BIS Working Papers
No 1087

Crypto carry

by Maik Schmeling, Andreas Schrimpf and
Karamfil Todorov

Monetary and Economic Department

April 2023 (revised October 2025)

JEL classification: G12, G13, G15

Keywords: crypto, carry, futures basis, bitcoin, ether

This publication is available on the BIS website (www.bis.org).

ISSN 1020-0959 (print)
ISSN 1682-7678 (online)

Crypto Carry*

Maik Schmeling
Goethe University & CEPR
schmeling@finance.uni-frankfurt.de

Andreas Schrimpf
BIS & CEPR
andreas.schrimpf@bis.org

Karamfil Todorov
BIS
karamfil.todorov@bis.org

This version: October 1, 2025

Abstract

We analyze the dynamics of carry in crypto markets – the difference between futures and spot prices – and document that it can reach exceptionally high levels, sometimes exceeding 40% per annum, with significant variation over time. This phenomenon reflects a substantial and volatile inconvenience yield associated with holding spot cryptocurrencies relative to futures. We trace the large and volatile crypto carry to the interplay of two main forces: (i) demand from smaller, trend-chasing investors seeking leveraged exposure, and (ii) the limited deployment of arbitrage capital due to regulatory and margin frictions. Our findings highlight how structural limits to arbitrage – especially severe in the case of crypto – can amplify price inefficiencies across financial markets, offering lessons for understanding asset pricing and market behavior more generally.

JEL Classification: G12, G13, G15
Keywords: Crypto, Carry, Futures basis, Bitcoin, Ether

*We are grateful for helpful comments to Jules van Binsbergen, Agostino Cappponi, Egemen Eren, Wenqian Huang, Albert Menkveld, Gabor Pinter, Ganesh Viswanath-Natra (discussant), Marina Niessner (discussant), Kenny Phua (discussant), Shihao Yu (discussant), as well as seminar participants at the Bank for International Settlements (BIS) and conference participants at the 15th Annual Conference on Advances in the Analysis of Hedge Fund Strategies at Imperial College London, the 18th Central Bank Market Microstructure Conference in Washington, D.C., Banque de France Digital Finance Research Conference, the 6th Sydney Market Microstructure and Digital Finance Meeting, and the SGF conference in Zurich. We also thank Giulio Cornelli for excellent research assistance. Maik Schmeling acknowledges financial support from the German Science Foundation (DFG). The views expressed in this paper are those by the authors and do not necessarily represent those by the BIS.

The ‘Risk-Free’ Crypto Trade Is Back In a Big Way [Bloomberg, 8 October 2021]

“The closest thing to a risk-free bet has reemerged in the cryptocurrency market as traders [...] bid up the price of futures.”

1. Introduction

Digital assets such as Bitcoin (BTC) and Ether (ETH) have grown rapidly in market capitalization and trading volume, effectively forming a new asset class. Cash and derivative crypto instruments are now actively traded both on crypto-native platforms and on traditional exchanges like the Chicago Mercantile Exchange (CME). Understanding these markets is thus crucial for decision makers in private and public organizations to identify opportunities and to mitigate risks in an increasingly digital and decentralized financial landscape. Despite the growth of crypto assets, regulatory frictions and uncertainty have long limited access to these assets for many professional investors. As a result, crypto markets remained segmented from traditional finance for years, especially during their early development.

This setting provides a unique opportunity to explore asset pricing and market efficiency during the emergence of a new asset class when institutional arbitrage capital was largely absent. In addressing these issues, we focus on one of the most prominent and persistent features of crypto markets: the significant gap between futures and spot prices, which occasionally exceeds 40% per annum (p.a.). This gap, often referred to as crypto carry or basis, is frequently portrayed in the media as a risk-free arbitrage opportunity.[1] We demonstrate empirically that crypto carry cannot be explained by fundamental factors alone, such as interest rate differentials between the crypto and traditional financial markets. Instead, we interpret the carry in crypto markets through the lens of the convenience yield, demonstrating how it is shaped by the interplay between buying pressure from smaller, trend-following investors and the limits to arbitrage caused by regulatory and other frictions. Additionally, we quantify the crypto convenience yield and provide causal evidence on the frictions that contribute to mispricing and impede the realization of ’risk-free’ arbitrage opportunities.

Although convenience yields matter in other asset classes as well, crypto markets serve as a particularly useful laboratory for studying them in a clean setting. Unlike futures in commodi-

[1] We use the terms “carry” and “basis” interchangeably to describe the difference between futures and spot prices.

ties, fixed income, or foreign exchange (FX), digital assets like bitcoin have negligible storage costs (unlike commodities), deterministic supply and no policy interventions (unlike FX or fixed income), and no industrial demand (unlike gold or other financial commodities). These features greatly simplify the identification of convenience yields and allow for a direct test of how regulation-induced market segmentation and other frictions affect asset prices.

Empirically, we show that the pricing distortions reflected in crypto carry are substantial. From April 2019 to July 2024, the average annualized carry across exchanges was approximately 7% p.a. In theory, a cash-and-carry arbitrage trader could earn this return by buying the cryptocurrency in the spot market and simultaneously selling a futures contract. By holding the spot position until the futures contract expires, the trader locks in the price difference and earns a return that is (in the absence of frictions) fully hedged against movements in the underlying asset and thus appears risk-free.

To understand crypto carry, we rely on the basic futures pricing equation as an organizing framework:

$$F_t = S_t \cdot e^{(r_t+u_t-y_t)(T-t)},$$

where F_t is the futures price for a contract maturing at date T, S_t is the spot price, r_t is the U.S. dollar risk-free rate, u_t denotes storage or other costs of holding the asset, and y_t refers to the convenience yield of the asset.

