophistication and the delegation gap.

This section provides evidence that delegation can reduce the disposition effect. And if the goal for practitioners is to reduce the disposition effect, which has been shown to be detrimental to investor outcomes, then delegation can be a useful tool except when the investment mandate is implemented only using single stocks. Next, I explore the delegation gap in subsets of accounts by demographic features of investors.

6 Demographic features

The dataset includes investors with a variety of demographic characteristics. Demographic data for a portion of accounts was captured upon account opening and includes discrete categories a number of characteristics, including: employment status, gender, marital status, nationality, financial dependents, time horizon (discrete bands), % of total wealth invested (discrete bands), and risk rating (a subjective assessment of risk capacity and risk attitude given by the firm). Coverage varies by field and the choice of demographic variables to include in analysis. Table 2 shows a sample of available categories per characteristic. The data shows noticeable clusters of accounts, with the majority of account holders being married, with no financial dependents, and nationality as either South African or British (with a long tail of other nationalities). 80 % of accounts are male and almost half had less than 10% of their total wealth in the account at inception (for the accounts where these metrics had been disclosed). Accounts have a median investment time horizon at 7 years and are evenly split between lower and higher risk rating categories.

For cross sectional analysis of the disposition effect across accounts, I calculate the account level disposition effect using the method in Odean (1998) and Barber et al. (2007), as the proportion of gains which an investor realised compared to the proportion of losses realised, over the life of the account (PGR-PLR).

$$PGR = \frac{RealisedGains}{RealisedGains + PaperGains} \tag{3}$$

$$PLR = \frac{RealisedLosses}{RealisedLosses + PaperLosses} \tag{4}$$

A subset of 2,306 have PGR-PLR and all demographic variables available. Demographic data are mapped into dummy variables as in Table 2: employed, male, married, British nationality, above 10% wealth invested, and above 7 year time horizon have a dummy = 1. Regressing PGR-PLR from this subset of accounts on their demographic dummy variables as per Equation 5, only time horizon is statistically significant, with a positive sign. This is inconsistent with the literature on the factors associated with the disposition effect, who have found relationships between characteristics such as gender and nationality with the disposition effect (Dhar and Zhu (2006), Frina, Lepone, and Wright (2015),Oreng, Yoshinaga, and Junior (2021)).

$$PGRPLR_i = \alpha + \beta Demo_k + \epsilon \tag{5}$$

6.1 Feature Selection - LASSO

Ahn (2021) combines survey data on South Korean individual investors with trading records to explore the use of feature selections methods including the least absolute shrinkage and selection operator (LASSO, Tibshirani (1996)) to determine the most relevant factors affecting the disposition effect. A standard OLS minimises the sum of squared residuals, producing the parameter estimates for β from Equation 5 in the first row of Table 7. Instead, LASSO will minimise the sum of squared residuals whilst minimising the absolute value of the parameters based on a chosen penalty λ, as per Equation 6.

$$\sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2 + \lambda_1 \sum_{j=1}^{p} |\beta_j| \tag{6}$$

The parameter estimates are reduced towards zero, leaving the most salient features. As in Ahn (2021), I use LASSO to select a key subset of relevant features that drive the disposition effect. Table 7 compares the results of the LASSO to the standard OLS regression in Equation 5, the ℓ_1 penalty shrinks the coefficients of Employment Status, Gender, and Marital Status to zero (using $\lambda = 0.005$). This is in contrast to Ahn (2021) who did find gender to be a key driver of the disposition effect. The difference in results may be due to the different variables used in that study, which included those derived from survey responses such as risk profile, loss aversion, knowledge, and experience.

Removing the variables that were shrunk to zero post regularisation, I rerun Equation 5 on the remaining three which shows statistically significant and negative coefficients for Percentage of Wealth Invested and Time Horizon, whilst Nationality is positive but insignificant. These results suggest a lower disposition effect for accounts which have more than 10 % of total wealth invested, and a time horizon longer than 7 years.

6.2 Delegation and Demographic features - Results

The features selected from the LASSO were based on the cross section of disposition effects at the account level as labels. Using the features selected, I look at the extent to which the effect of delegation on the disposition effect differs between investor groups. Applying Equation 2 to transaction subsets (separated by the key features) provide insight into the impact of delegation on trading behaviour for the following groups: British investors vs other nationalities, investors with low vs high percentage of wealth invested, and investors with short vs long term time horizons.

Table 8 shows the regression results for the model in Equation 2 on subsets of accounts by each demographic feature selected by the LASSO - I look for the significance of the delegation gap first in accounts where the relevant dummy variable is equal to zero, and then in accounts where it is equal to one.

The disposition effect for direct accounts ($Gain$) is positive and significant for all subsets. The delegation gap (represented by the coefficient on $Gain \times Del$) is also significant in all but one of the subsets - the extent to which investors realise winners more than losers is reduced for accounts where investment decisions are delegated. However, the delegation gap is insignificant for accounts with shorter time horizons - delegation does not reduce the disposition effect significantly for shorter term investors. One explanation could be the tighter compliance restrictions for shorter term mandates (time horizon is a direct input into risking profiling tools for UK wealth management firms), which may lead to more homogeneous trading behaviour between direct and delegated accounts in this subset. Another explanation is that shorter term investment mandates induce behaviours in intermediaries that are not conducive to reducing the disposition effect. However, the positive and significant coefficient on the Del variable for shorter term accounts indicates that delegation led to a higher propensity to sell losing securities, which suggests that even in the subset of shorter term accounts there were differences in trading behaviour.

Given the disposition effect has been documented as being costly for investors (Odean (1998)), these results suggest that delegation of investment decisions (which is intended to improve investor outcomes) should be focused on investors with longer time horizons.

A caveat to this section of analysis is that the subset of accounts with both a measurable disposition effect and all the demographic features available might not be representative - a large proportion of the accounts in the overall dataset are missing at least one feature and thus not included ($\sim$ 80% of accounts are ineligible). However, the included accounts in this section cover a third of total transactions, in which the results are mostly consistent with the results across all transactions in Table 4 (a significant disposition effect and negative sign for the delegation gap). In the following section, I look at time variation in the delegation gap and its relationship with economic news sentiment.

7 Sentiment and the Delegation Gap

7.1 The Disposition Effect in Time Series

Though the dataset began in 1998, the number of observations were only meaningful from 2003 [3]. Grouping transactions by calendar year highlights a number of insights. Qualitatively, the magnitude of the disposition effect has decreased over time. The disposition effect on a univariate basis has been consistently positive and significant in each year (with the exception of 2020) - this is in line with the time series results in Barber et al. (2007). When assessing the effect of delegation, there is time variation in the delegation gap across calendar year - the sign and significance of the delegation gap varies dramatically (Table 9). The delegation gap is noticeably large and positive in the years immediately following 2008 - delegation doubled the size of the disposition effect in 2009 and increased by over a third in 2010 and 2011. Delegated accounts were also significantly more likely to sell losing securities (the *Del* coefficient is positive and significant in all years between 2006 and 2012). In contrast, during 2020, delegation reduced the disposition effect to almost zero - whilst direct accounts exhibited similar trading behaviour to other years (*Gain* was large and significant), delegated accounts saw a dramatic reduction in the propensity to sell winners compared to losers in that calendar year ($Gain \times Del$ was negative, significant, and a similar magnitude to the *Gain* coefficient). There was no significant difference between how delegated accounts traded losing positions in 2020.

The source of time variation is not obvious - qualitatively, the calendar year results do not appear to be related to equity market volatility - for example, the delegation gap has different signs in 2008 and 2020. This is inconsistent with both Cheng, Lee, and Lin (2013) and Bernard, Loos, and Weber (2021), who find a higher disposition effect during and following periods of negative equity returns. The negative and significant delegation gap in recent years also coincide with an increase in the number of transactions in the dataset over time, as the platform grew in size. Similarly, the proportion of trades in funds and index investments increased over time, whilst the proportion of total transactions in single stocks declined.

Table 10 shows the model in Equation 2 run on 12 subsets, each constructed as the full transaction dataset ex transactions in each calendar month. There is little evidence of the tax motivated trading as documented in the literature (Buhlmann et al. (2020), Bazley, Moore, and Vosse (2021)), where the disposition effect is reduced when investors are involved in tax-loss selling and holding winning securities to reduce capital gains tax (Odean (1998)). The delegation gap is robust to calendar month effects and the $Gain \times Del$ coefficient is significant in each subset - there is no evidence that the gap is driven by more tax awareness from delegated accounts since the effect is not concentrated in any one calendar month.

To more directly test for time varying effects, I use look at a more granular analysis to utilise the daily frequency of the transaction data.

[3] prior years saw less than 10,000 total transactions per calendar year

7.2 Sentiment Index

Shapiro, Sudhof, and Wilson (2020) introduce the Daily News Sentiment Index, now published by the Federal Reserve of San Francisco, which uses text sentiment analysis tools to develop a time-series measure of economic sentiment based on economic and financial newspaper articles. Using a combination of lexicons, they rate the sentiment of articles from 16 major newspapers between 1985 and 2015, where the topic was economic or the economy. By aggregating individual article scores they create U.S. national daily and monthly time series of news sentiment.

They find survey based consumer sentiment correlated to their index time series and also that economic activity is responsive to their measure of news sentiment - their index falls during recessions but also during months of key historical events, which may not be captured by broad economic indicators (Figure 1). I use this novel and publicly available dataset to test the impact of news sentiment on the relationship between delegation and the disposition effect.

The $SENT$ variable is the news sentiment index normalised to have mean 0 and standard deviation of 1, averaged over the previous 5 business days. I interact $SENT$ with $Gain$,$Gain \times Del$, and Del on days the transactions were made. ω in Equation 7 measures the relationship between the disposition effect in direct accounts and news sentiment during the previous week. ψ measures the relationship between the delegation gap (differences in the disposition effect between delegated and direct accounts) and news sentiment during the previous week, whilst η measures the relationship between news sentiment during the previous week and the difference in propensity to sell losing securities between delegated and direct accounts. The choice of 5 business days is to capture the broad prevailing sentiment level during which a trade decision is made.

$$\begin{aligned} Sales_{i,j,t} = \alpha + \beta Gain_{i,j,t} + \gamma(Gain_{i,j,t} \times Del_j) + \delta Del_j + \boldsymbol{\omega}(Gain_{i,j,t} \\ \times SENT_t) + \boldsymbol{\psi}(Gain_{i,j,t} \times Del_j \times SENT_t) + \boldsymbol{\eta}(Del_j \times SENT_t) + \epsilon \end{aligned} \tag{7}$$

7.3 Delegation and Sentiment - Results

Table 11 shows the regression results for Equation 7 looking at the impact of news sentiment on the relationship between delegation and the disposition effect.

Assessing the results across all transactions, the coefficient on $Gain \times SENT$ is positive and significant. This indicates that the disposition effect in direct accounts is positively related to prevailing news sentiment at the time of transaction - positive (negative) sentiment increases (decreases) the disposition effect. This relationship is stronger in single stocks than active funds but is insignificant in index investments. These results point in a different direction to other studies on the sensitivity of the disposition effect to market conditions. Cheng, Lee, and Lin (2013) find evidence amongst futures traders in Taiwan that the disposition effect is higher during bear markets. Bernard, Loos, and Weber (2021) find evidence amongst German equity investors that the disposition effect increases (decreases) during busts (booms). However, the difference may be due to the time horizons - the two studies define bear markets / bust periods as a 3 month / 24 month period of negative historic equity market returns respectively, whereas in this analysis prevailing news sentiment is being measured over a 5 day window. The sentiment index constructed by Shapiro, Sudhof, and Wilson (2020) is a more holistic measure of market sentiment than equity market returns since these transactions span other asset classes - equity market movements are not relevant for all investors in this data set.

The coefficients on $Gain \times Del \times SENT$ is positive across all transactions, suggesting a positive relationship between the delegation gap and news sentiment - delegated accounts are more (less) likely to sell winning securities over losing securities following weeks with positive (negative) news sentiment than direct accounts. The disposition effect in direct accounts is positively related to news sentiment and the impact of delegation is to enhance this sensitivity. Whilst delegation reduces the disposition effect in an unconditional setting, the delegation gap is less negative during positive news sentiment and more negative in periods of negative sentiment. The benefits of delegation in reducing the disposition effect is enhanced during more pessimistic environments.

In contrast, none of the $Del \times SENT$ coefficients are significant, suggesting an asymmetry in the differences between trading behaviour of direct and delegated accounts during periods of positive and negative news sentiment - though delegated accounts were more likely to sell winners during periods of positive sentiment, there was no difference in how they traded losers relative to direct accounts.

One interpretation of these results is that investors managing their own money are more sensitive to loss aversion during periods of negative short term news flows compared to an intermediary, since the act of delegation reduces the salience of losses and thus the trading behaviour of delegated accounts is less affected. Another interpretation might be that professional investors are less impacted during these periods as they are more financially sophisticated, which would be consistent with Dhar and Zhu (2006), Calvet, Campbell, and Sodini (2009b), and Ahn (2021). In this interpretation, these results suggest a more nuanced relationship between investor sophistication and the disposition effect - the extent to which sophisticated investors are less prone to the disposition effect differs depending on the prevailing sentiment.

Another explanation might be that all investors are more prone to the disposition effect during periods of positive sentiment but professional investors are more sensitive to short term news, which is reflected in their trading behaviour.

Therefore, in looking to reduce the disposition effect, intermediation is more effective during periods of negative sentiment than during more optimistic times. A caveat is that the index measuring news sentiment is US centric and won't capture the effect of local shocks to sentiment for the investors. Further research could include local measures of sentiment matched with investor location based on address data to assess this effect.

8 Conclusion

Though the literature on the disposition effect is vast, the majority of studies have focused on trading behaviour of investors in various data sets without taking into account whether the trading decisions are made by the account holder or delegated to a third party.

This paper highlights the following: 1) The disposition effect is prevalent across account types, asset classes, and asset types (with the exception of actively managed funds) 2) A Delegation Gap exists between the disposition effect in accounts of investors investing their own money compared to accounts in which they are making decisions on behalf of someone else. This is also true for professional investors investing in their personal accounts. 3) The delegation gap varies depending on the type of asset being traded, whether single stocks, active funds, or index investments. 4) The delegation gap also differs by investor cohorts - delegation did not reduce the disposition effect in accounts of shorter term investors as it did in other groups. 5) Both the disposition effect itself and the delegation gap are affected by short term news sentiment - delegation is less effective during periods of positive sentiment but further reduces the disposition effect during negative new sentiment.

The results of this study can aid practitioner firms in assessing the contexts in which delegating investment decision making can improve investor outcomes - though investors can reduce the impact of the disposition effect by delegating investment decisions, investment mandates which are shorter term and/or more sensitive to recent news sentiment may not benefit from delegation.

This study also highlights further areas of research to extend the literature. The extent to which the delegation gap varies with characteristics of delegatees, demographic differences between principal and agent, whether it explains cross sectional differences in the disposition effect in other data sets, its sensitivity to other measures of sentiment, and its impact on investment performance and investor outcomes more broadly, are all potential extensions of this research.

Table 1: Summary statistics for transactions in Delegated accounts, where decisions making is delegated to an investment manager, and Direct accounts, where the account holder makes the investment decisions.

	Delegated			**Direct**		
	Transactions	**Unique Accounts**	**Unique Securities**	**Transactions**	**Unique Accounts**	**Unique Securities**
All Transactions	613,094	4,589	9,946	286,635	10,746	16,293
Single Stocks	256,181	2,941	4,455	183,618	5,361	8,294
Active Funds	252,132	3,451	2,457	40,427	5,879	3,150
Index Funds & ETFs	45,275	2,780	596	15,117	2,066	959

Table 2: Sample categories of demographic characteristics available for investors in the dataset, with count of unique accounts in each.