Based on this framework, we show empirically that the main observable "fundamentals" posited by this standard pricing equation cannot explain the level and fluctuations in crypto carry. The variation in interest rates r bears little to no relation to crypto carry, and the storage costs for holding spot crypto are negligible ($u \approx 0$).[2] This implies a large and *negative convenience yield* from holding spot crypto $y < 0$. In other words, investors prefer to be long the futures contract as opposed to being long the spot asset. This finding is in stark contrast to the case of commodities, where the convenience yield typically accrues when holding the physical asset (e.g., Gorton and Rouwenhorst, 2006; Koijen et al., 2018).[3] Interestingly, though,

[2]Interest rate variation cannot account for crypto carry even if we take into account potential income generated from holding the spot position, e.g., by lending out the spot asset through decentralized finance protocols such as Aave. We return to these issues in more detail in Section 2.

[3]In the commodities futures literature, the convenience yield is traditionally understood as the benefit of physically holding the spot asset rather than holding a futures contract. This immediate availability is valued by investors, for example, to prevent production disruptions or to meet urgent market demands.

a similar *inconvenience* of spot over derivatives has been documented in U.S. Treasury markets (e.g., Schrimpf et al., 2020; Duffie, 2020; He et al., 2022), where investors prefer interest rate swaps or futures over holding the underlying bond, due to balance sheet considerations.

What could explain this large inconvenience of spot positions relative to futures in crypto? *First*, futures allow investors to obtain leveraged exposure to crypto assets with little or no upfront capital, making them especially attractive for smaller, potentially unsophisticated investors who would otherwise face leverage constraints (see, e.g., the discussion in Frazzini and Pedersen, 2022).[4] Such investors could push futures prices higher relative to spot, driving up carry in the presence of limits to arbitrage. *Second*, and most importantly, various constraints and frictions, explored in detail below, limit the ability of sophisticated market participants to execute cash-and-carry trades in crypto markets. These limitations stem either from the inability to hold spot positions or from margin requirements. As a result, (i) constrained investors seeking exposure to crypto may exhibit a stronger preference for (regulated) crypto futures over crypto-native spot positions, and (ii) arbitrage activity to close the basis remains limited, as effective arbitrage requires the ability to trade both spot and futures contracts. More broadly, these frictions highlight the significant segmentation between the crypto ecosystem and the traditional financial system. We exploit several regulatory changes affecting this segmentation in our causal inference.[5]

In support of the first point above, i.e., the importance of less sophisticated investors, we document that a higher crypto carry is associated with a rise in net long positions of smaller, presumably less sophisticated traders, as inferred from the Commodity Futures Trading Commission's (CFTC) commitment of traders (COT) reports. We show that these types of investors increase their futures positions in times of strong price trends and heightened attention to crypto markets. On the other side of these investors are dealer intermediaries and leveraged funds who tend to take the opposite position by being short crypto futures. To identify the causal impact of higher demand from smaller investors for crypto carry, we exploit the introduction of micro bitcoin futures on the CME. The reduced size of these futures contracts made it easier for smaller

[4]While leverage via futures applies to other assets and is not unique to crypto, the maximum leverage attainable on crypto-native exchanges is much larger than the one for other assets like commodities or the S&P 500, with maximum leverage sometimes even exceeding 100 as seen from Table A.1.

[5]Aside from regulatory barriers, differential tax treatments for spot holdings of crypto vs. futures may also lead to pricing distortions. Indeed, crypto returns have been extraordinarily high, so any tax advantages of holding futures positions relative to spot are likely to be particularly valuable. However, as we discuss in more detail in Section 2, we do not find support for tax-induced distortions to be an important driver of crypto carry.

investors to trade bitcoin futures. In a difference-in-differences (DiD) regression setting, we show that the introduction of micro futures significantly increased crypto carry on the CME relative to other exchanges, which supports the hypothesis that higher demand from smaller investors increases crypto carry.

While these findings highlight the significant role of demand pressure from smaller investors, a crucial question remains: why do sophisticated market participants not arbitrage the basis using a standard cash-and-carry trade? Our central argument points to the strong segmentation between the traditional financial system and the crypto-native ecosystem, as well as margin frictions on crypto-native exchanges. In practice, arbitrageurs encounter several obstacles when trying to execute the cash-and-carry strategy. Chief among these is the regulatory environment, which has long made it challenging for many professional market participants to hold spot bitcoin as part of an arbitrage position. Additionally, even those arbitrageurs capable of holding spot positions face margin-related frictions, such as the absence of cross-margining on traditional exchanges and the complexity and lack of transparency in margin rules on crypto-native exchanges. For example, the lack of cross-margining means that a long spot position cannot be used as collateral for a short futures position on traditional exchanges like the CME, which are more accessible to larger institutions. All this means that implementing a long-short cash-and-carry arbitrage trade exposes the trader to funding risk (Brunnermeier and Pedersen, 2009). When the basis moves unfavorably before maturity, a spike in margins can force a liquidation of the position before convergence of spot and futures prices at maturity.

We provide several pieces of evidence in support of the limits to arbitrage explanation. To start off, we find that the futures leg of the cash-and-carry strategy exhibits severe drawdowns. The lack of cross-margining implies that such drawdowns can easily lead to futures contract liquidations of carry trade positions if margin calls cannot be met. In line with this mechanism, we find that carry is a significant predictor of liquidations in short futures positions. For example, an increase in the standardized carry by 10% predicts an increase in liquidations of short futures positions by 22% of total open interest over the next month. Hence, even patient carry traders are exposed to significant risk of forced liquidations due to the existence of margin frictions. Faced with such risk, arbitrageurs might be reluctant to deploy capital on a large scale to "lean against" a widening of carry during crypto price booms.

To further examine the impact of regulatory constraints that limit the deployment of arbitrage capital, we analyze the introduction of the spot bitcoin exchange-traded fund (ETF) in 2024. Although the presence of the ETF per se does not resolve the cross-margining friction, it at least allows institutional investors in the traditional financial system to hold (a proxy for) spot bitcoin to hedge bitcoin futures positions on the CME. Thus, the introduction of the ETF should decrease crypto carry. This is indeed what we find: in a DiD setting, we show that the introduction of the ETF significantly decreased crypto carry across exchanges by about three percentage points and by an additional five percentage points on the CME. In economic terms, these are very large declines of 36% and 97% of the mean crypto carry, respectively. This result provides causal evidence that margining frictions, which are central to our limits to arbitrage story, have an important bearing on crypto carry.

Related literature. Our paper contributes to the growing literature on crypto assets as well as the literature on convenience yields in other asset classes. Makarov and Schoar (2020) analyze arbitrage deviations in crypto assets and find that mere transaction costs cannot explain the size of arbitrage spreads, similar to our findings for crypto futures. Makarov and Schoar (2021) study the bitcoin ecosystem, disaggregated by individual wallets and different types of users over time. Liu and Tsyvinski (2021) find that attention is useful for predicting cryptocurrency returns, in line with our findings that investor attention is strongly correlated with the demand for leveraged crypto exposure through futures. Kogan et al. (2023) study retail traders and show that these investors tend to be trend followers in crypto assets, which is also consistent with our results and interpretations.

Recent papers study so-called perpetual crypto futures instead of the standard fixed-term futures analyzed in our paper, and complement our findings. He et al. (2022) document that perpetual futures allow investors to use leverage and typically trade at a positive funding rate, that is, the long side pays the short side. Christin et al. (2022) also use perpetual futures' eight-hour funding periods to show that cryptocurrency exchanges facilitate leverage for long-side traders and that a carry trade that is short perpetual futures generates high Sharpe ratios. Our paper provides a different perspective on basis trades, since we study fixed-date contracts with proper price convergence of spot and futures on the settlement date. Moreover, we provide a detailed account of the frictions that limit arbitrage in these longer-dated contracts and show why these trades are risky from an arbitrageur's point of view. Angeris et al. (2022) derive model-free expressions for the funding rate and the discount rate of perpetual futures. In contrast to fixed-maturity futures, perpetual ones are not guaranteed to converge to the spot price, since the contracts have no expiration date to strictly enforce arbitrage.

2. Stylized facts about crypto carry

This section starts with a brief description of the data sources and the construction of the crypto futures basis. We then present some stylized facts about crypto carry, examining its time series and cross-sectional (i.e., across exchanges) characteristics.

2.1. Data sources

We collect daily data from March 2019 to July 2024 on BTC and ETH spot prices, as well as futures and options market characteristics from the data providers 'Skew' and 'Coinmetrics'. The data on futures contracts contain the futures basis, trading volume, open interest, buy and sell liquidations. Options data include one-month implied volatility of BTC based on variance swaps and one-month realized volatility.

The futures basis data are annualized and refer to constant maturities of one and three months from a set of crypto trading platforms: Binance, OKEx, FTX, Huobi, BitMEX, Deribit, and the CME. The data for some exchanges are shorter (e.g., the CME basis starts in August 2020). From these data, we calculate the prices of one- and three-month fixed-maturity futures $F_{t,T}$, where T is the expiration date of the futures.

We complement these data with daily borrowing and lending rates from Aave (one of the largest decentralized finance (DeFi) lending platforms) and Binance (one of the largest crypto-native exchanges), and with weekly net trader futures positions from the CFTC's Commitments of Traders Reports. We also use data from Google trends on searches related to bitcoin futures. When examining the cash-and-carry strategy, we also use data on BTC CME futures and spot prices from Bloomberg starting in January 2018.

2.2. Crypto carry over time and across exchanges

A notable feature of crypto markets is that futures trade on a range of crypto-native exchanges such as OKEx and Binance, and more recently, also on regulated exchanges such as the CME.

These venues operate in markedly different ways. Futures contract specifications differ across exchanges, both within the crypto ecosystem as well as between crypto-native exchanges and the CME. For example, futures on the CME settle in cash against U.S. dollars (USD) and have to be collateralized by high-quality assets in USD. Crypto exchanges, in contrast, allow traders to post spot BTC (or other crypto assets) as collateral when opening a futures position and can be settled in crypto currencies instead of USD.

Attainable leverage on crypto-native exchanges is very high and significantly exceeds that on the CME. Table A.1 in the Online Appendix documents differences between exchanges and a range of maximum leverage from two (CME) to 125 on OKEx.[6] The higher leverage attainable in turn might attract more smaller retail traders to these venues, which we will come back to in our empirical analysis below.

[Figure 1 about here]

In terms of the size of crypto futures markets, Figure 1 shows that Binance has the largest dollar open interest on average for both BTC and ETH futures. The total dollar open interest on all BTC exchanges in our sample was USD 20.8 billion as of July 2024.[7] Most of the open interest in BTC futures contracts (78%) is traded on crypto-native exchanges, and the regulated CME accounts for the residual 22%. In terms of trading volume, crypto-native exchanges account for an even larger share of 93% of total volume.

[Table 1 about here]

Table 1 shows summary statistics for crypto carry for a representative crypto-native exchange (OKEx) and the only regulated one in our sample (CME).[8] We chose OKEx here and in other

[6]That said, maximum leverage has come down on many of these exchanges as well in the past years. For example, Binance and FTX reduced the maximum leverage on their platforms to $\times 20$ (from $> \times 100$ before) in July 2021. Although the maximum leverage was greater than 100 on these exchanges, anecdotal evidence suggests that the average leverage employed by traders was less than 20, and the reduction of the maximum leverage did not significantly affect the basis or trading volume. See also the tweets by FTX CEO Sam Bankman-Fried at the time as reported in https://www.coindesk.com/markets/2021/07/25/ftx-cuts-leverage-limit-to-20x-from-100x-as-criticism-of-margin-trading-in-crypto-grows/.