Sample categories (mapping)	**# Accounts**
Employment Status (Employed = 1)	
Employed	6,485
Self-Employed	742
Unemployed	547
Gender (Male = 1)	
Female	1,157
Male	5,235
Marital Status (Married = 1)	
Divorced / Separated	201
Married	10,332
Single	651
Widowed	155
Nationality (UK = 1)	
Australia	93
South Africa	8,156
Switzerland	85
United Kingdom	5,106
Risk Rating (Growth & Adventurous = 1)	
Cautious	1,017
Moderate	2,102
Growth	2,219
Adventurous	1,535
Time Horizon (> minimum 7 years = 1)	
Less Than 3 Years	526
5-7 Years	508
7-10 Years	817
10 Years +	1,379
% Wealth Invested (>10% = 1)	
0-10	2,228
10-20	414
20-30	251
30-40	134

Figure 1: News Sentiment Index from Shapiro, Sudhof, and Wilson (2020) 1980 to 2020, normalised to a mean 0 and standard deviation of 1

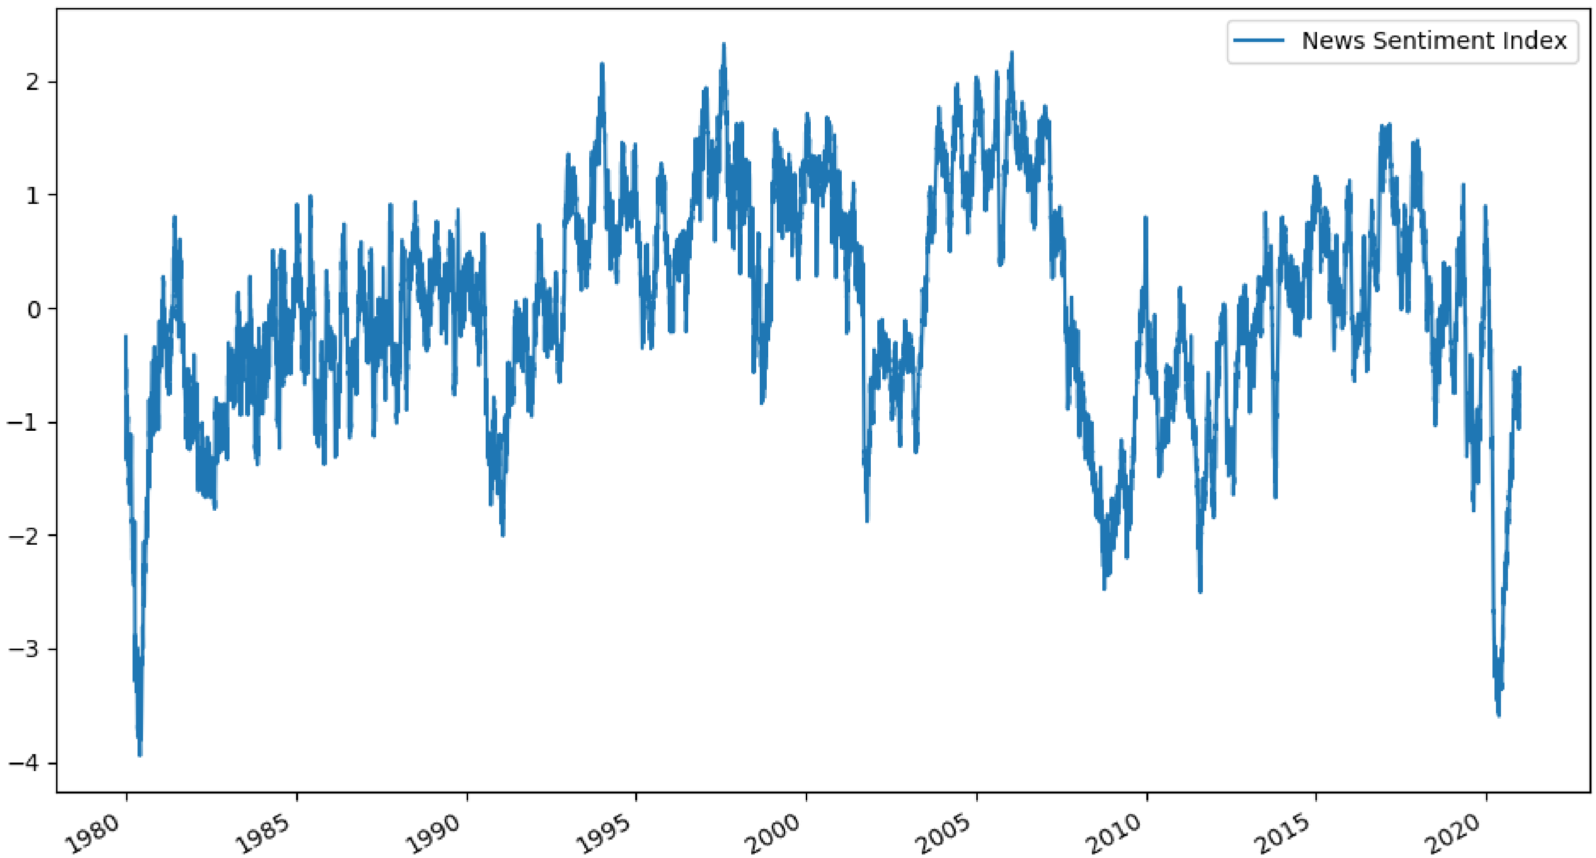

Table 3: Regressions run on all transactions. Dependent variable is a dummy equal to 1 when a sale occurred, Gain is a dummy equal to when the transaction price is above the historic average purchase price. T-statistics in brackets are based on standard errors two-way clustered by account and date.

	Transactions	**Gain**	**Constant**
All Transactions	899,729	**0.195****	**0.187****
		(7.083)	(20.204)
Single Stocks	439,799	**0.266****	**0.186****
		(16.731)	(29.249)
Active Funds	292,559	0.095	**0.180****
		(1.831)	(7.734)
Index Funds & ETFs	60,392	**0.127****	**0.163****
		(3.32)	(8.59)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 4: Regressions run on all transactions in direct accounts (investor makes the investment decisions) and delegated accounts (decisions are delegated to an investment manager). Dependent variable is a dummy equal to 1 when a sale occurred, Gain is a dummy equal to when the transaction price is above the historic average purchase price, Del is a dummy equal to 1 for transactions in delegated accounts. T-statistics in brackets are based on standard errors two-way clustered by account and date.

	Transactions	**Delegated**	*Gain*	*Gain* × *Del*	*Del*	*Constant*
All Transactions	899,729	613,094	0.268***	-0.098***	0.020*	0.174***
			(22.675)	(-2.971)	(1.66)	(31.139)
Single Stocks	439,799	256,181	0.258***	0.011	0.014	0.178***
			(19.342)	(0.513)	(1.564)	(28.999)
Active Funds	292,559	252,132	0.219***	-0.139**	0.054**	0.136***
			(8.886)	(-2.406)	(2.016)	(14.568)
Index Funds & ETFs	60,392	45,275	0.226***	-0.116**	-0.012	0.171***
			(7.709)	(-2.133)	(-0.505)	(8.609)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 5: Regressions run on transactions made by retail and non retail investors in direct accounts (investor makes the investment decisions) and delegated accounts (decisions are delegated to an investment manager). Dependent variable is a dummy equal to 1 when a sale occurred, Gain is a dummy equal to when the transaction price is above the historic average purchase price, Del is a dummy equal to 1 for transactions in delegated accounts. T-statistics in brackets are based on standard errors two-way clustered by account and date.

Retail	**Transactions**	**Delegated**	**Gain**	**Gain x Del**	**Del**	**Constant**
All Transactions	517,009	395,340	0.264***	-0.115***	0.007	0.179***
			(22.059)	(-3.312)	(0.534)	(28.624)
Single Stocks	220,783	146,454	0.251***	0.016	0.018*	0.182***
			(18.319)	(0.687)	(1.826)	(26.464)
Active Funds	173,640	155,063	0.205***	-0.112**	0.042	0.161***
			(7.944)	(-2.003)	(1.781)	(13.143)
Index Funds & ETFs	4,032	1,194	0.198***	0.062	0.108***	0.136***
			(4.716)	(0.779)	(2.672)	(8.280)

Non Retail	**Transactions**	**Delegated**	**Gain**	**Gain x Del**	**Del**	**Constant**
All Transactions	382,720	217,754	0.271***	-0.062**	0.037***	0.170***
			(21.965)	(-2.101)	(3.807)	(31.434)
Single Stocks	219,016	109,727	0.263***	0.009	0.006	0.175***
			(19.268)	(0.421)	(0.720)	(29.850)
Active Funds	118,919	97,069	0.224***	-0.163**	0.054	0.117***
			(7.054)	(-2.566)	(1.582)	(13.104)
Index Funds & ETFs	56,360	44,081	0.237***	-0.130**	-0.022	0.179***
			(7.838)	(-2.419)	(-0.878)	(8.280)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 6: Regressions run on transactions made by investment professionals in their personal accounts and delegated accounts (decisions are delegated to an investment manager). Dependent variable is a dummy equal to 1 when a sale occurred, Gain is a dummy equal to when the transaction price is above the historic average purchase price, Del is a dummy equal to 1 for transactions in delegated accounts. T-statistics in brackets are based on standard errors two-way clustered by account and date.

	Transactions	**Delegated**	*Gain*	*Gain* × *Del*	*Del*		*Constant*
Direct PA & Delegated	639,475	625,204	0.280***	-0.110***	-0.016		0.210***
			(9.477)	(-2.638)	(-1.011)		(18.934)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 7: Compare coefficients of regression of difference between proportion of gains realised and proportion of losses realised across investor accounts on dummy variables based on investors' Employment Status, Gender, Marital Status, Nationality, Percentage of Total Wealth Invested, and Time Horizon to the coefficients when using LASSO shrinkage

	Constant	Employment Status	Gender	Marital Status	Nationality 1	Percentage of Total Wealth	Time Horizon of Investment
Model 1	0.365	-0.006	-0.026	-0.009	0.054	-0.039	-0.045
LASSO	0.000	0.000	0.000	0.000	0.019	-0.018	-0.023

Table 8: Regressions run on subsets of accounts based on differences in Nationality, Percentage of Total Wealth Invested, and Time Horizon. Coefficients using the model in Equation 2, with Sales on Gains, Delegated accounts, and $Gain \times Del$ which shows the difference in disposition effect between direct and delegated accounts in that subset. T-statistics in brackets are based on standard errors two-way clustered by account and date.

	Transactions	*Gain*	$Gain \times Del$	*Del*		*Constant*
TimeH=0	61,961	0.259***	-0.078	0.043***		0.158***
		(6.668)	(-1.462)	(2.967)		(15.756)
TimeH=1	227,108	0.289***	-0.132***	0.002		0.177***
		(17.754)	(-3.800)	(0.181)		(26.709)
Wealth=0	68,877	0.270***	-0.092**	0.037**		0.162***
		(8.599)	(-2.001)	(2.532)		(16.898)
Wealth=1	220,192	0.286***	-0.128***	0.003		0.176***
		(16.099)	(-3.544)	(0.219)		(25.564)
Nat=0	150,524	0.267***	-0.103***	0.019		0.167***
		(11.752)	(-2.640)	(1.434)		(20.228)
Nat=1	138,545	0.298***	-0.139***	0.003		0.179***
		(15.416)	(-3.700)	(0.190)		(25.328)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 9: Regressions run on transactions grouped by calendar year in direct accounts (investor makes the investment decisions) and delegated accounts (decisions are delegated to an investment manager). Dependent variable is a dummy equal to 1 when a sale occurred, Gain is a dummy equal to when the transaction price is above the historic average purchase price, Del is a dummy equal to 1 for transactions in delegated accounts. T-statistics in brackets are based on standard errors two-way clustered by account and date.

	Transactions	**Delegated**	**Gain**	**Gain x Del**	**Del**	**Constant**
2001	9,681	5,539	**0.465****	0.043	**0.077****	**0.148***
2002	9,794	5,697	**0.260****	**0.150****	0.055	**0.257****
2003	19,774	13,949	**0.417****	0.049	**0.073****	**0.226****
2004	21,172	14,560	**0.425****	0.017	0.074	**0.229****
2005	17,627	10,771	**0.434****	0.044	-0.008	**0.236****
2006	17,146	9,622	**0.511****	-0.020	**0.085****	**0.174****
2007	14,451	6,156	**0.284****	0.118	**0.114****	**0.163****
2008	11,583	5,206	**0.308****	0.041	**0.074****	**0.182****
2009	15,449	6,195	**0.230****	**0.212****	**0.045****	**0.139****
2010	21,233	7,660	**0.316****	**0.136****	**0.056***	**0.122****
2011	22,738	8,185	**0.246****	**0.158****	**0.078****	**0.151****
2012	24,666	9,241	**0.260****	0.061	**0.096****	**0.140****
2013	31,134	14,956	**0.232****	-0.005	0.028	**0.129****
2014	39,797	19,724	**0.207****	0.055	0.024	**0.121****
2015	56,915	34,943	**0.247****	0.044	-0.018	**0.159****
2016	64,225	44,095	**0.218****	0.028	0.026	**0.186****
2017	76,661	50,633	**0.257****	0.025	-0.008	**0.162****
2018	93,009	68,557	**0.293****	-0.062	-0.017	**0.201****
2019	137,869	115,512	**0.229****	-0.087	-0.028	**0.213****
2020	194,728	161,889	**0.184****	**-0.176****	-0.021	**0.213****

******, ***** - significant at the 1% and 5% level, respectively.

Table 10: Regressions run on transaction subsets which each exclude transactions made in a calendar month. Transactions include those made in direct accounts (investor makes the investment decisions) and delegated accounts (decisions are delegated to an investment manager). Dependent variable is a dummy equal to 1 when a sale occurred, Gain is a dummy equal to when the transaction price is above the historic average purchase price, Del is a dummy equal to 1 for transactions in delegated accounts. T-statistics in brackets are based on standard errors two-way clustered by account and date.

	Transactions	**Delegated**	**Gain**	**Gain x Del**	**Del**	**Constant**
Ex January	829,306	564,360	**0.265*** **	**-0.103*** **	0.0194	**0.173*** **
			(21.881)	(-3.006)	(1.513)	(29.201)
Ex February	832,952	571,712	**0.262*** **	**-0.104*** **	0.0185	**0.174*** **
			(22.154)	(-3.062)	(1.459)	(31.466)
Ex March	826,872	565,660	**0.271*** **	**-0.094*** **	0.0118	**0.171*** **
			(23.003)	(-2.793)	(1.068)	(34.107)
Ex April	805,727	542,648	**0.267*** **	**-0.115*** **	**0.028** **	**0.176*** **
			(22.104)	(-3.271)	(2.608)	(30.698)
Ex May	831,785	568,907	**0.268*** **	**-0.101*** **	0.0214	**0.173*** **
			(22.293)	(-2.915)	(1.704)	(29.436)
Ex June	835,448	573,410	**0.271*** **	**-0.106*** **	0.0203	**0.174*** **
			(22.181)	(-3.053)	(1.580)	(29.298)
Ex July	834,091	570,422	**0.268*** **	**-0.094*** **	0.0195	**0.173*** **
			(22.011)	(-2.664)	(1.541)	(29.225)
Ex August	816,504	553,012	**0.270*** **	**-0.095** **	0.0174	**0.174*** **
			(22.337)	(-2.607)	(1.361)	(29.658)
Ex September	831,123	567,067	**0.269*** **	**-0.095*** **	0.0215	**0.174*** **
			(21.976)	(-2.691)	(1.691)	(29.425)
Ex October	812,994	552,304	**0.269*** **	**-0.102*** **	0.0242	**0.174*** **
			(21.968)	(-2.876)	(1.814)	(28.941)
Ex November	817,811	556,615	**0.265*** **	**-0.074*** **	0.0201	**0.175*** **
			(21.891)	(-2.921)	(1.543)	(29.158)
Ex December	822,406	557,917	**0.270*** **	**-0.096*** **	0.0188	**0.175*** **
			(22.381)	(-2.674)	(1.455)	(30.057)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 11: Regressions run on all transactions as in Equation 7. Dependent variable is a dummy equal to 1 when a sale occurred, *Gain* is a dummy equal to when the transaction price is above the historic average purchase price, *Del* is a dummy equal to 1 for transactions in delegated accounts. *Gain* and *Del* dummies are interacted with *SENT* (the 5 day change in the Daily News Sentiment Index Shapiro, Sudhof, and Wilson (2020)). T-statistics in brackets are based on standard errors two-way clustered by account and date.

	#Transactions	**#Delegated**	**Gain**	**Gain x Del**	**Del**	**Gain x SENT**	**Gain x Del x SENT**	**Del x SENT**	**Constant**
All Transactions	899,729	613,094	0.271***	-0.068***	0.022**	0.037***	0.054***	0.005	0.174**
			(23.7)	(-2.597)	(2.395)	(4.167)	(2.681)	(0.336)	(31.138)
Single Stocks	439,799	256,181	0.257***	0.005	0.014	0.045***	-0.001	0.004	0.178**
			(20.574)	(0.219)	(1.601)	(3.902)	(-0.046)	(0.557)	(28.999)
Active Funds	292,559	252,132	0.227***	-0.078	0.059***	0.029**	0.074*	0.007	0.136**
			(9.605)	(-1.519)	(3.189)	(2.024)	(1.854)	(0.267)	(14.568)
Index Funds & ETFs	60,392	45,275	0.213***	-0.074	-0.022	-0.022	0.057	-0.014	0.171**
			(6.276)	(-1.097)	(-0.794)	(-1.112)	(1.311)	(-0.858)	(8.609)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

References

Ahn, Y. (2021). “The anatomy of the disposition effect: Which factors are most important?” In: *Finance Research Letters* 38, p. 101818.

An, L., J. Engelberg, M. Henriksson, B. Wang, and J. Williams (2019). “The Portfolio-Driven Disposition Effect”. In: *PBCSF-NIFR Research Paper* 06, pp. 1–51.

Asness, C., T. Moskowitz, and L. J. Pedersen (2013). “Value and Momentum Everywhere”. In: *Journal of Finance* 68.3, pp. 929–985.

Barber, B. M., Y. Lee, Y. Liu, and T. Odean (2007). “Is the Aggregate Investor Reluctant to Realise Losses? Evidence from Taiwan”. In: *European Financial Management* 13.2, pp. 423–447.

Barber, B. M., Y. Lee, Y. Liu, and T. Odean (2009). “Just how much do individual investors lose by trading?” In: *Review of Financial Studies* 22.2, pp. 609–632.

Barber, B. M. and T. Odean (2002). “Online investors: Do the slow die first?” In: *Review of Financial Studies* 15.2, pp. 455–488.

Barberis, N. and W. Xiong (2009). “What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation”. In: *Journal of Finance* 64.3, pp. 751–784.

Barberis, N. and W. Xiong (2012). “Realization Utility”. In: *Journal of Financial Economics* 105.2, pp. 251–275.

Bashall, J., G. Willows, and D. West (2018). “The Extent to Which Professional Advice Can Reduce the Disposition Effect: An Emerging Market Study”. In: *Journal of Emerging Market Finance* 17.3, pp. 272–288.

Bazley, W., J. Moore, and M. M. Vosse (2021). “Taxing the Disposition Effect: The Impact of Tax Awareness on investor Behavior”. In: *Journal of Financial and Quantitative Analysis* 56.2, pp. 461–495.

Bergsma, K., A. Fodor, and E. Tedford (2019). “A Closer Look at the Disposition Effect in U.S. Equity Options Markets”. In: *Journal of Behavioral Finance* 12.4, pp. 499–516.

Bernard, S., B. Loos, and M. Weber (2021). “The Disposition Effect in Boom and Bust Markets”. In: *SAFE Working Paper* 309, pp. 1–60.

Birru, J. (2015). “Confusion of Confusions: A Test of the Disposition Effect and Momentum”. In: *Review of Financial Studies* 28.1, pp. 118–152.

Buhlmann, F., P. Doerrenberg, B. Loos, and J. Voget (2020). “How do taxes affect the trading behavior of private investors? Evidence from individual portfolio data”. In: *ZEW - Centre for European Economic Research Discussion Paper No. 20-047* 20.47, pp. 1–44.

Calvet, L., J. Campbell, and P. Sodini (2007). “Down or out: Assessing the welfare costs of household investment mistakes”. In: *Journal of Political Economy* 115.5, pp. 707–747.

Calvet, L., J. Campbell, and P. Sodini (2009a). “Fight or flight? Portfolio rebalancing by individual investors.” In: *Quarterly Journal of Economics* 124.1, pp. 301–348.

Calvet, L., J. Campbell, and P. Sodini (2009b). “Measuring the financial sophistication of households”. In: *American Economic Review* 99.2, pp. 393–398.

Chang, T. Y., D. H. Solomon, and M. M. Westerfield (2016). “Looking for someone to blame: delegation, cognitive dissonance, and the disposition effect”. In: *Journal of Finance* 71.5, pp. 2671–2706.

Cheng, T., C. Lee, and C. Lin (2013). “An examination of the relationship between the disposition effect and gender, age, the traded security, and bull-bear market conditions”. In: *Journal of Empirical Finance* 20.4, pp. 764–786.

Chevalier, J. and G. Ellison (1997). “Risk Taking by Mutual Funds as a Response to Incentives”. In: *Journal of Political Economy* 105.6, pp. 1167–1200.

Choe, H. and Y. Eom (2009). “The Disposition Effect and Investment Performance in the Futures Market”. In: *Journal of Futures Markets* 29.11, pp. 2929–2942.

Da Silva Rosa, R., H. To, and T. Walter (2005). “Evidence contrary to the disposition effect among UK managed funds.” In: *Proceedings of Accounting and Finance Association of Australia and New Zealand Conference 2005* 40, pp. 29–38.

Dhar, R. and N. Zhu (2006). “Up close and personal: Investor sophistications and the disposition effect”. In: *Management Science* 52.5, pp. 726–740.

Dorn, D. and G. Huberman (2005). “Talk and action: what individual investors say and what they do.” In: *Review of Finance* 60.1, pp. 223–261.

DØskeland, T. and H. Hvide (2011). “Do individual investors have asymmetric information based on work experience?” In: *Journal of Finance* 66.2, pp. 571–604.

Eom, Y. (2018). “The opposite disposition effect: Evidence from the Korean stock index futures market”. In: *Finance Research Letters* 25, pp. 142–146.

FCA (2019). *Investment Platforms Market Study*. https://www.fca.org.uk/publication/market-studies/ms17-1-3.pdf.

Feng, L. and M. Seasholes (2008). “Individual investors and gender similarities in an emerging stock market.” In: *Pacific-Basin Finance Journal* 14.1, pp. 69–97.

Frina, A., G. Lepone, and D. Wright (2015). “Investor characteristics and the disposition effect”. In: *Pacific-Basin Finance Journal* 34, pp. 227–246.

Frino, A., D. Johnstone, and H. Zheng (2004). “The propensity for local traders in futures markets to ride losses: Evidence of irrational or rational behavior?” In: *Journal of Banking Finance* 28.10, pp. 2485–2510.

Genesove, D. and C. Mayer (2001). “Loss aversion and seller behavior: Evidence from the housing market”. In: *Quarterly Journal of Economics* 116.4, pp. 1233–1260.

Grinblatt, M. and M. Keloharju (2001). “What makes investors trade?” In: *Journal of Finance* 56.2, pp. 589–616.

Grinblatt, M., M. Keloharju, and J. Linnainmaa (2012). “IQ, trading behavior, and performance”. In: *Journal of Financial Economics* 104.2, pp. 339–362.

Henriksson, M. (2021). “Disastrous Selling Decisions: The Disposition Effect and Natural Disasters”. In: *Household Finance eJournal* 3.58, pp. 1–31.

Hincapie-Salazar, J. and D. Agudelo (2019). “Is the disposition effect in bonds as strong as in stocks? Evidence from an emerging market”. In: *Global Finance Journal* 40, pp. 19–29.

Hirschleifer, D. and I. Ben-David (2012). "Are investors really reluctant to realize their losses? Trading responses to past returns and the disposition effect". In: *Review of Financial Studies* 25.3, pp. 774–792.
Ivkovich, Z. and S. Weisbenner (2009). "Individual Investor Mutual Fund Flows". In: *Journal of Financial Economics* 92.2, pp. 223–237.
Kahneman, D. and A. Tversky (1979). "Prospect theory: An analysis of decisions under risk". In: *Econometrica* 47.2, pp. 263–291.
Kaustia, M. (2010). "Prospect theory and the disposition effect". In: *Journal of Financial and Quantitative Analysis* 45.1, pp. 109–133.
Odean, T. (1998). "Are investors reluctant to realize their losses". In: *Journal of Finance* 53.5, pp. 1775–1798.
Oreng, M., C. Yoshinaga, and W. Junior (2021). "Disposition effect, demographics, and risk taking". In: *RAUSP Management Journal* 56.1, pp. 22–37.
Richards, D., J. Rutterford, D. Kodwani, and M. Fenton-O'Creevy (2015). "Stock Market Investors' Use of Stop Losses and the Disposition Effect". In: *The European Journal of Finance* 21.11, pp. 1567–1580.
Schlarbaum, G., R. Lease, and W. Lewellen (1978). "The common-stock-portfolio performance record of individual investors". In: *Journal of Finance* 33.2, pp. 471–483.
Shapira, Z. and I. Venezia (2001). "Patterns of behavior of professionally managed and independent investors". In: *Journal of Banking & Finance* 25.8, pp. 1573–1587.
Shapiro, A., M. Sudhof, and D. Wilson (2020). "Measuring news sentiment". In: *Journal of Econometrics* 220.2, pp. 374–395.
Shefrin, H. and M. Statman (1985). "The disposition to sell winners too early and ride losers too long: theory and evidence." In: *Journal of Finance* 40.3, pp. 777–790.
Thaler, R. (1985). "Mental accounting and consumer choice". In: *Marketing Science* 4.3, pp. 199–214.
Tibshirani, R. (1996). "Regression shrinkage and selection via the lasso". In: *Journal of the Royal Statistical Society: Series B (Statistical Methodological)* 58.1, pp. 267–288.
Woo, J. S. (2016). "Do Institutional Investors Show the Disposition Effect? : Evidence from South Korea's Bond Market". In: *Asia-Pacific Journal of Financial Studies* 45.3, pp. 466–485.

Magical Internet Money? On-chain Cashflows and the Cross-section of Cryptocurrency Returns

Ainsley To

6 July 2023

Abstract

I find that crypto valuation measures derived from on-chain fundamental cashflow characteristics, analogous to valuation metrics used in equity markets, are priced in the cross-section of token returns. A cashflow-based value factor constructed from these measures is not spanned by crypto factor models in the literature. I test different measures of cashflow and find that revenues retained by protocols show the strongest results, whilst token incentives as a cost of revenue measure have little pricing power. I also find evidence that different characteristics are significant for native tokens of a blockchain compared to tokens issued by decentralised applications. Lastly, I test a set of novel crypto native characteristics, unique to public blockchains, that proxy for capital gains overhang, insider ownership, and investor sophistication.

The author gratefully acknowledges helpful comments from Laurent Calvet,

The views or opinions expressed in this article are those of the author only and do not necessarily represent the views and opinions of any other organization or any of their affiliates.

10 Introduction

Aside from the operational overheads to trading digital assets, a significant barrier for many institutional investors is that there is little consensus over valuation methodologies in crypto. A key issue is the common misconception that cryptocurrencies have no cash flows, a central input for fundamental investors when building valuation models. The cross-section of tokens in the crypto space displays more heterogeneity than other asset classes, which is intuitive given the open source nature of protocol development and token economic design. Many cryptocurrencies are hybrid securities, exhibiting properties from a number of different traditional asset classes. Because of their hybrid nature, it is no surprise that the early research within the nascent crypto asset pricing literature has also yet to reach consensus on fundamental drivers of token prices. Having identified cashflow characteristics for cryptocurrencies, I document that equivalent measures to valuation metrics used in equity markets have efficacy in pricing the cross-section of token returns.

Earlier research in the cryptocurrency space used crypto-specific metrics to construct "Network Factors" based on the usage and size of the network. Athey et al. (2016) model bitcoin exchange rates against fiat currencies as a function of its usage as a payment vehicle. Using transaction volume as a measure of demand Pagnotta and Buraschi (2018) attempt an equilibrium valuation of bitcoin using the hashrate as a measure of production - so the hashrate measures the amount of computing power being used to process transactions. In Biais et al. (2023), they create proxies for net transactional benefit of bitcoin and build a general equilibrium model around that. Cong, Y. Li, and N. Wang (2021) build a model for tokens that also considers transactional benefits but with network externalities, since tokens enable users to capitalize on platform growth, which reduces their transactions costs on the platform, so demand is dependent on demand from others as the network grows. Sockin and Xiong (2023) consider the impact of speculator sentiment on the crowding out of users of a platform. Variables such as the number of active addresses and the total number of transactions have also been used Cong et al. (2022), who find total addresses with a balance span other metrics related to network adoption. However, one weakness of using address metrics is that they measure ownership and interactions of users with a token or token address and not the interactions of users with the underlying protocol itself, thus providing a limited gauge of true adoption.

There are also studies linking the evolution of cryptocurrency prices to the marginal cost of production, namely mining costs Cong, He, and J. Li (2018), which includes costs of electricity and the price of hardware used in cryptocurrency mining.

Hu, Parlour, and Rajan (2019) is an empirical study and present stylized facts about a number of smaller coins, with a main conclusion that all other cryptocurrencies are strongly correlated with Bitcoin returns. Borri (2019) highlights the conditional tail risk for the larger cryptocurrencies to the overall crypto market but also finds that cryptocurrencies are not exposed to tail risk in other asset classes. Liu and Tsyvinski (2021) is a broad study considering a number of macro factors, some of the network and production factors discussed above, as well as 150 equity factors from Andrew Chen's factor zoo website - they find time series momentum constructed on cryptocurrency returns to be one of the key factors. Liu, Tsyvinski, and Xi (2022) is a key study within crypto asset pricing, which looks at cross-sectional factors within cryptocurrency, constructing crypto Market, Size, and Momentum factors from token prices to explain the cross-section of cryptocurrency returns. They find that momentum behaves more consistently with the theory of investor overreaction since they also observe long term reversals, however they notably do not identify a crypto value factor despite testing a number of other characteristics.

Other studies look for a proxy of crypto fundamentals to construct a crypto value factor, including Cong et al. (2022) who use a number of crypto native characteristics and test their ratio to market value in portfolio sorts, including number of on-chain transactions, cumulative number of addresses, and number of addresses with a balance. However they find that none generate long-short portfolios with significant excess returns and instead define their crypto value factor using price reversal, analogous to specifications for value factors constructed in currency and commodity markets (Asness, Moskowitz, and Pedersen (2013)).

In this paper I refute the common notion that cryptocurrencies do not have cashflows and document evidence these are priced in the cross-section of token returns. Earlier studies were focused on a narrow subset of coins representing a specific type of cryptocurrency, namely native tokens of Layer-1 blockchains. But since the advent of smart contract platforms such as Ethereum, the token universe has exploded both in number and in the breadth of intended use cases. Many of the tokens issued by decentralised applications (a more detailed description below) that are built on the application layer of smart contract platforms are distinct from the native tokens issued as block rewards and used for transaction fees - Dapp tokens often more closely resemble securities in other asset classes such as equities, earning fees from protocol activity which can accrue to tokenholders.