[7]Together with CoinFlex, ByBit and Bitfinex, which are not in our sample, the total open interest is around USD 25.5 billion. For ether (ETH), Binance, FTX and OKEx have the largest dollar open interest. The total dollar open interest on all exchanges in our sample was USD 8.5 billion as of July 2024.

[8]We use the longest available time series for all regression variables in our analysis and report summary statistics for basis on CME after August 2020 since that is our main regression period for CME.

parts of the paper as a representative nonregulated exchange, because it has complete data for both one-month and three-month carry, for BTC and ETH.[9] Table A.1 shows the cross-sectional summary statistics on how carry differs across exchanges.

On average, the one-month bitcoin basis is about 8% on OKEx and 6.4% on the CME, with maximum values reaching approximately 55% and 45%, respectively. The same statistics for three-month contracts show a similar pattern and are only slightly lower. In sum, crypto carry is high and suggests a high inconvenience of spot positions over futures contracts.

Carry is very persistent at the first lag (daily frequency) with autoregressive coefficients only slightly below one. The volatility of carry tends to decline in the maturity of the contract, with one-month carry being more volatile than three-month carry. Interestingly, carry is right-skewed for both BTC and ETH and for both maturities (except CME for the one-month maturity), i.e. the distribution of carry features a long tail of large positive observations. For cash-and-carry arbitrageurs, which typically are in a long spot position and short the futures contract, this asymmetry in the distribution of carry implies a high risk of large drawdowns, an issue to which we return to below. In addition, carry on OKEx and other crypto-native exchanges is much more volatile than on the CME, possibly driven by the higher maximum leverage of nonregulated crypto exchanges.

[Figure 2 about here]

Figure 2 shows the evolution of carry for OKEx and CME. The figure illustrates the high average level of crypto carry of about 6–8% p.a., and frequently exceeding 20%. Such levels are higher than the carry of other assets such as equities, fixed income, currencies and commodities (see Koijen et al. (2018) and Section 2.4 below). The concurrent work of He et al. (2022) independently documents that bitcoin futures prices are typically higher than spot prices using perpetual bitcoin futures. Figure 2 also illustrates the large, sudden spikes in crypto carry that occur several times during our sample period. For example, there are three periods for the one-month BTC futures on OKEx (early 2019, early 2020, and March 2021), during which crypto carry reached or exceeded 40% p.a. before declining significantly in a short period of time. CME carry is also volatile, ranging from below -50% during the FTX collapse in November 2022 to above 45% before the launch of the spot BTC ETF in January 2024. These features make a

[9] The other such exchange is Huobi but it has fewer observations.

tax-based explanation of crypto carry unlikely since tax advantages of futures over spot are hard to square with a highly volatile basis that shows frequent large spikes.

[Table 2 about here]

The upper two panels in Figure 2 illustrate another striking feature, namely that the carry for BTC and ETH is similar in magnitude and highly correlated on a given exchange for maturities of 1 and 3 months. Table 2 expands on this observation and reports correlation coefficients for the carry of BTC and ETH and for both maturities on OKEx (Panel A), as well as correlations of one-month BTC carry across different exchanges (Panel B). The table shows that carry is highly correlated across nonregulated exchanges with correlation coefficients in excess of 90%. That said, it is especially noteworthy that the CME carry is the *least correlated* with that of other crypto exchanges, which indicates some degree of market segmentation: see, e.g. Makarov and Schoar (2020) for an empirical analysis of price inefficiencies in crypto markets. Possible explanations include differences in contract specifications, differences in trading hours, and the difference in settlement currency between unregulated exchanges and the CME discussed above.

Before moving on to dissecting crypto carry into different factors, a potential explanation for the preference of futures over spot crypto could be due to taxes. Specifically, futures contracts can, in principle, be held in tax-advantaged accounts, which is not true for spot crypto. This tax benefit, combined with the historically high returns of cryptocurrencies, might create a higher demand for futures contracts relative to spot, contributing to a significant crypto carry. However, this rationale primarily applies to U.S. investors and CME-traded contracts and does not apply to those on unregulated overseas exchanges like OKEx or Binance, which are not eligible for U.S. tax-advantaged accounts. If tax considerations were a key factor, one would expect higher carry on the CME compared to unregulated exchanges. Interestingly, the data does not support this hypothesis. In fact, the carry on CME is on average lower than on crypto-native, unregulated exchanges as we shoed above.

2.3. The crypto (in)convenience yield

How can we understand the magnitude and large swings in crypto carry? For a simple framework guiding our subsequent analysis, consider again the basic futures pricing equation:

$$F_t = S_t \cdot e^{(r_t+u_t-y_t)(T-t)},$$

where F_t is the price of a futures contract maturing at date T, S_t is the spot price, r_t is the U.S. dollar risk-free rate, u_t denotes storage or other costs of holding the asset, and y_t refers to the convenience yield of the asset. To accommodate the special features of crypto assets, we further specify the convenience yield y_t as the sum of two components $y_t = r_t^\star + \delta_t$ where $r_t^\star$ captures the yield an investor can earn on spot crypto holdings (e.g., interest earned through lending protocols such as Aave or staking yields on ether), and δ_t captures any additional convenience from holding the spot asset beyond direct interest income. One can think of δ as a "net convenience yield", i.e., the convenience of holding spot crypto assets net of earning a direct (monetary) yield.