Whilst previous studies such as (Cong et al. (2022)) attempt to model cryptocurrencies similarly to commodity and currency markets and measure the value factor with long term reversal (defined as the negative of 1 year returns, vis-à-vis Asness, Moskowitz, and Pedersen (2013) who use the negative of the average log price between 4.5 and 5.5 years ago) or other proxies for demand (such as number of users, addresses, and volume), I find that using fundamental cashflow characteristics more analogous to those used in equity cross sectional factor models (dating back to Fama and French (1992)) can price the cross section of token returns.

A brief description of decentralised applications (Dapps) should highlight the similarities between tokens issued by Dapps and traditional securities. There is a large and heterogenenous universe of Dapps built on the application layer of smart contract platforms. Dapps are smart contracts that utilise the cryptographic security of the underlying blockchain they are built on for authenticating transactions made with and by the application. Tokens issued by Dapps are non-native to the blockchain they are built on (in contrast to a native token such as ETH, which is used to pay transaction fees on the Ethereum blockchain), and are themselves smart contracts, often built following a specific token standard (such as ERC-20 on Ethereum) that enable them to interact with other smart contracts. Each token has their own design features for how they interact with their underlying Dapp and other smart contracts, as well as their own "token economics" for how they accrue value to tokenholders - common features include a share of revenues earned by a Dapp and voting rights for participating in protocol governance. In this paper, I also document evidence that the ability of on-chain revenues to explain the cross-section of token returns is more significant in tokens issued by Dapps than for native tokens issued by Layer-1 blockchains. Though it should be noted that not all applications issue tokens, and not all tokens earn revenues - the extent to which the existence of a token impacts application adoption is beyond the scope of this study and a fruitful area for further research.

I contribute to the literature in a number of ways. Firstly, I test a new set of factors using crypto native fundamental data derived from public blockchains. Secondly, I show that value factors based on on-chain cashflow characteristics are priced in the cross section of token returns. Thirdly, I provide evidence that these factors are not spanned by the crypto factors in the current literature. Finally, I shed light on the hybrid nature of cryptocurrencies and how a value factor similar to those used in equities is more applicable for token resembling securities, such as those linked to decentralised applications (and less so for tokens native to Layer-1 blockchains).

The rest of the paper is laid out as follows. Section 11 provides an overview of the on-chain fundamental characteristics in this study and details the nuances between crypto fundamentals in Layer-1s and Dapps. Section 12 outlines the data sources used in the analysis and also the unique aspects of cryptocurrencies that make it distinct from similar analysis in traditional markets. Section 13 covers the methodology and results, including for portfolio sorts, subsets by token type, Fama-MacBeth regressions, and spanning regressions. Section 14 concludes.

11 On-chain Fundamentals

11.1 Crypto Revenues

Between June 2017 and December 2022, over $35 billion in value has been generated by cryptocurrencies through on-chain activity (Figure 2). Which participants in the ecosystem capture this value, whether it be tokenholders or other actors, varies depending on the nature of the protocol in question, as well as its specific token economics. For instance, in the case of ETH, the native token of the Layer-1 smart contract platform Ethereum, fees are earned from users transacting on the Ethereum network. A proportion of these fees are paid to supply side users, specifically in the case of Ethereum, tips are paid to validators for providing security to the blockchain (and previously paid to miners during Ethereum's Proof-of-Work era). ETH tokenholders also receive a portion of the transaction fees but in the form a token burn, in which tokens are removed from circulation and the outstanding supply of ETH is reduced. The amount burned varies depending on the number of transactions on the network and has at times been significant enough to offset ETH's token inflation (discussed in subsection 11.2).

Whilst the above describes a common token economic model for Layer-1s, the fee accrual mechanism is dramatically different for tokens issued by Dapps, which are themselves smart contracts built on platforms such as Ethereum. The Dapp universe is more heterogenous and revenue models are often idiosyncratic to the individual protocol. In decentralized exchanges (DEXs) for example, users swapping tokens on the DEX pay a transaction fee, a share of which can be retained by the protocol itself, with the remainder captured by supply side participants such as liquidity providers supplying their tokens for market making. In the DEX Uniswap for instance, 100% of the fees accrued to the liquidity providers[1], though there have been governance proposals in the Uniswap DAO (the decentralised autonomous organisation governing the protocol) to explore increasing the share of fees that would flow to UNI tokenholders. Another prominent type of Dapp are lending protocols, which facilitate borrowing and lending of tokens between users. A common revenue model for lending protocols is to earn a share of the interest paid on loans outstanding and also liquidation fees paid by borrowers whose collateralisation ratio (the spread between deposited collateral and amount borrowed) falls below the liquidation threshold. As with DEXs, the share of these fees earned by the lending protocol versus suppliers of capital (lenders) varies. The proportion of value retained by the protocol is known colloquially as the "take rate" - just under half of protocols in the data sample (described in Section 12) had either a 0% or 100% average take rate in 2022[2], with 61% of protocols capturing less than half of total fees vis-à-vis supply side participants (Figure 3). The choice of take rate appears to be an idiosyncratic choice at the protocol level, as cross-sectional correlation with market cap and total fees generated across tokens is negligible.

In terms of economic magnitude, the cumulative revenue accrued in Figure 2 is against a total crypto market cap that has averaged $820 billion over the same period, having peaked just shy of $3 trillion in November 2021 (Figure 7). It is also worth noting that the subset of tokens which earned non zero fees only represent 70% of total market cap as at the end of 2022.

From a valuation perspective, the median ratios of market capitalisation over revenues for the token universe as at the end of 2022 were 9.2, 3.8, and 3.2 depending on the scope of cashflows included. There is also a wider dispersion for the more granular protocol revenue and supply side revenue measures (Figure 4) compared to the global equity universe.

[1] as at the end of 2022

[2] Averaged over daily Protocol and Supply Side Revenue data in calendar year 2022

11.2 On-chain Cost of Revenue and Expenses

Measuring the recurring costs associated with running a crypto protocol is more challenging, given the decentralised and autonomous nature of blockchains and Dapps - prima facie there are no employee related expenses or costs associated with running fixed assets in the traditional sense, though for some projects there may be an associated legal entity that formally employs the development team and incurs other expenses similar to a traditional corporation. A common practice for many protocols is to mint new tokens as economic incentives for users to perform certain functions - most famously for Bitcoin and other Layer-1s, new native tokens are minted as block rewards to incentivise miners or validators to add new blocks and maintain the network. These maintenance costs paid through token inflation (newly minted tokens increase outstanding supply) are more analogous to the process used by central banks in operating fiat currency systems. Many Dapps use similar mechanisms to incentivise activity within their protocol, such as minting new governance tokens as additional rewards to attract liquidity providers, who can earn the rewards by supplying tokens on the DEX for users to trade. Some protocols have a hard-coded inflation schedule for how the rate of incentives will be distributed over time - Bitcoin famously has a fixed halving schedule, where block rewards halve in value as the blockchain reaches preset block height thresholds (occurring approximately every 4 years).

The value of on-chain token incentives that are paid to participants in a protocol ecosystem are openly available for public blockchains. Given the novel nature of crypto token economic models, there is no current consensus over the most relevant comparison for token inflation rates. Figure 6 compares the average daily token inflation rates (token incentives divided by market capitalisation) to inflation rates for countries within the IMF World Economic Outlook database and also to the percentage change in shares outstanding for global equities over calendar year 2022.

There is a wide dispersion amongst inflation rates within the cross-section of tokens, with a median of 6.7% and a maximum of 140% in 2022. Though an imperfect comparison (further discussion below), this was not significantly different from the distribution of global CPI inflation rates for fiat currency based economies in 2022, which saw a median and maximum of 8.1% and 201% over the same period - the largest coins BTC and ETH (the native tokens for the Bitcoin and Ethereum networks respectively) saw 2022 inflation rates of 1.7% and 3.2% respectively, which were at the lower end of developed market currency blocs over the period and also comparable to long term central bank inflation targets.

However the equity market comparison paints a different picture, with the median firm actually contracting their share count by 6.9% in 2022. Changes in shares outstanding also have a significantly tighter range than either token or fiat currency inflation rates. As stated previously, given their hybrid nature, there is still little consensus on the most relevant comparisons for cryptocurrencies within the universe of traditional asset classes and interpretations of token inflation rates relative to metrics in either currency or listed equities should be heavily caveated.

A deeper analysis of the specific trade-offs of either comparison for token inflation specifically are beyond the scope of this study but potential grounds for further research.

One methodology for calculating the earnings of a protocol is to take the net value of Protocol Revenue (discussed in section 11.1) and Token Incentives (discussed above). There are a number of issues with this intepretation. Firstly, revenue is often received as a nominal amount whereas the costs as measured by token incentives are often borne in real terms in the form of inflation. Secondly, many protocols allow token holders to stake their tokens and receive a share of new token issuance, for example following Ethereum's transition to a Proof-of-Stake consensus mechanism, ETH holders can stake their tokens to participate in network security, allowing them to earn a share of token incentives (staking rewards in this case) and blurring the line between tokenholders (beneficiaries of protocol revenues) and recipients of incentives. For completeness, I also test the efficacy of this measure of earnings and find no evidence it can price the cross-section of token returns.

11.3 Characteristics of Tokenholder Base

In addition to data on protocol fundamentals, the transparency of public blockchains also provides a degree of visibility into token ownership. The pseudonomous nature of wallet addresses makes a thorough analysis at the tokenholder level more complex - the same entity could have multiple addresses. There are both heuristic based and statistical clustering methods to aggregate addresses belonging to the same entities (Meiklejohn et al. (2016),Möser and Narayanan (2022)), as well as open source attempts to tag wallets belonging to larger well known entities (*Etherscan* n.d.), however they each come with subjective design choices.

Above caveats not withstanding, I test a number of tokenholder metrics that proxy for significant investor characteristics in the literature and measure their relationship with the cross section of token returns. As a proxy for investor sophistication (which has been documented to be inversely related to poor investment decisions Dhar and Zhu (2006),Calvet, Campbell, and Sodini (2009b)), I use the dollar value of wallet balances and aggregate the percentage of a token's outstanding supply held by wallets with large balances ($\geq$ \$1 million).

Additionally, I look at addresses that hold a large percentage of the outstanding token supply ($\geq$ %1) as a proxy for insider and/or institutional ownership of a token - the literature is mixed on this area within the equity universe (Nofsinger and Sias (1998) finds a positive relationship between changes in institutional ownership and stock returns, Cai and L. Zheng (2004) find a negative relationship when looking at lagged institutional trading with some evidence for reverse causality).

Thirdly, I test market value to realised value (MVRV), a variable unique to the cryptocurrency asset. Realised value measures the total value of a token using the last price at which each token was traded on-chain by each tokenholder address. A higher ratio between market value and realised value is a proxy for aggregate unrealised gains for holders of a token, which has been documented as a driver of future stock returns (Grinblatt and Han (2005),Barberis, Jin, and B. Wang (2021)).

Finally, I also use address count data for each token, similar to the NET factor in Cong et al. (2022), as a control for effects associated with network adoption.

I find mixed results within the tokenholder characteristics set. Whilst tokens with a larger proportion of supply held by "whales" have outperformed, MVRV is insignificant and the extent to which tokens with a larger proportion of sophisticated investors (as proxied by \$ wallet ballance) is dependent on whether I control for market capitalisation.

11.4 Token Subsets

As discussed in subsections 11.1 and 11.2, there are nuances to the nature in which different types of protocol accrue revenues and distribute incentives. For Layer-1 blockchains, whose native token is commonly used to pay for transaction fees for securing the network, incentives are mining or staking rewards that flow to miners or stakers for providing network security, whilst revenues come in the form of transaction fees that are burnt to reduce the supply of the native token. For tokens issued by Dapps such as decentralised exchanges for example, token incentives are often distributed to liquid providers who have deposited in the liquidity pools on the DEX, and revenues accruing to the protocol are the share of trading fees that flow to the protocol treasury.

In an attempt to control for these differences, which mainly exist between Layer-1 tokens and Dapp tokens, I rerun the analysis on the two separate subsets and find that results for revenue measures are more significant for the Dapp subset and are no longer significant for the Layer-1 subset. There have been a number of initial attempts to formally classify the token universe into sectors (Datanomy, Cong et al. (2022)). For the purposes of this analysis, I define Layer-1s tokens as native tokens of protocols which are not reliant on another protocol for network security (e.g. projects which are not built on top of another blockchain) - this definition includes both protocols that are primarily used for payments such as Bitcoin, and smart contract platforms such as Ethereum. The distinction between Layer-1s and Dapps is relevant since the interpretation of protocol revenues and costs are most different between these two types of protocols as noted above.

12 Data

A prominent property of public blockchains is the transparency of on-chain activity. A growing number of data providers aggregate a variety of on-chain metrics that are unique to the crypto asset class.

Token Terminal is a data aggregation platform that have compiled fundamental data across a universe of over 170 blockchains and decentralised applications. They aggregate on-chain data to a daily frequency and make each variable available via their API as far back as 2017. As discussed in sections 11.1 and 11.2, the specific variables I use are: total fees earned by protocols ($FEES$), the fees retained by the protocol ($PREV$), fees earned by supply side actors ($SREV$), and token incentives distributed ($TOKINC$). I also use their formulation of earnings ($EARN$) which is the net of ($PREV$) minus ($TOKINC$).

I use token price and market capitalisation data gathered from CoinGecko, which aggregates the data for tokens which fulfil their listing criteria across crypto exchanges that meet their API standards - prices are a volume weighted average across eligible exchanges. Not all the protocols with fundamental data in Token Terminal have prices available - the NFT platform Opensea for example, has accumulated over $2.7bn in fees to the end of 2022 but have yet to issue a token. Of the sub sample of 170 protocols with fundamental data, 145 have tokens with prices on CoinGecko, which represent 75% of total crypto market cap. From these, I exclude any token which has any individual daily price change over 10,000% (resulting in two tokens being excluded as they appear to have erroneous price data).

The token universe is sparse prior to 2020, with the number of projects increasing exponentially following the advent of “DeFi Summer” in the summer of

2020, when a number of new primitives gained traction on Ethereum, spurring dramatic growth in new token launches. Though there are some select fundamental metrics available from on-chain data prior to 2020, these are only available for a small number of tokens and limited mostly to payment tokens and smart contract platforms, with few Dapps available before this date. Thus, for data availability reasons the results in this paper are for the three year period from 2020 to the end of 2022. However, it is worth noting that similar results can be obtained on the smaller subset where longer term data on specific metrics are available. For the additional analysis on Layer-1 and Dapp subsets (discussed in section 11.4), the limited number of Dapp tokens with fundamental data available makes portfolio sorts for that subset difficult before 2021, hence that specific analysis is run from 2021 to 2022 for both subsets for a fair comparison.

Separately, crypto data platform IntoTheBlock have compiled a wide variety of on-chain metrics including on tokenholder characteristics available at a daily and weekly frequency. As outlined in section 11.3, the specific metrics I use are market value to realised value ($MVRV$), the number of addresses with a non-zero token balance ($ADDNZ$), the proportion of addresses which hold more than %1 of token supply (WHA), and the proportion of addresses which hold more than \$1m in the token (BAL). These metrics are all available for the set of tokens with fundamental data from Token Terminal, with the exception of $MVRV$, which as of March 2023 was only available for 100 of the protocols in the fundamental set.

A unique aspect of asset pricing studies in crypto is that data on public blockchains is more transparent and frequently available than for traditional securities - rather than earnings figures being released at quarterly or semi annual intervals, live up to date protocol fundamental data is available with every new block added (approximately every 14 seconds on Ethereum) and in a transparent and consistent form that has been pre-programmed within the smart contracts (free from the discretions of earnings management practices). However, despite the higher frequency data, some fundamental metrics show a degree of persistence over time. Table 13 shows the regression of future fundamentals on current fundamentals across various time lags - a large majority of the coefficients are economically large and statistically significant with high R^2 values. A noticeable pattern is the meaningful drop off in explanatory power that many metrics experience beyond 60 days. Non-zero addresses, the percentage of supply held by whales (addresses with more than 1% of supply), and percentage of supply held by large balances ($\geq$\$1m) show strong persistence as far out as 180 days.

An exploratory data analysis also reveals an intuitive cross-sectional correlation between fundamental characteristics and market capitalisation - the revenue metrics, token incentives, and number of addresses show a positive relationship with token size (Table 12), since the data is in US dollar terms. Total wallet addresses with nonzero balance also increase with market cap. Interestingly, the proportion of wallet addresses holding more than 1% of outstanding supply has a moderately negative correlation to market capitalisation, suggesting that larger tokens are more widely distributed and have less concentrated ownership, though there is no relationship between size and the proportion of wallet addresses with more than \$1 million. To control for the effect of size, in the analysis I divide each metric by market capitalisation, with the exception of *MVRV*, *WHA*, and *BAL*.

13 Fundamentals in Crypto

Using fundamental characteristics derived from on-chain data, I construct factor sorted token portfolios similar to Fama and French (1992). On a daily basis, tokens are sorted on the value of the characteristic from the previous day and split into three portfolios ("Lo", "Med", "Hi"), with the factor return constructed as the spread between the Hi and Lo portfolios.

Liu, Tsyvinski, and Xi (2022) use value-weighted (VW) portfolio returns in their analysis. However, the consequence of value weighting is more significant within the crypto asset class, given the historic concentration of weights into BTC and ETH (Figure 7). The two largest tokens have average 66% of total crypto market cap between 2020 and 2022. This is important when interpreting the VW results, which can be driven to a large extent by which sorted portfolio the largest tokens fall into between rebalancing periods. For this reason, going forward I report results using equally-weighted (EW) portfolios. The choice of EW may limit the extent to which the factor can be implemented in practice and implementation costs may be an interesting area for further research.

As per Table 13, the persistence of the fundamental measures make the analysis less sensitive to the precise measurement period for the characteristics used for the portfolio sorts. Going forward, I report the figures for portfolios sorted on a daily basis using fundamental characteristics smoothed over 7 days as certain tokenholder characteristics from IntoTheBlock were only available at weekly frequency for select tokens. It is worth noting that the broader results are robust for measurement windows less than 60 days and also directionally similar whether using daily or weekly portfolio formation.

Table 14 reports the mean returns and t statistics for the full token subset. Portfolio sorts are positive and statistically significant across *FEES PREV*, and *SREV*, each with a monotonic increase in mean return and significant from Lo to Hi, suggesting a broad relationship between revenues and future token returns that is robust to the specific revenue measure used. However, the spread between the Hi and Lo portfolios formed on token incentives is small and not statistically different from zero. Consistent with Cong et al. (2022), addresses with non-zero addresses are statitically significant, other characteristics of the investor base show mixed results - *MVRV* and *WHA* are not statistically significant, and Hi-Lo fpr *BAL* was negative and significant, suggesting that tokens with high percentage of addresses holding $\geq \$1m$ of the token outperformed those with a lower percentage.

13.1 Double Sorts with Size

I also look at independent double sorts of each variable with size. Tables 15 shows the results for 2x3 independent portfolio sorts on size and the revenue characteristics, with Big representing the top 50% largest tokens by market capitalisation and Small the bottom 50%, whilst Table 15 shows the same for tokenholder characteristics. Whilst *PREV* is significant in both size brackets, *SREV* and *FEES* are only significant in the bottom 50% of tokens. As discussed in section 13.2, the distinction between results in Big vs Small may be driven by the differences in portfolios composition of big and small portfolios - across the token universe, Layer-1s have tended to be higher in market cap and are over represented in the Big portfolios, with Dapps underrepresented (Figure 8). Monegro (2016) provides a theory (known colloquially as the "Fat Protocol Thesis") for why Layer-1 tokens should accrue more value than Dapps built on top. Providing a formal test for the Fat Protocol Thesis is beyond the scope of this study, though I find evidence that revenue metrics are more applicable for pricing Dapp tokens than Layer-1s (section 13.2).

In addition, *TOKINC* is negative and significant in the Big sample, with only the portfolio containing tokens issuing the lowest amount of incentives having average returns statistically different from zero. Within the Small sample, the return difference between portfolios with the lowest and highest incentives is not statistically significant.

Table 16 shows results for the double portfolio sorts on tokenholder characteristics. These variables do not show significant results with the exception of *ADDNZ* in the Small token subset, where addresses with a non zero balance are positive related to token returns.

13.2 Layer-1s vs Decentralised Applications

As discussed in section 11.4 and section 13.1, the nature of revenue accrual and issuance of token incentives differs across the token universe, with a particularly stark difference between native Layer-1 tokens and tokens issued by Dapps. This section compares results for the analysis within these two token subsets. Due to data limitations for Dapps with fundamental characteristics before summer of 2020 as discussed in section 12, the analysis on the two subsets are instead run from the beginning of 2021 to the end of 2022. The portfolio sorts in the Layer-1 and Dapp subsets show that whilst revenue characteristics are strongly significant within tokens issued by Dapps (Table 17), these are not significant within Layer-1 native tokens. This is in line with intuition that equity like metrics are more effective at pricing Dapp tokens which more closely resemble securities, whereas native tokens in smart contract platforms serve different purposes such as to pay for transaction fees or for staking to secure the network in the case of Proof-of-Stake chains.

Within the Layer-1 subset (Table 18), *FEES*, which for this subset represent the sum of transaction fees flowing to validators plus the reduction in supply from transaction fees burnt, remains significant. *ADDNZ* and *WHA* are significant for both subsets, suggesting that network adoption and holder concentration are equally important for both types of tokens. Interestingly, token incentives are not significant for the Layer-1 subset, despite the notion that bitcoin's fixed supply might make it appealing relative to fiat currencies - there is no evidence that differences in inflation rate between tokens can explain differences in token returns within this sample. The insignificance of *TOKINC* across the full sample and the subsets is consistent with the lack of significant of *EARN* across samples - despite the intuitive appeal of this formulation for token "earnings", there is no evidence that modelling token issuance as a negative cashflow that offsets fee income helps to further explain the cross-section of token returns.

13.3 Fama-MacBeth

To further differentiate between the revenue measures, I run Fama-MacBeth regressions (Fama and MacBeth (1997)) to analyse whether any are subsumed for the tokens which have all three metrics available. I follow the same methodology in Liu, Tsyvinski, and Xi (2022) in sorting tokens by each characteristic into three portfolios each day and using the portfolio rank numbers as the explanatory variable.

Table 19 reports the Fama-MacBeth OLS regressions and shows that whilst all three variables are significant on a univariate basis, when run in the same model, only *PREV* is significant at the 5% level, with *FEES* adopting a negative sign. These results are consistent with protocol revenue subsuming the other measures. It is noted that the standard Fama-MacBeth OLS considers each observation equally, consistent with the focus on equally weighted portfolios in this study as discussed in section 13.

13.4 Spanning Regressions

Following on from secton 13.3, I test the extent to which *PREV* is spanned by other cryptocurrency factors in the literature. For the crypto Market, Size, and Momentum factors (CMKT, CSMB, and CMOM from Liu, Tsyvinski, and Xi (2022)), I use the authors data from their website which provides factor returns for the period up to the end of 2021. I reconstruct their factors using CoinGecko data to update the series to the end of 2022. To keep the analysis consistent with their returns series, which is only available for a weekly frequency, I follow the same methodology they use in aggregating the daily returns into the 52 weeks (7 days for the first 51 weeks, with the remaining days included in the final week).

For the spanning regressions I include all the fundamental factors discussed in section 11 with the exception of *FEES* and *EARN* since they are combinations of the others. In addition, I include factor returns from portfolio sorts on the negative of the past 365-day returns, a mean of long term reversal (R365d), which together with *ADDNZ* covers the two factors documented in the C-5 model of Cong et al. (2022).

Table 20 shows the results when running the regressions for the crypto three-factor model in Liu, Tsyvinski, and Xi (2022) with R365d and *PREV*. Each row is a regression of the factor on all the other factors between December 2019 and December 2022. The spanning regressions show statistically significant intercepts for CMKT, CSMB, and *PREV*, consistent with the notion that these factors are distinct from the others. Whilst Cong et al. (2022) use reversal as a VALUE factor in their C-5 model, the methodology used in Asness, Moskowitz, and Pedersen (2013) for commodity and currency markets, I show evidence that *PREV*, which is based on fundamental characteristics more akin to the value factors used in equity factor models, is distinct from reversal and in fact has a negative (though not statistically significant) loading on long term reversal. *PREV* also has no statistically significant loading on any of the three factors documented in Liu, Tsyvinski, and Xi (2022). Notably the intercept is not statistically significant for CMOM over this period, which is also the case when running the the same test on only CMKT, CSMB, and CMOM.

For completeness I also include table 21 which shows the results of regressing of each factor on all the other factors. Intercepts for CSMB and *PREV* continue to be statistically significant with the inclusion of all the factors, however the intercept for CMKT is no longer significant whilst the *TOKINC* is negative and not spanned. Other observations include intuitive negative loadings for CMKT on CSMB and for CMOM on R365d.

13.5 GRS

Tables o Summary Stats – Mean, SD, Skew, Kurt, Min, Percentiles o Correlation between different measurement methodologies used o Persistence – Correlation of metric using non-overlapping data e.g. time series average of cross sectional correlations o Portfolio Analysis – As per Empirical Methods Dependent and Independent Sorts. Dependent sort on first X1 then X2 gives you relationship between X2 and Y conditional on X1. E.g. X2 sort happens withing the X1 slices. EW and VW NYSE and CRSP (NYSE/AMEX/NASDAQ) breakpoints o Fama MacBeth Newey West Standard Errors Simultaneously control for many effects o GRS tests whether the Sharpe of a portfolio is spanned by Sharpe of factor model (simultaneously tests if alphas of each asset in the portfolio are jointly equal to 0 F test shows whether statistical significance between RSS of an unconstrained portfolio Vs a constrained portfolio Constrained RSS always higher than unconstrained RSS so ratio of two always positive thus lognormal distribution and use F distribution to test Wald stat is (Su^2-Sr^2)/(1+Sr^2), for Sharpe of unrestricted (assets+factors) and restricted (just factors) Ftest is T(T-N-1)/N(T-2)*Wald for N factors and T observations of factors

14 Conclusion

Value and momentum are asset pricing anomalies which are pervasive across traditional asset classes such as equities, bonds, currencies, and commodities (Asness, Moskowitz, and Pedersen (2013)). Though the presence of momentum has been well documented in the crypto asset pricing literature, consensus around a value factor within crypto has proved elusive, owing to the difficulty of pinning down cashflows for cryptocurrencies. As such, some studies favour long term reversal measures to construct a crypto value factor (Cong et al. (2022)), more consistent with a view of crypto tokens as commodities or currencies.

This paper identifies cashflows within the universe of cryptocurrency tokens derived from on-chain data. Using crypto native fundamental characteristics inferred from these cashflows, I document a value factor in the cross-section of token returns analogous to valuation characteristics used in the equity market. This cashflow based value factor is not spanned by prevailing crypto factor models and is distinct from long term reversal, highlighting the hybrid nature of crypto tokens as displaying features from a breadth of traditional asset classes. I find this factor is more applicable in tokens issued by decentralised applications than for native tokens of Layer-1 blockchains, consistent with the intuition that tokens of Dapps more closely resemble equities, whereas native Layer-1 tokens have properties more akin to currencies. However, I find no evidence that on-chain token incentives are priced in the cross section and a further exploration of cost of revenue measures for crypto protocols is one of many potential extensions for this study. Finally, I test an additional set of crypto native metrics unique to public blockchains that proxy for capital gains overhang, insider ownership, and investor sophistication, though I find inconclusive results for these tokenholder characteristics to price the cross-section.

The cryptocurrency market, and the universe of decentralised applications in particular, is still nascent and rapidly evolving. The extent to which fundamental characteristics such as protocol revenue will continue to be priced as the market matures will be an interesting area for further research. Given the value premium in crypto appears significantly larger than analogous asset pricing anomalies in traditional asset classes (Liu, Tsyvinski, and Xi (2022) make a similar observation for their three crypto factors), the expectation would be an attenuation over time. These caveats notwithstanding, this study provides a clean out of sample test for the value factor within a novel asset class.

Figure 2: Cumulative on-chain fees earned from June 2017 to December 2022, broken down by revenues accruing to protocols and to supply side users.

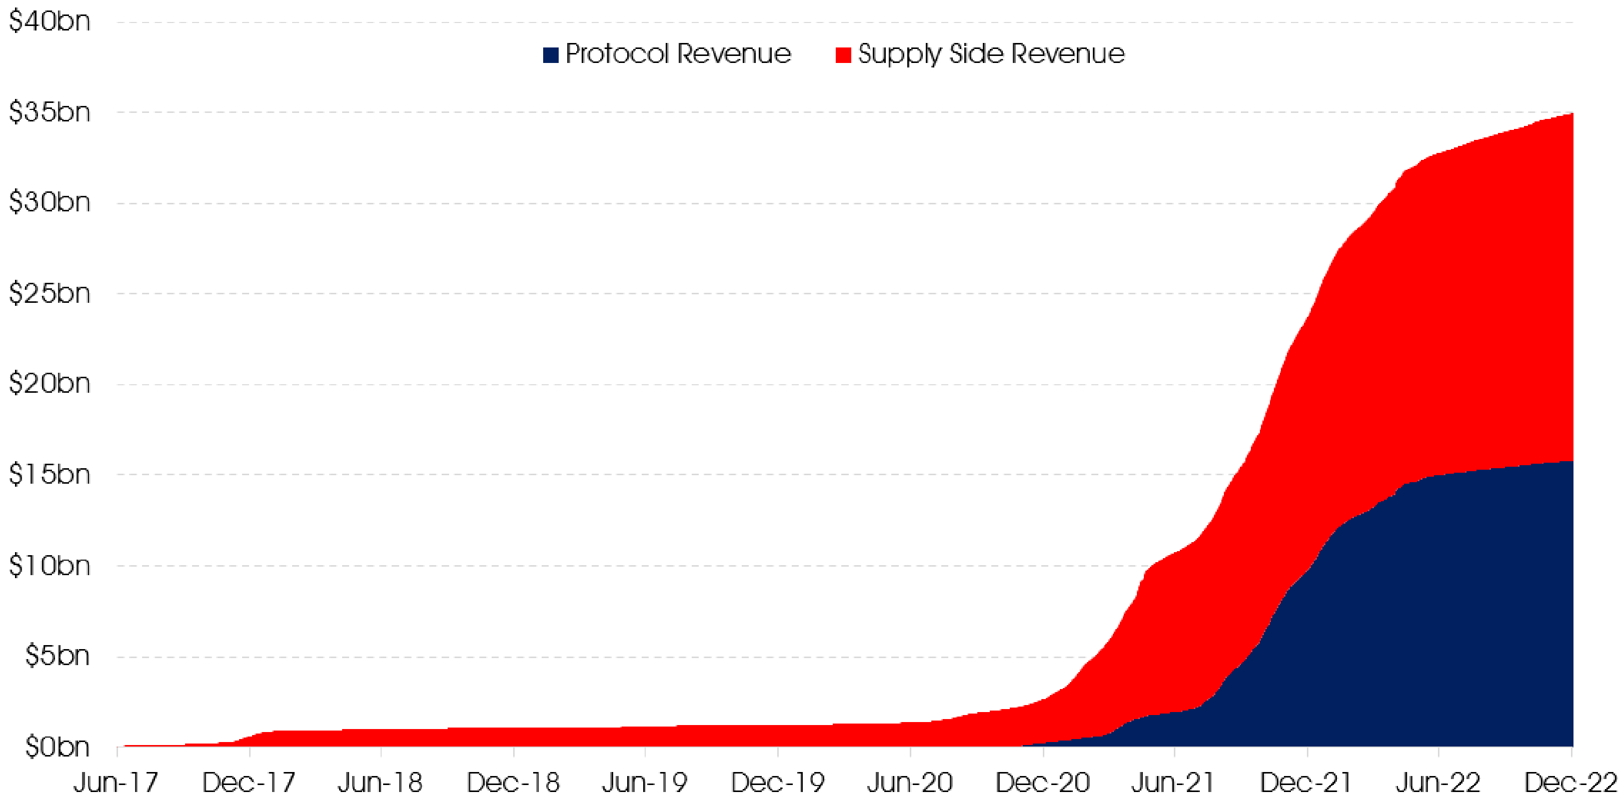

Figure 3: Cross-sectional distribution of average take rates (% of total fees captured by protocol versus supply side user) during 2022 by protocol