Taking logs of the pricing equation, setting $T = t + 1$ for ease of exposition, and adding an exchange-specific pricing error ε^i, we have that:

$$\underbrace{f_t - s_t}_{\text{Carry}} = \underbrace{(r_t - r_t^\star)}_{\text{Interest differential}} + \underbrace{u_t}_{\text{Storage cost}} - \underbrace{\delta_t}_{\text{Net convenience yield}} + \underbrace{\varepsilon_t^i}_{\text{Pricing error}}, \qquad (1)$$

where f_t and s_t now denote log futures and spot prices, respectively.[10] Importantly, δ_t denotes the *crypto convenience yield* of spot over futures positions net of any direct monetary yield of the crypto asset (which is captured by $r_t^\star$). As Eq. (1) makes clear, the observed fluctuations in carry could stem from (i) variation in interest rate differentials, (ii) variation in storage costs, (iii) variation in pricing errors across exchanges i, or (iv) variation in (net) convenience yields.

To shed light on the relative importance of the basic drivers posited by this framework, we conduct two diagnostic tests. First, we dissect the contribution of time-specific variation (t) and exchange-specific (i) variation to the total variability of carry in a panel of all exchanges in our sample. Table 3 shows results for panel regressions in which one-month carry is the dependent

[10] We add an exchange-specific pricing error here since it was not straightforward to arbitrage prices on different crypto-native exchanges in the early days of crypto, i.e. before the rise of stablecoins.

variable. The first two specifications simply regress carry on exchange fixed effects (FEs) or time FEs, respectively. The R^2 with exchange FEs is less than 0.5%, while the R^2 with time FEs is 88%. We find similar results for ETH, where exchange-specific factors capture less than 0.2% of the variation, while time-specific factors capture 93% (these regressions are not reported for brevity). These results are well in line with the visual impression from Figure 2 above and show that exchange-specific factors do not explain a significant share of the variation of crypto carry.[11]

[Table 3 about here]

The second test concerns the role of interest rate differentials, which, according to Eq. (1), could be considered the main observable "fundamental" driver of crypto carry. There are two types of crypto interest rates that can matter for carry. First, if we take the perspective of an investor in the crypto world, the appropriate risk-free rate depends on whether the investor already has an endowment of spot BTC or not. If the investor has the asset, then the opportunity cost of deploying it in a cash-and-carry arbitrage is to lend out the asset to earn a yield. As a rough approximation for that rate, we use one-month BTC lending rates on the DeFi platform Aave.[12] Second, in case the investor does not have the spot asset, she can borrow it either from the same pool (if she has some other crypto-asset to pledge as collateral), or she can source the asset from other platforms, paying higher rates. To proxy for these larger rates, we use Binance borrow rates for the top client ("VIP") category (assuming carry traders are able to borrow at the most attractive rates). Results with the rate for the lowest client category are similar. Finally, one can also take the perspective of an investor without any prior endowment of crypto assets who can use funding in fiat currency to finance the collateral pledged on a regulated exchange such as the CME. The appropriate interest rate can be approximated by LIBOR (plus some funding spread depending on the creditworthiness of the investor). Using other fiat currency benchmark rates, e.g. a repo-based reference rate such as SOFR instead of LIBOR, yields very similar results.

[11] In column (3) of Table 3, we also include contemporaneous changes in VIX to proxy for potential pricing errors due to (time-varying) financial constraints that may affect arbitrageurs who also operate in the traditional financial system. Such constraints could drive variation in δ_t and explain the (co-)variation in carry observed in the data. However, the results in column (3) do not suggest an important role for such constraints in explaining crypto carry.

[12] Aave is a large decentralized finance protocol that allows users to lend and borrow several crypto-assets at variable and fixed interest rates.

The plot of these interest rates in the third panel of Figure 2 shows that they are not variable enough to explain the more volatile carry of crypto futures – regardless of whether we take interest rates in the crypto or in the fiat world. In addition, even the riskiest borrowing rate from Binance is still too low compared to the one implied by crypto carry, on average. Overall, the evidence from the plots suggests that interest rate differentials cannot explain the variation in carry observed in the data.

To test more formally whether interest rate differentials span carry, we also regress carry on the two measures of interest rate differentials described above, in columns (4) and (5) of Table 3. Interestingly, while the R^2 is low in both regressions, the two interest rate spreads are negative and statistically significant in each case. While the benchmark futures pricing equation in Eq. (1) would prescribe a coefficient of minus one in these regressions in the absence of frictions, both coefficients are far from this predicted theoretical value and confirm that carry cannot be explained by interest rate differentials alone.

These results leave storage costs and convenience yields as sole drivers that could explain fluctuations in crypto carry. Storage costs are an unlikely explanation since holding spot bitcoin or ether at an exchange (like Coinbase) is essentially cost-free. Likewise, self-hosting spot positions in a wallet does not incur significant costs (except maybe for some upfront hardware costs). In any case, even if these costs were non-zero, they would not vary much over time and thus cannot explain the large fluctuations observed in crypto carry. This leaves crypto (in)convenience yield, i.e. δ in Eq. (1), as the only remaining driver of the basis. Section 3 studies what gives rise to this convenience yield and addresses the crucial question why arbitrageurs do not step in on a scale that is large enough to take advantage of the large crypto basis.

2.4. Comparing carry across asset classes

To set the stage for the following analysis, it is instructive to compare crypto carry to the carry of other asset classes. This comparison seems especially relevant in light of an earlier literature showing an important role for (in)convenience yields in driving the basis, e.g., for commodities or for derivatives linked to government bonds.