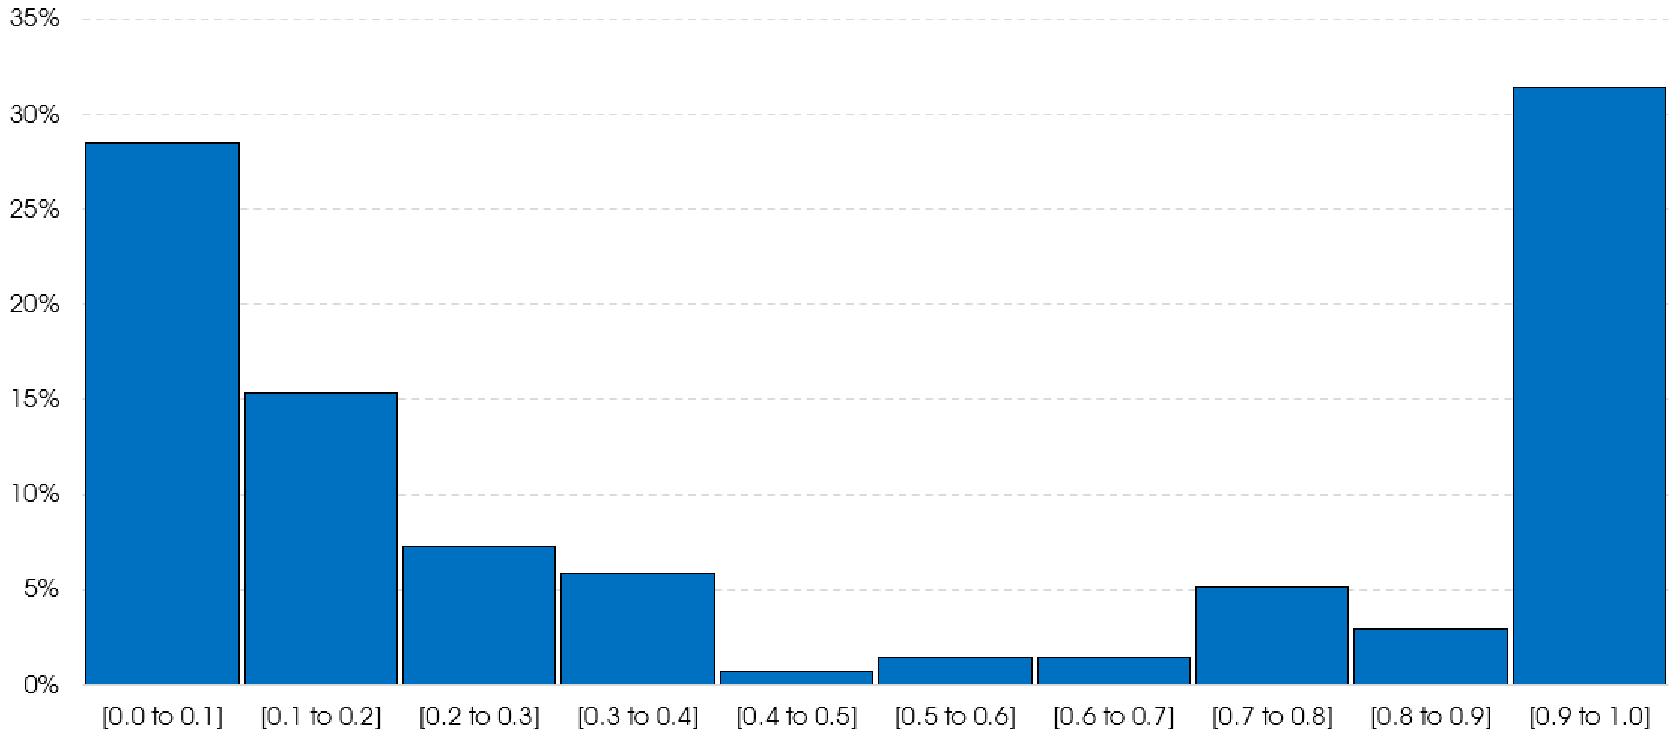

Figure 4: Boxplots of token market capitalisation divided by protocol revenue (blue), supply side revenue (orange), and total fees (grey) as at 31st December 2022, compared to price-to-sales ratios of listed equities in the MSCI All-Country World (yellow) on the same date.

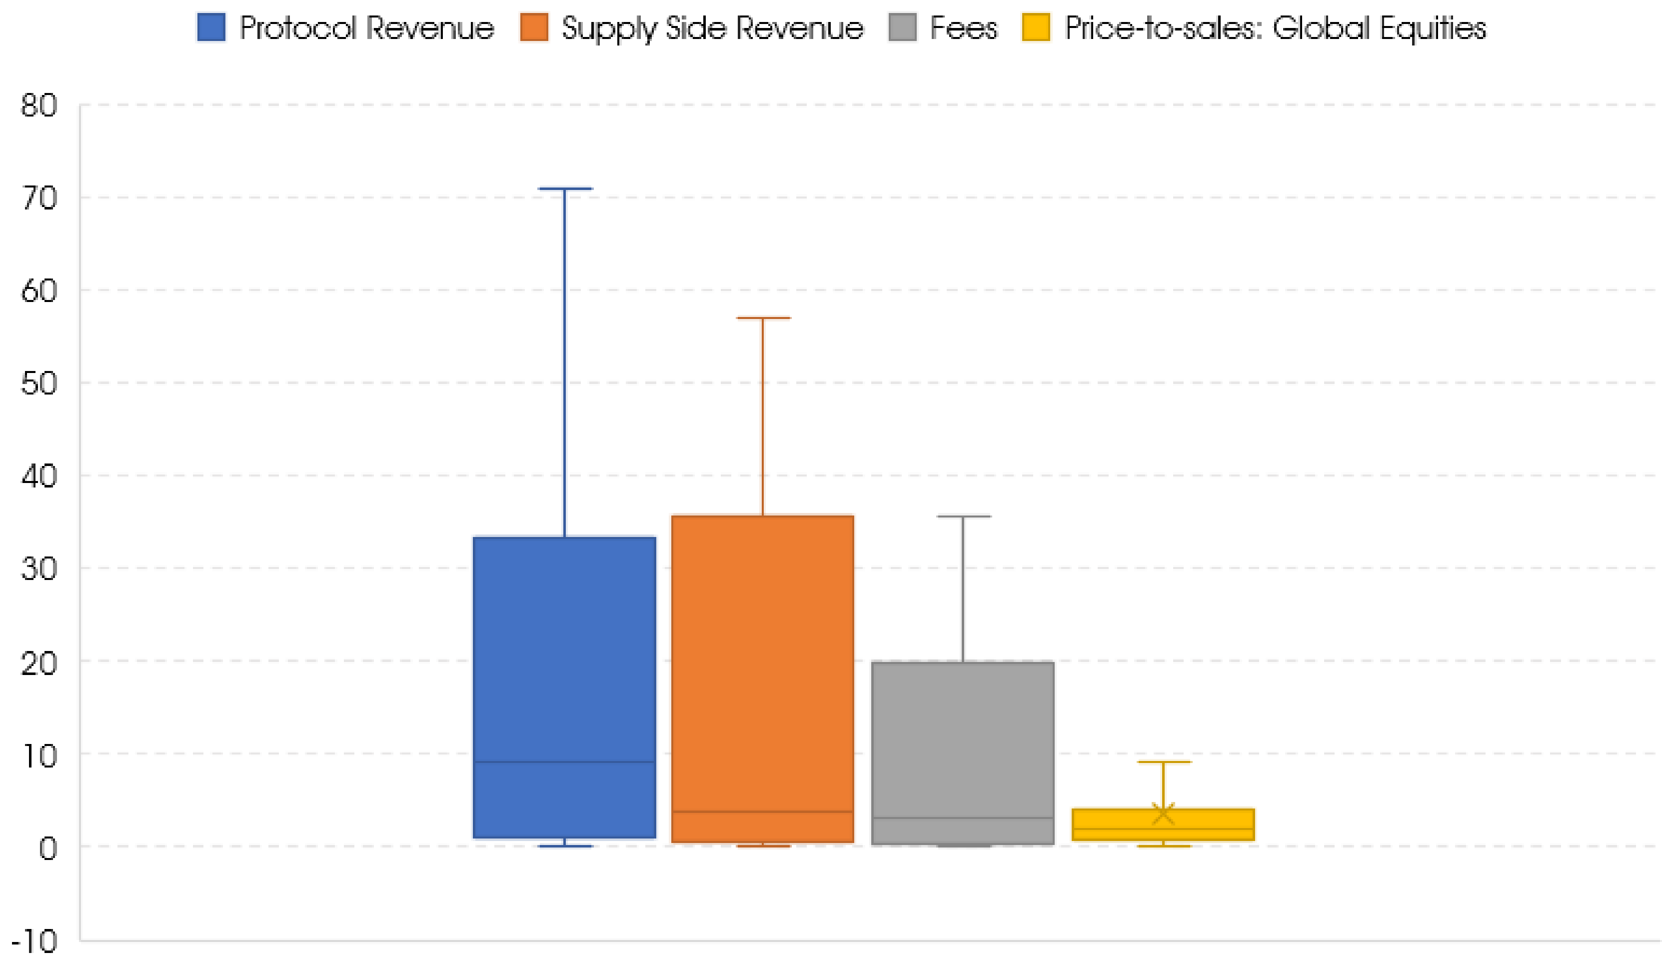

Figure 5: Cumulative token incentives paid by protocols (including block rewards paid to miners and validators) between June 2017 and December 2022.

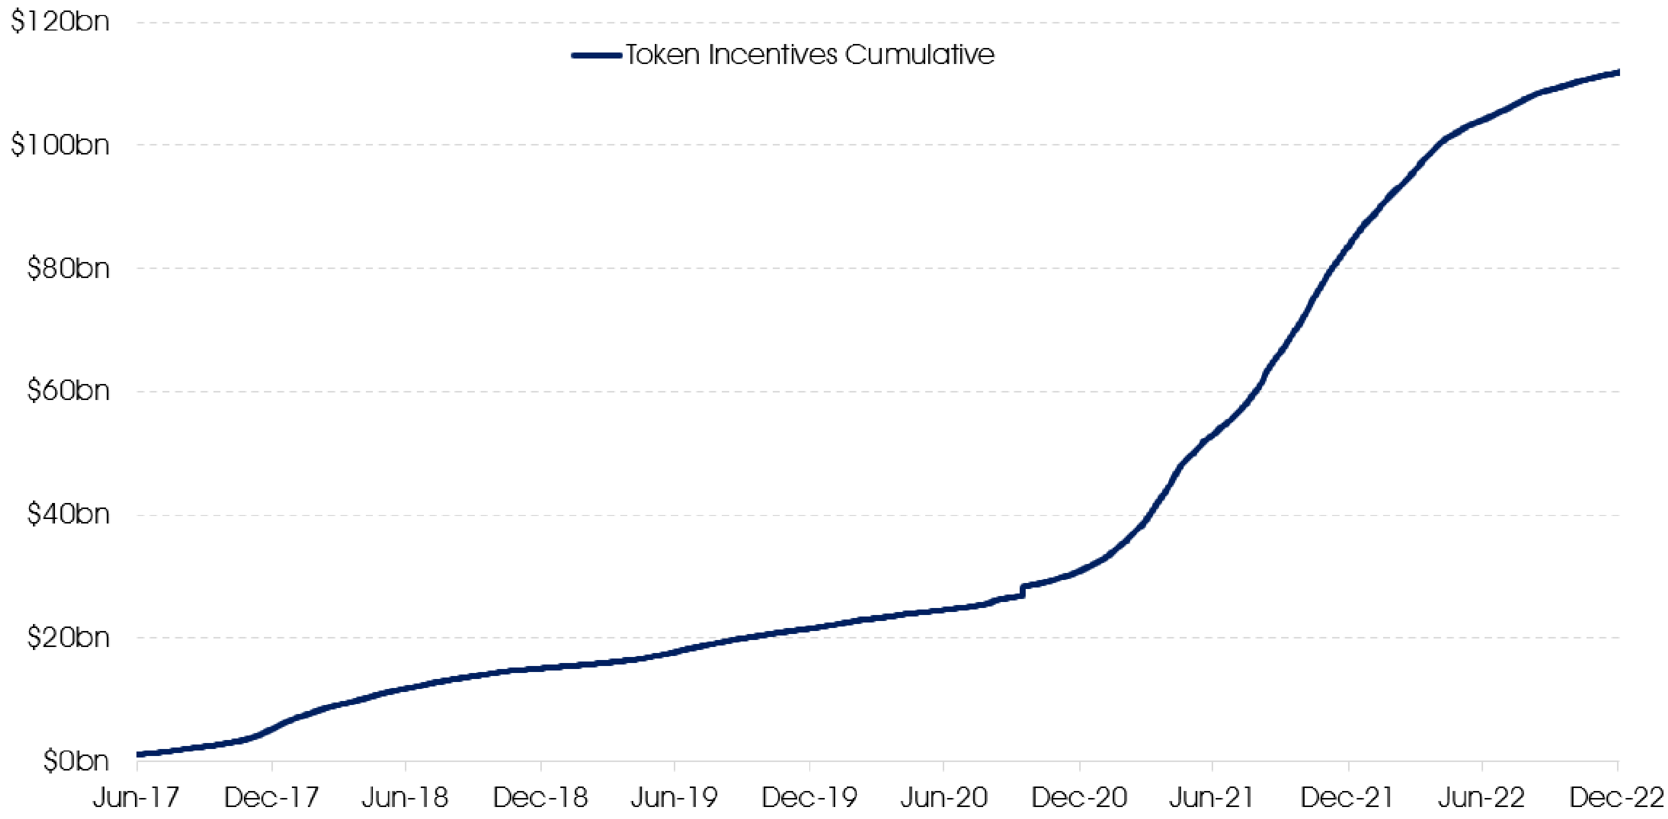

Figure 6: Boxplots of token incentives divided by market cap (blue) over calendar year 2022, compared to country inflation rates (orange) and changes in shares outstanding for listed equities in the MSCI All-Country World (yellow) over the same period.