First of all, the sign of the crypto basis implies a large and time-varying inconvenience of holding spot positions relative to futures. This fact is surprising in light of earlier research in commodity markets (e.g. Gorton and Rouwenhorst, 2006; Koijen et al., 2018) which documents the opposite, i.e., a convenience yield for holding the spot asset. This could be due to, e.g., the convenience of holding physical gold or having access to oil in storage instead of a futures claim. Unlike commodities, negative convenience yields can be observed in some government bond markets at different times (Schrimpf et al., 2020; Duffie, 2020; He et al., 2022), driven by demand and supply imbalances that cannot be fully arbitrated due to funding and balance sheet constraints. In U.S. Treasury markets, this issue stems from securities holdings consuming dealer banks' balance sheet capacity, whereas derivatives, being off-balance sheet, impose fewer constraints. Inconvenience yields can hence arise when regulatory constraints limit intermediaries' ability to make their balance sheet available to absorb shocks, leading to persistent price dislocations. Similarly, in crypto markets, regulated investors face challenges in holding the spot asset but can more easily hold derivatives, especially on regulated exchanges like the CME. As with U.S. Treasuries, inconvenience yields in crypto arise because regulatory constraints reduce the relative attractiveness of the spot asset compared to that of futures contracts.

Table A.2 in the Online Appendix provides some perspective on carry across asset classes and shows summary statistics of carry for the sample period from February 2018 to July 2024 for equities (S&P 500), fixed income (U.S. Treasuries), oil, and gold. While the average carry is hard to compare across asset classes, due to differences in storage costs and yield spreads, an interesting observation is that crypto assets have by far the largest carry of all asset classes in the table: at roughly 7% p.a., crypto carry is around *10 times larger* than the carry of the S&P 500 and more than *12 times larger* than the one of U.S. Treasuries. Carry in commodities markets is also much smaller and even negative in the case of oil.[13] The carry of the S&P 500, U.S. Treasuries and commodities is also much less volatile (2-7 times less) and also exhibits less extreme positive values compared to crypto carry as seen from the table. Moreover, the last row reports correlations of an asset's carry with crypto carry, and we find these correlations to be close to zero. We thus conclude that crypto carry is distinct from carry in other assets and requires an explanation that is tied to the specifics of the asset class.

[13] For oil, we exclude the period of negative oil prices in April 2020.

3. Understanding crypto carry

Our results so far indicate that fluctuations in crypto carry largely owe to movements in convenience yields. We now go a step further and investigate the *economic drivers* of crypto carry. In particular, we explore why convenience yields in crypto markets fluctuate so strongly and why smart capital does not take full advantage of the high yields implied by crypto carry.

We first study the investor composition of crypto futures markets and the positioning of different classes of investors. This allows us to get a better sense of the different counterparties trading crypto futures and their likely motives. We then analyze the role of less sophisticated investors that seek leveraged exposure to crypto, before turning to factors that might limit the deployment of capital leaning against the trades of such investors. In both cases, we provide causal evidence based on the introduction of new financial instruments that facilitated the participation of smaller unsophisticated investors (micro bitcoin futures) as well as institutional arbitrageurs (spot bitcoin ETFs), respectively.

3.1. Positioning of different types of investors

To study how the trading actions of different investors impact crypto carry, we run regressions of carry on net CME futures positions of various trader groups. We identify those positions based on the COT reports by the CFTC and run:

$$\text{Carry}_t = \alpha + \beta_D \Delta X_{D,t} + \Delta\beta_I \Delta X_{I,t} + \Delta\beta_L \Delta X_{L,t} + \beta_N \Delta X_{N,t} + \epsilon_t, \tag{2}$$

where $\Delta X_{D,t}, \Delta X_{I,t}, \Delta X_{L,t}, \Delta X_{N,t}$ are the changes in positions of dealer intermediaries, institutional investors, leveraged funds, and nonreportables, respectively, in week t. The group of other reportables is left out and thus serves as a reference category, since net futures positions sum to zero across all investor groups. In these regressions, we use one-month BTC carry (the results for ETH and three-month carry are similar) at a weekly frequency to match the CME position data.

[Figure 3 about here]

The left panel of Figure 3 illustrates how the positions of different types of investors evolve over the sample period. Nonreportable traders are net long futures on average, whereas leveraged funds and dealers are net short on average. Positions of other reportables are fluctuating around zero, whereas institutional investors' positions were flat and close to zero before late 2021 but increased substantially afterwards. The spike in institutional investor positions towards the end of 2021 and the rise of their long positions thereafter are likely related to the launch of the first U.S. bitcoin ETF in October 2021 (Todorov, 2021) and subsequent inflows into this futures-based ETF. There also appears to be some migration of positions from nonreportables and other reportables to institutional investors after the launch of the ETF.

The right panel of Figure 3 reports regression estimates for Eq. (2) above, all of which are significant at the 5% level except the one for dealer intermediaries. More interestingly, we find that increases in net long positions by nonreportables are associated with a larger basis, whereas leveraged funds' net long positions are negatively correlated with it. The coefficient for dealer intermediaries is smaller in magnitude and insignificant.

Since the group of nonreportables in the COT data captures the activity of smaller investors (e.g., family offices, smaller trading shops, and proprietary traders), these initial results square well with the notion that a high carry, and thus inconvenience of spot positions, typically comes along with a rise in positions of smaller and likely less sophisticated investors. In contrast, larger leveraged institutions typically take the opposite side, effectively providing liquidity. Aside from outright short futures positions, such liquidity provision could take the form of collecting the returns on a cash-and-carry trade.

The composition of traders in bitcoin futures shows that the largest long positions come from smaller investors before 2022, in contrast to most other futures markets, which are typically dominated by commercial entities and institutions (according to data from the CFTC). Thus, our analysis of crypto carry also provides useful insights about a futures market dominated by smaller (potentially less sophisticated retail) investors on the long side.

3.2. Demand for leveraged upside exposure by small and trend-chasing investors

We now seek to establish more formally the link between the demand for futures by smaller trend-chasing investors and increases in crypto carry. To better understand the drivers of crypto carry – and to assess whether crypto carry aligns with the attention-driven, trend-following crypto speculation often depicted in the media – we explicitly test for the role of trend-following and attention-based trading behavior. We do so by running regressions of carry on potential drivers of the convenience yield. To keep the presentation concise we present the regression of carry using OKEx as a representative crypto exchange in our results below. We choose OKEx since it has one of the largest open interest and the longest data available in our sample. Results for other crypto exchanges are similar. In addition, we also present results for carry on the only regulated exchange, that is, the CME.

[Table 4 about here]

In the first column of Table 4 we regress one-month crypto carry on an index for Google searches for "BTC", "bitcoin", "bitcoin futures", "bitcoin price", "bitcoin leverage" from Google Trends. The main motivation behind the Google search variable is to capture the interest in the trading of smaller and potentially less sophisticated traders.[14] For OKEx carry, we find a strong and positive link between changes in the number of Google searches, and crypto carry, with a relatively large R^2 of 12%. The explanatory power of the Google variable for the CME basis is much smaller with an R^2 of only 1%, and the estimate is insignificant, as seen from column (7) of Table 4 (this is also true if we run the OKEx regressions for the same sample period as for CME). The difference between these two exchanges could be explained by prices on nonregulated exchanges being more heavily influenced by retail traders whose sentiment is likely captured by the Google search variable. In contrast, on regulated exchanges, such as the CME, which are populated by a relatively larger share of institutional investors, the impact of changes in the attention of retail investors seems to be less pronounced.

In the specifications of columns (2), (3), (8) and (9) of Table 4, we include BTC returns over the past week and month. The idea here is to analyze how possible trend-following behavior is related to crypto carry. Consistent with a trend-chasing argument, the estimates in the table show that higher past crypto returns are strongly associated with a larger crypto carry.

[14]In a previous version of the paper, we used changes in social media users on Reddit and showed that the results are similar to using the Google Trends variable. The motivation behind the Reddit variable was to capture the interest in trading of smaller, potentially less sophisticated traders that coordinate on social media platforms. Unfortunately, the data on the Reddit measure is less reliable after 2022. In addition, the measure suffers from asymmetry in crypto booms vs. busts. Whereas the number of users increases in boom periods as more people subscribe to Reddit groups dedicated to BTC and ETH, the number of users does not drop substantially during crypto busts since investors do not unsubscribe from these groups. In contrast to the Reddit variable, the number of Google searches fluctuates much more between periods of booms and busts, potentially reflecting the changing sentiment of smaller traders.

3.2.1. *Price pressure*

To address the concern that trend-chasing variables simply capture general price pressure effects, we also add proxies for price pressure to the regression and report the results in columns (4) and (10). The first measure is the signed trading volume (futures trading volume multiplied by +1 for positive futures returns and −1 for negative futures returns), and the second is futures open interest. We find a positive and significant effect of both measures for OKEx and CME, with a large R^2. However, in the richer specifications (columns (5) and (11)), only open interest remains statistically significant. These results support the idea that price pressure effects (or "downward sloping demand curves") have a bearing on crypto carry.[15] More importantly, the effects of the Google search variable and trend-chasing variables remain significant and positive also after controlling for the price pressure variables in columns (5) and (11).

The carry and trend-chasing variables (number of Google searches and past returns) could also reinforce each other (as highlighted, for example, by Plantin and Shin (2014)): As carry increases, the trend-chasing variables would also increase, putting further upward pressure on carry. To account for these effects, we also control for lags of the basis and verify that the estimates remain positive and significant in columns (6) and (12). In unreported results, we also verify that the main results of the trend-chasing hypothesis hold in a VAR setting and after accounting for several lags of the dependent and independent variables (like in a linear projection model or vector autoregression).

[Table 5 about here]

3.2.2. *Difference-in-differences regressions around the CME micro BTC futures introduction*

The regressions so far show a strong positive relation between carry and appetite for crypto exposure by smaller investors. To establish this link more formally, we run difference-in-differences (DiD) regressions around an event that made it easier for smaller investors to trade bitcoin futures. The event is the introduction of micro bitcoin futures by the CME on 3 May 2021. Whereas the standard-sized bitcoin futures contract on CME is equivalent to five bitcoins, the

[15] As shown in Hong and Yogo (2012), open interest can not only be a measure of price pressure, but it could also reflect an information signal as in other commodity markets. However, the latter interpretation seems unlikely in our context, since we also show that high carry forecasts future price crashes and that it is positively correlated with attention measures of smaller and potentially less sophisticated investors.

micro bitcoin futures contract is equivalent to one-tenth of one bitcoin, i.e., 1/50 the size of the standard bitcoin futures contract. Other exchanges (OKEx, Huobi, Deribit, and Kraken) did not experience similar changes around the event. Therefore, the difference between carry on CME and these other exchanges should reflect the effect of the contract size decrease on CME. To test this, we run DiD regressions using a two-months window around the introduction of the micro futures contract:

$$Carry_{t,i} = \gamma \times Treat_i \times Post_t + \beta_1 \times Treat_i + \beta_2 \times Post_t + \epsilon_{t,i}, \tag{3}$$

where $Treat_i$ is a dummy equal to one for CME, and $Post_t$ is a dummy equal to one for the period after the introduction of the micro contract. The coefficient of interest is γ, and the estimate should be positive if the introduction of the micro contract increased the relative difference in the basis on the CME versus that on other exchanges. The results in column (1) of Table 5 show that this is indeed the case: the estimate of γ is positive and highly significant, and illustrates that the relative difference between carry on the CME and that on other exchanges increased by 11%, after the introduction of the micro contract. The coefficient is unaffected by controlling for the return on ETH (which captures general trends in crypto prices around the event) as seen from column (2). Figure 4 shows that the assumption of parallel trends is largely satisfied.[16] These results provide causal evidence that higher demand by smaller, presumably less sophisticated investors pushes up carry.