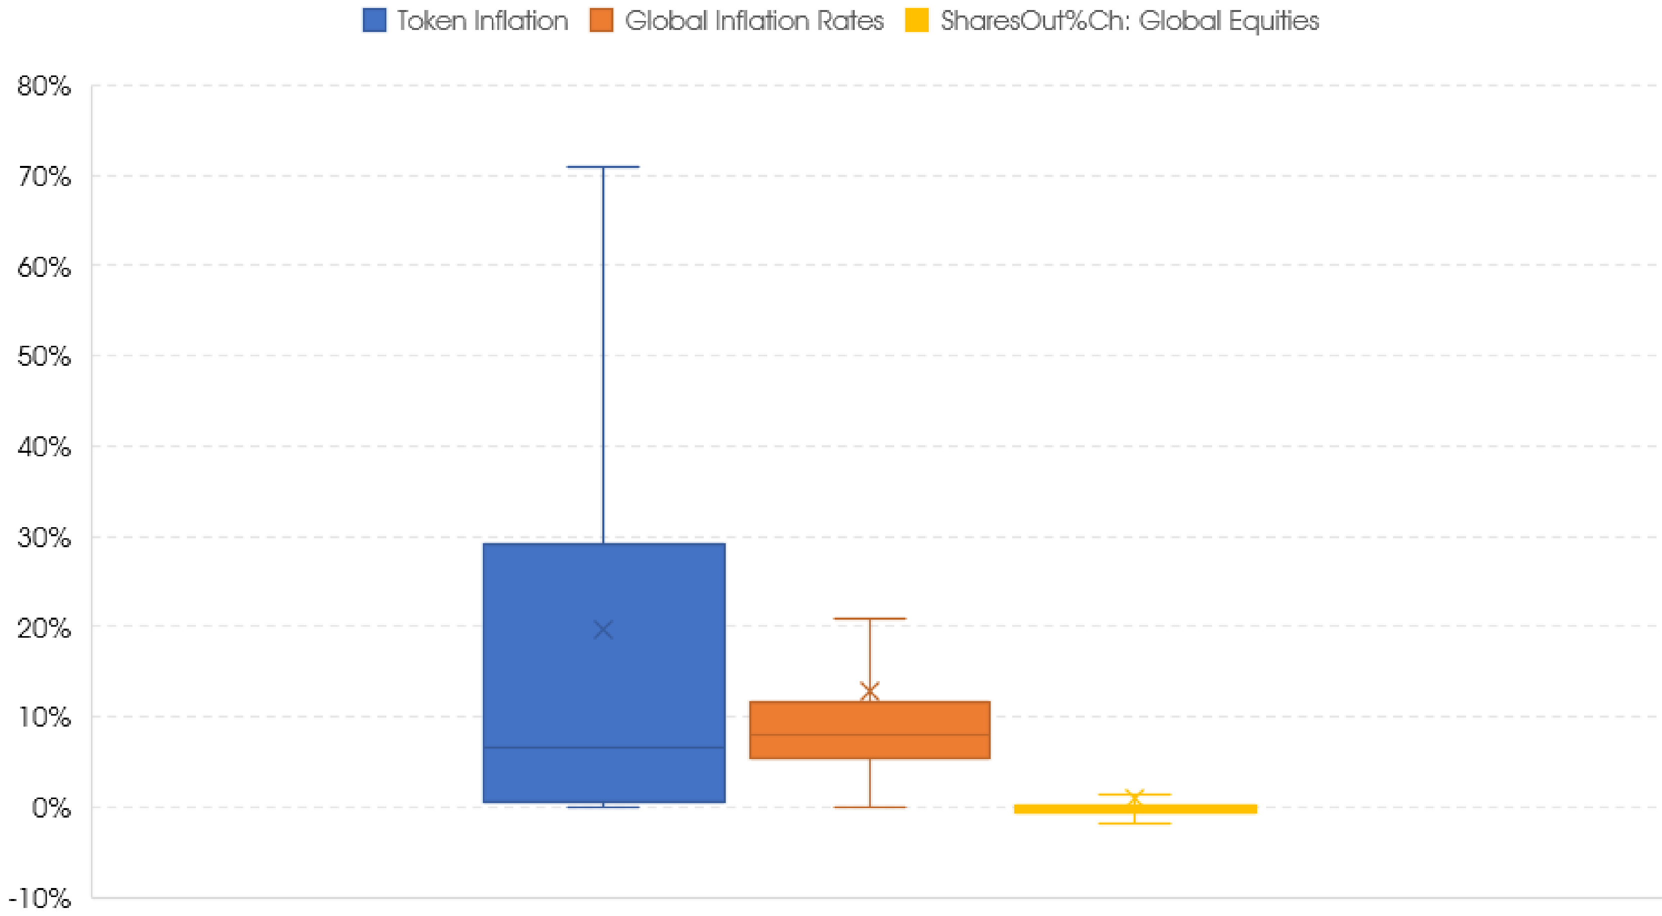

Figure 7: Total market capitalisation of crypto asset class broken down by BTC, ETH, and all other tokens, from June 2018 to December 2022.

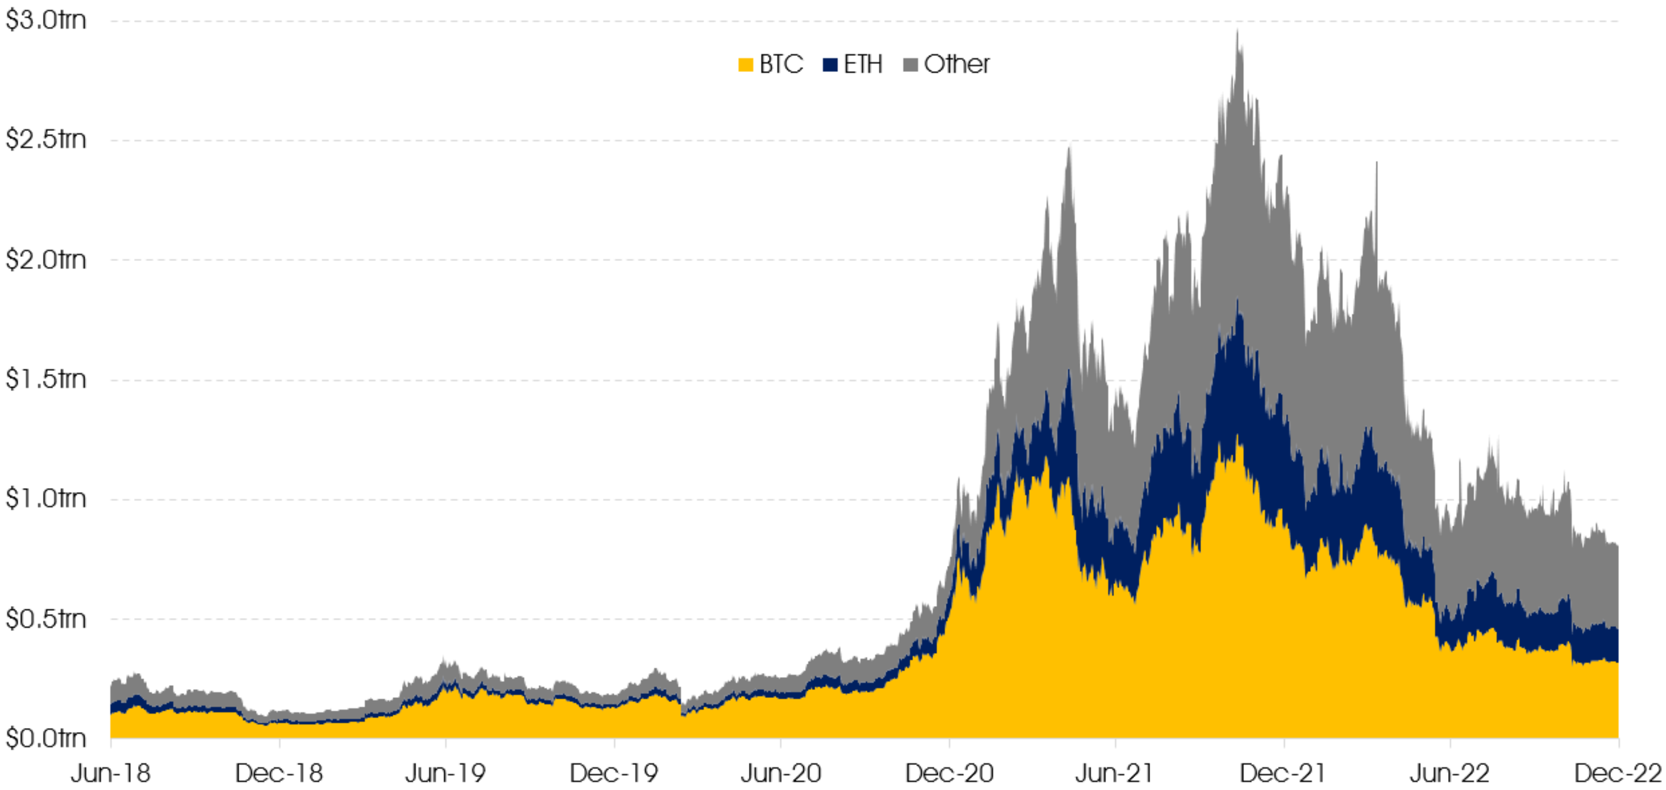

Figure 8: Median market capitalisation of Layer-1 tokens (red) and Dapp tokens (blue) in data sample, from December 2019 to December 2022

Table 12: Cross-sectional correlation of crypto fundamental characteristics across tokens using average values between December 2019 and December 2022. Characteristics include: Revenue retained by the protocol ($PREV$), Revenue earned by supply side users ($SREV$), Total fees paid to protocols ($FEES$), Token incentives distributed ($TOKINC$), Market value to realised value ($MVRV$), Number of addresses with a non-zero token balance ($ADDNZ$), Proportion of addresses that hold more than %1 of token supply (WHA), Proportion of addresses which hold more than \$1m in the token (BAL)

	PREV	SREV	TOKINC	FEES	MVRV	ADDNZ	WHA	BAL	Market Cap
PREV	1.0000	0.7715	0.9570	0.6440	-0.0504	0.7831	-0.2572	-0.0129	0.3003
SREV	0.7715	1.0000	0.9229	0.6218	-0.0077	0.8224	-0.3985	0.0195	0.4772
TOKINC	0.9570	0.9229	1.0000	0.6679	-0.0338	0.8386	-0.3328	0.0009	0.3993
FEES	0.6440	0.6218	0.6679	1.0000	-0.0579	0.9297	-0.5936	-0.0646	0.8584
MVRV	-0.0504	-0.0077	-0.0338	-0.0579	1.0000	-0.0557	0.1308	0.1452	-0.0287
ADDNZ	0.7831	0.8224	0.8386	0.9297	-0.0557	1.0000	-0.5481	-0.0437	0.7498
WHA	-0.2572	-0.3985	-0.3328	-0.5936	0.1308	-0.5481	1.0000	0.2759	-0.5104
BAL	-0.0129	0.0195	0.0009	-0.0646	0.1452	-0.0437	0.2759	1.0000	-0.0154
Market Cap	0.3003	0.4772	0.3993	0.8584	-0.0287	0.7498	-0.5104	-0.0154	1.0000

Table 13: Regression of t days lagged fundamental characteristics on current value, across token sample 2019-2022. Characteristics include: Revenue retained by the protocol (*PREV*), Revenue earned by supply side users (*SREV*), Total fees paid to protocols (*FEES*), Token incentives distributed (*TOKINC*), Earnings (*EARN*) defined as *FEES* - *TOKINC*, Market value to realised value (*MVRV*), Number of addresses with a non-zero token balance (*ADDNZ*), Proportion of addresses that hold more than %1 of token supply (*WHA*), Proportion of addresses which hold more than $\$1m$ in the token (*BAL*)

		1	7	30	60	180
PREV	Beta	0.886	0.838	0.721	0.648	0.229
	Tstat	586.528	468.685	311.102	248.880	61.737
	r2	0.784	0.701	0.515	0.415	0.050
SREV	Beta	0.889	0.752	0.595	0.509	0.302
	Tstat	597.274	349.498	223.174	174.811	84.703
	r2	0.790	0.566	0.354	0.259	0.089
FEES	Beta	0.886	0.818	0.703	0.632	0.453
	Tstat	587.313	434.954	296.551	239.586	134.189
	r2	0.785	0.668	0.491	0.396	0.198
TOKINC	Beta	0.952	0.932	0.898	0.836	0.744
	Tstat	264.424	249.716	228.628	197.627	147.012
	r2	0.501	0.474	0.439	0.380	0.293
EARN	Beta	0.877	0.837	0.800	0.743	0.639
	Tstat	223.532	206.325	189.698	166.992	123.793
	r2	0.397	0.361	0.329	0.284	0.206
MVRV	Beta	1.068	0.999	1.769	1.873	1.641
	Tstat	721.694	230.411	196.841	138.208	69.194
	r2	0.828	0.330	0.268	0.157	0.050
ADDNZ	Beta	0.999	0.995	0.979	0.957	0.872
	Tstat	92094.422	37756.516	12407.637	7557.330	3325.003
	r2	1.000	1.000	0.999	0.997	0.988
WHA	Beta	1.000	1.000	1.001	0.999	0.993
	Tstat	12049.288	4658.710	2408.544	1695.661	958.414
	r2	0.999	0.996	0.985	0.971	0.923
BAL	Beta	0.999	0.995	0.977	0.952	0.873
	Tstat	6568.998	2742.011	1330.397	924.733	513.117
	r2	0.997	0.982	0.928	0.864	0.685

Table 14: Average returns and t-statistics for equally-weighted token portfolios sorted daily by fundamental characteristics constructed from on-chain data between December 2019 and December 2022. Characteristics include: Revenue retained by the protocol (*PREV*), Revenue earned by supply side users (*SREV*), Total fees paid to protocols (*FEES*), Token incentives distributed (*TOKINC*), Earnings (*EARN*) defined as *FEES* - *TOKINC*, Market value to realised value (*MVRV*), Number of addresses with a non-zero token balance (*ADDNZ*), Proportion of addresses that hold more than %1 of token supply (*WHA*), Proportion of addresses which hold more than \$1m in the token (*BAL*).

	Lo	Med	Hi	Hi-Lo
PREV	0.003*	0.003*	0.007***	0.004***
	(1.81)	(1.838)	(3.627)	(3.017)
SREV	0.003**	0.004**	0.006***	0.003***
	(2.121)	(2.177)	(3.564)	(2.699)
FEES	0.003*	0.003**	0.006***	0.003***
	(1.94)	(2.105)	(3.644)	(3.223)
TOKINC	0.004**	0.003	0.005**	0.001
	(2.355)	(1.607)	(2.53)	(0.87)
EARN	0.005***	0.003**	0.004**	-0.001
	(2.682)	(2.036)	(2.505)	(-0.582)
MVRV	0.005***	0.004**	0.004**	-0.001
	(2.899)	(2.433)	(2.314)	(-1.256)
ADDNZ	0.003**	0.004***	0.008***	0.005***
	(2.099)	(2.787)	(4.177)	(3.718)
WHA	0.004**	0.004**	0.007***	0.003*
	(2.432)	(2.533)	(3.191)	(1.957)
BAL	0.007***	0.004***	0.004**	-0.003***
	(3.736)	(2.732)	(2.339)	(-2.597)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 15: Average returns and t-statistics for 2x3 independent portfolio sorts by size and revenue based fundamental characteristics constructed from on-chain data between December 2019 and December 2022. Characteristics include: Revenue retained by the protocol (*PREV*), Revenue earned by supply side users (*SREV*), Total fees paid to protocols (*FEES*), Token incentives distributed (*TOKINC*), Earnings (*EARN*) defined as *FEES* - *TOKINC*.

		Lo	Med	Hi	Hi-Lo
PREV	Big	0.003*	0.003	0.006***	0.004**
		(1.693)	(1.605)	(3.083)	(2.536)
	Small	0.003	0.004**	0.008***	0.005***
		(1.392)	(2.074)	(3.679)	(2.935)
SREV	Big	0.003	0.003**	0.002	-0.0
		(1.612)	(1.98)	(1.247)	(-0.322)
	Small	0.004**	0.004**	0.008***	0.004**
		(2.111)	(1.991)	(4.043)	(2.227)
FEES	Big	0.003*	0.003*	0.003	0.0
		(1.653)	(1.702)	(1.489)	(0.09)
	Small	0.003*	0.004**	0.008***	0.004***
		(1.918)	(2.149)	(4.184)	(2.929)
TOKINC	Big	0.003**	0.002	0.001	-0.003**
		(2.28)	(1.306)	(0.375)	(-2.26)
	Small	0.005***	0.003*	0.005***	0.0
		(2.653)	(1.722)	(2.653)	(0.046)
EARN	Big	0.002	0.003*	0.003*	0.001
		(0.898)	(1.903)	(1.818)	(1.169)
	Small	0.006***	0.004*	0.005***	-0.001
		(3.014)	(1.9)	(2.761)	(-0.489)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 16: Average returns and t-statistics for 2x3 independent portfolio sorts by size and tokenholder based fundamental characteristics constructed from on-chain data between December 2019 and December 2022. Characteristics include: Market value to realised value ($MVRV$), Number of addresses with a non-zero token balance ($ADDNZ$), Proportion of addresses that hold more than %1 of token supply (WHA), Proportion of addresses which hold more than \$1m in the token (BAL).

		Lo	Med	Hi	Hi-Lo
MVRV	Big	0.003*	0.003*	0.003*	-0.0
		(1.847)	(1.916)	(1.896)	(-0.186)
	Small	0.006***	0.005***	0.004**	-0.002
		(3.36)	(2.842)	(2.381)	(-1.336)
ADDNZ	Big	0.002*	0.003*	0.003	0.0
		(1.664)	(1.914)	(1.607)	(0.166)
	Small	0.006***	0.006***	0.01***	0.005**
		(2.747)	(3.258)	(4.411)	(2.064)
WHA	Big	0.003*	0.003*	0.004*	0.001
		(1.8)	(1.788)	(1.654)	(0.653)
	Small	0.006***	0.005***	0.007***	0.001
		(2.935)	(2.723)	(3.411)	(0.556)
BAL	Big	0.002*	0.003**	0.002	-0.001
		(1.65)	(2.134)	(1.202)	(-1.23)
	Small	0.01***	0.005***	0.006***	-0.004
		(3.964)	(3.052)	(3.084)	(-1.623)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 17: Average returns and t-statistics for equally-weighted token portfolios for the **decentralised application** subset, sorted daily by fundamental characteristics constructed from on-chain data between December 2020 and December 2022. Characteristics include: Revenue retained by the protocol (*PREV*), Revenue earned by supply side users (*SREV*), Total fees paid to protocols (*FEES*), Token incentives distributed (*TOKINC*), Earnings (*EARN*) defined as *FEES* - *TOKINC*, Market value to realised value (*MVRV*), Number of addresses with a non-zero token balance (*ADDNZ*), Proportion of addresses that hold more than %1 of token supply (*WHA*), Proportion of addresses which hold more than $\$1m$ in the token (*BAL*).

	Lo	Med	Hi	Hi-Lo
PREV	0.002	0.001	0.006***	0.004***
	(0.953)	(0.398)	(2.827)	(3.283)
SREV	0.003	0.002	0.006***	0.004***
	(1.266)	(0.799)	(2.626)	(2.582)
FEES	0.002	0.001	0.006***	0.004***
	(1.02)	(0.494)	(2.807)	(3.53)
TOKINC	0.004*	0.002	0.005**	0.002
	(1.708)	(0.732)	(2.206)	(1.07)
EARN	0.004*	0.003	0.003	-0.001
	(1.792)	(1.35)	(1.575)	(-0.584)
MVRV	0.004*	0.002	0.002	-0.001
	(1.776)	(1.265)	(1.076)	(-1.572)
ADDNZ	0.002	0.004*	0.008***	0.006***
	(0.849)	(1.888)	(2.931)	(3.252)
WHA	0.002	0.003	0.003	0.001
	(1.197)	(1.449)	(1.464)	(0.604)
BAL	0.007***	0.003	0.003	-0.004*
	(2.692)	(1.477)	(1.606)	(-1.958)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 18: Average returns and t-statistics for equally-weighted token portfolios for the **Layer-1** native tokens subset, sorted daily by fundamental characteristics constructed from on-chain data between December 2020 and December 2022. Characteristics include: Revenue retained by the protocol (*PREV*), Revenue earned by supply side users (*SREV*), Total fees paid to protocols (*FEES*), Token incentives distributed (*TOKINC*), Earnings (*EARN*) defined as *FEES* - *TOKINC*, Market value to realised value (*MVRV*), Number of addresses with a non-zero token balance (*ADDNZ*), Proportion of addresses that hold more than %1 of token supply (*WHA*), Proportion of addresses which hold more than $\$1m$ in the token (*BAL*).

	Lo	Med	Hi	Hi-Lo
PREV	0.003*	0.003	0.004**	0.001
	(1.678)	(1.254)	(1.989)	(0.637)
SREV	0.003	0.002	0.006**	0.002
	(1.524)	(1.05)	(2.273)	(1.362)
FEES	0.003	0.002	0.005**	0.003**
	(1.276)	(1.096)	(2.439)	(2.045)
TOKINC	0.004*	0.004	0.003	-0.0
	(1.825)	(1.533)	(1.537)	(-0.328)
EARN	0.003	0.003	0.005**	0.002
	(1.386)	(1.251)	(2.088)	(1.278)
MVRV	0.005**	0.003	0.003	-0.002
	(2.36)	(1.56)	(1.437)	(-1.559)
ADDNZ	0.002	0.004**	0.011***	0.009***
	(1.365)	(2.194)	(2.96)	(2.603)
WHA	0.001	0.007**	0.004	0.003
	(0.553)	(2.415)	(1.51)	(1.64)
BAL	0.009***	0.004**	0.005**	-0.005
	(2.635)	(1.971)	(2.142)	(-1.369)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 19: Fama-MacBeth regression results for fundamental revenue characteristics constructed from on-chain data between December 2019 and December 2022. For each characteristic, tokens are sorted into three portfolios each day, with the portfolio rank number used as the explanatory variable. Characteristics include: Revenue retained by the protocol ($PREV$), Revenue earned by supply side users ($SREV$), Total fees paid to protocols ($FEES$)

PREV	0.0018***			0.0025**
	(2.903)			(2.153)
SREV		0.0011**		0.0022
		(2.011)		(1.443)
FEES			0.0012**	-0.0030*
			(2.459)	(-1.763)

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 20: Spanning regressions across crypto factors from December 2019 to December 2022. Each row is a regression of the factor on all the other factors. Crypto factors include Market (CMKT), Size (CSMB), Momentum (CMOM), LT Reversal (R365d), and Revenue retained by the protocol ($PREV$).

	Intercept	CMKT	CSMB	CMOM	R365d	Protocol Revenue	R2
CMKT	0.019**		-0.357***	0.115	0.136	0.045	0.039
	(2.201)		(-3.095)	(0.837)	(0.818)	(0.618)	
CSMB	0.026***	-0.168***		0.095	0.184	0.056	0.055
	(4.517)	(-3.095)		(1.013)	(1.632)	(1.12)	
CMOM	0.006	0.04	0.071		-0.185*	-0.037	0.006
	(1.132)	(0.837)	(1.013)		(-1.896)	(-0.867)	
R365d	0.002	0.033	0.095	-0.127*		-0.051	0.022
	(0.483)	(0.818)	(1.632)	(-1.896)		(-1.43)	
PREV	0.023**	0.056	0.148	-0.133	-0.264		-0.004
	(2.38)	(0.618)	(1.12)	(-0.867)	(-1.43)		

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 21: Spanning regressions across crypto factors from December 2019 to December 2022. Each row is a regression of the factor on all the other factors. Crypto factors include Market (CMKT), Size (CSMB), Momentum (CMOM), LT Reversal (R365d), and factors constructed from on-chain fundamental characteristics, including: Revenue retained by the protocol ($PREV$), Revenue earned by supply side users ($SREV$), Token incentives distributed ($TOKINC$), Market value to realised value ($MVRV$), Number of addresses with a non-zero token balance ($ADDNZ$), Proportion of addresses that hold more than %1 of token supply (WHA), Proportion of addresses which hold more than $\$1m$ in the token (BAL).

	Intercept	CMKT	CSMB	CMOM	R365d	PREV	TOKINC	MVRV	ADDNZ	WHA	BAL	R2
CMKT	0.012		-0.466***	0.053	0.096	0.024	0.037	-0.209*	-0.013	0.065	-0.318***	0.156
	(1.398)		(-4.152)	(0.397)	(0.601)	(0.297)	(0.477)	(-1.878)	(-0.165)	(1.092)	(-3.21)	
CSMB	0.019***	-0.228***		0.071	0.115	0.019	0.026	-0.103	0.063	0.086**	-0.091	0.137
	(3.317)	(-4.152)		(0.77)	(1.024)	(0.333)	(0.479)	(-1.317)	(1.153)	(2.096)	(-1.278)	
CMOM	0.007	0.021	0.057		-0.18*	-0.067	0.041	0.122*	0.001	-0.032	-0.089	0.031
	(1.309)	(0.397)	(0.77)		(-1.821)	(-1.354)	(0.861)	(1.757)	(0.021)	(-0.873)	(-1.393)	
R365d	0.002	0.026	0.063	-0.124*		-0.064	0.051	-0.02	0.094**	-0.015	0.082	0.048
	(0.43)	(0.601)	(1.024)	(-1.821)		(-1.545)	(1.286)	(-0.349)	(2.352)	(-0.491)	(1.554)	
PREV	0.023***	0.026	0.041	-0.185	-0.255		0.363***	0.481***	-0.117	0.112*	-0.178*	0.243
	(2.603)	(0.297)	(0.333)	(-1.354)	(-1.545)		(4.926)	(4.397)	(-1.453)	(1.843)	(-1.691)	
TOKINC	-0.018**	0.043	0.061	0.123	0.222	0.395***		-0.342***	0.19**	0.214***	0.205*	0.286
	(-1.965)	(0.477)	(0.479)	(0.861)	(1.286)	(4.926)		(-2.887)	(2.281)	(3.454)	(1.875)	
MVRV	-0.01	-0.114*	-0.115	0.171*	-0.041	0.244***	-0.159***		0.134**	-0.023	0.105	0.142
	(-1.604)	(-1.878)	(-1.317)	(1.757)	(-0.349)	(4.397)	(-2.887)		(2.358)	(-0.516)	(1.401)	
ADDNZ	0.01	-0.015	0.144	0.003	0.393**	-0.123	0.182**	0.276**		0.098	-0.909***	0.54
	(1.057)	(-0.165)	(1.153)	(0.021)	(2.352)	(-1.453)	(2.281)	(2.358)		(1.565)	(-11.627)	
WHA	0.022*	0.126	0.341**	-0.161	-0.11	0.203*	0.355***	-0.081	0.169		0.279**	0.19
	(1.807)	(1.092)	(2.096)	(-0.873)	(-0.491)	(1.843)	(3.454)	(-0.516)	(1.565)		(1.983)	
BAL	0.001	-0.209***	-0.122	-0.149	0.2	-0.109*	0.115*	0.127	-0.531***	0.095**		0.548
	(0.125)	(-3.21)	(-1.278)	(-1.393)	(1.554)	(-1.691)	(1.875)	(1.401)	(-11.627)	(1.983)		

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

Table 22: Spanning test across LTW factors and fundamental revenue characteristics constructed from on-chain data between 2020 to 2022

	const	CMKT	CSMB	CMOM	PREV	SREV	FEES	R2
CMKT	0.022**		-0.282**	-0.09	-0.072	0.063	0.088	0.013
	(2.24)		(-2.151)	(-0.604)	(-0.667)	(0.436)	(0.503)	
CSMB	0.027***	-0.12**		0.045	-0.024	-0.147	0.233**	0.034
	(4.383)	(-2.151)		(0.457)	(-0.338)	(-1.579)	(2.067)	
CMOM	0.014**	-0.031	0.035		-0.046	0.03	-0.015	-0.026
	(2.462)	(-0.604)	(0.457)		(-0.73)	(0.362)	(-0.148)	
PREV	0.011	-0.046	-0.036	-0.087		-0.003	0.754***	0.491
	(1.352)	(-0.667)	(-0.338)	(-0.73)		(-0.025)	(6.077)	
SREV	0.001	0.023	-0.126	0.033	-0.002		0.921***	0.72
	(0.154)	(0.436)	(-1.579)	(0.362)	(-0.025)		(13.22)	
FEES	0.0	0.022	0.135**	-0.011	0.29***	0.619***		0.784
	(0.041)	(0.503)	(2.067)	(-0.148)	(6.077)	(13.22)		

***,**, * - significant at the 1%, 5%, and 10% level, respectively.

References

Athey, S., V. Parashkevov, V. Sarukkai, and J. Xia (2016). "Bitcoin pricing, adoption, and usage: Theory and evidence". In: *Working Paper*.

Barberis, N., L. Jin, and B. Wang (2021). "Prospect Theory and Stock Market Anomalies". In: *Journal of Finance* 76, pp. 2639–2687.

Biais, B., C. Bisiere, M. Bouvard, and C. Casamatta (2023). "Equilibrium bitcoin pricing". In: *Journal of Finance* 78, pp. 967–1014.

Borri, N. (2019). "Conditional tail-risk in cryptocurrency markets". In: *Journal of Empirical Finance* 50, pp. 1–19.

Cai, F. and L. Zheng (2004). "Institutional Trading and Stock Returns". In: *Finance Research Letters* 1, pp. 178–189.

Cong, L. W., Z. He, and J. Li (2018). "Decentralized mining in centralized pools". In: *Working Paper*.

Cong, L. W., G. A. Karolyi, K. Tang, and W. Zaho (2022). "Value Premium, Network Adoption, and Factor Pricing of Crypto Assets". In: *Working Paper*.

Cong, L. W., Y. Li, and N. Wang (2021). "Tokenomics: Dynamic adoption and valuation". In: *Review of Financial Studies* 34, pp. 1105–1155.

Etherscan (n.d.). https://info.etherscan.com/address-tag-note/.

Fama, E. and K. French (1992). "The cross-section of expected stock returns". In: *Journal of Finance* 47, pp. 427–465.

Fama, E. and J. MacBeth (1997). "Risk, return, and equilibrium: Empirical tests". In: *Journal of Political Economy* 81, pp. 607–636.

Grinblatt, M. and B. Han (2005). "Prospect Theory, Metnal Accounting, and Momentum". In: *Journal of Financial Economics* 78, pp. 311–339.

Hu, A., C. A. Parlour, and U. Rajan (2019). "Cryptocurrencies: Stylized facts on a new investible instrument". In: *Financial Management* 48, pp. 1049–1068.

Liu, Y. and A. Tsyvinski (2021). "Risks and Returns of Cryptocurrency". In: *Review of Financial Studies* 34, pp. 2689–2727.

Liu, Y., A. Tsyvinski, and W. Xi (2022). "Common Risk Factors in Cryptocurrency". In: *Journal of Finance* 77, pp. 1133–1177.

Meiklejohn, S., M. Pomarole, G. Jordan, K. Levchenko, D. McCoy, G. Voelker, and S. Savage (2016). "A fistful of Bitcoins: characterizing payments among men with no names". In: *Communications of the Association for Computing Machinery* 59, pp. 86–93.

Monegro, J. (2016). *Fat Protocols*. https://www.usv.com/writing/2016/08/fat-protocols/.

Möser, M. and A. Narayanan (2022). "Resurrecting Address Clustering in Bitcoin". In: *Financial Cryptography and Data Security* 13411.

Nofsinger, J. and R. Sias (1998). "Herding and Feedback Trading by Institutional and Individual Investors". In: *Journal of Finance* 54, pp. 2263–2295.

Pagnotta, E. and A. Buraschi (2018). "An equilibrium valuation of bitcoin and decentralized network assets". In: *Working Paper*.

Sockin, M. and W. Xiong (2023). "A model of cryptocurrencies". In: *Management Science*.