[Figure 4 about here]

3.3. Limits to the deployment of carry trade capital

Our results so far suggest that the buying pressure by smaller trend-following investors in futures markets generates upward pressure on crypto carry. However, it is important to emphasize that a fully functioning arbitrage mechanism should neutralize such demand pressure. We therefore

[16] The figure shows that it took less than a week for the difference in bases between control and treated to decrease, which might be because it takes a few days for smaller investors to move funds to CME and open a futures account at the exchange. Running the regressions at a weekly level shows that the treatment effect takes place in the first week. The figure also illustrates that there was a downward trend for the basis before the event date that flattened for the CME after the introduction of micro bitcoin futures but continued for other unaffected exchanges.

now turn to the question why cash-and-carry traders do not arbitrage away price dislocations generated by smaller trend-following crypto investors.

In Section 3.3.1, we first provide a concise overview of how the cash-and-carry strategy works and then describe key frictions that arise in the context of crypto markets. Building on this, we then present empirical evidence that illustrates these limits to arbitrage in Section 3.3.2.

3.3.1. *Key frictions in crypto cash-and-carry trades*

To understand the incentives and risks faced by arbitrageurs, we first describe a textbook cash-and-carry strategy applied to the crypto basis. Specifically, an arbitrageur would take an offsetting position in spot and futures markets to benefit from the price discrepancy and bet on their convergence until maturity of the futures contract.

If carry is positive at time t and large enough to cover transaction and financing costs, a cash-and-carry arbitrageur would take a long spot position and a short futures position. Hence, if $F_{t,T} > S_t$, the payoff for the arbitrageur is positive at time t and zero at maturity T since $F_{T,T} - S_T = 0$. Likewise, if the current carry is negative ($F_{t,T} < S_t$), the carry trader would sell spot and buy the futures so the strategy earns a positive payout of $S_t - F_{t,T}$ at time t and has a zero cash flow at T.

This simple setup thus allows, in theory, for 'risk-free' convergence trades in both directions, depending on the sign of the carry when the trade is initiated. In both cases, the spot position hedges the futures position and there is no risk for the arbitrageur. However, there are at least two key frictions stemming from direct regulatory barriers or, more indirectly, from regulatory uncertainty in crypto markets that make cash-and-carry trades risky and less profitable than in this stylized textbook example.

First, on regulated exchanges (CME), there is no way of *cross-margining* between the spot and the futures position because traders cannot hold the spot position at the exchange but need to pledge liquid assets in fiat currency as collateral (e.g., Treasury securities). This is a crucial friction that would prevent more sophisticated carry traders from the traditional finance world from shorting futures when carry is positive: losses on the short futures position are not automatically offset by gains in the long spot position because the spot asset cannot be held with the CME. In essence, carry traders have to pledge capital twice: once to buy the spot position (held in a separate account elsewhere), and again in the form of fiat collateral on the futures exchange.[17]

Second, the opacity of margining rules and the absence of cross-margining on nonregulated exchanges can also act as a deterrent for carry traders to deploy more capital. To give an example, FTX (before its collapse) specified a maximum loss of USD 30,000 on a futures position. This means that a one million USD short futures position would be liquidated after a 3% rise in BTC. Other exchanges use USD as the numeraire when calculating margin balances, similar to CME, and thus require liquidations of the spot asset when the arbitrageur faces large mark-to-market losses. However, this liquidation removes the hedge asset used in the cash-and-carry strategy, and also leads to premature liquidation of the position. For example, if the trader is short one futures contract and long one spot bitcoin, liquidating part of the spot position to raise cash and cover the loss in the futures contract leaves the trader with less than one unit of spot and leads to an imperfect hedge with exposure to the underlying.

The upshot of these frictions, as well as the resulting market segmentation between crypto and traditional finance, is that arbitrageurs face a heightened risk of forced liquidations before maturity. Therefore, the cash-and-carry trade in crypto is far from being a free lunch. Given the large spikes in carry observed in Figure 2, there is substantial risk that a cash-and-carry position, with spot and futures held in separate accounts, will be subject to a forced liquidation before convergence of the trade.

3.3.2. Empirical evidence on the importance of limits to arbitrage

Cash-and-carry: Risk and return. The frictions described above make it clear that potential arbitrageurs have to fund and manage the two legs (spot and futures) of the cash-and-carry strategy separately. To provide some perspective on their risk and return, we start by computing the mean returns, standard deviations, and drawdowns for the spot leg and the futures leg, implemented for CME one-month futures contracts on BTC. For these results, we assume the collateral is financed at LIBOR but the results are similar if we use other benchmark rates like SOFR. Table 6 reports the results for two periods: 2018–2024 (longest available CME sample)

[17] The cross-margining friction described here applies to market participants that can hold spot crypto positions in the first place. Regulatory barriers have prevented some market participants from holding spot crypto positions completely so that these entities would not be able to engage in cash-and-carry trades in the first place.

and 2019–2024 (our main sample used in the rest of the paper).

[Table 6 about here]

Table 6 shows that the mean excess returns on the futures leg of the basis trade are sizable at 2–3% *per month*, which is much higher than the futures returns in traditional asset classes such as equities, fixed income, currencies, and commodities (see Heston and Todorov (2023)). However, most importantly, Table 6 also illustrates that crypto futures returns are highly volatile (about 17% per month).

Risk of forced liquidation. We now turn to empirical evidence to shed light on the key mechanisms that present a hurdle to arbitrage activity. We first quantify the risk of liquidation for carry traders before turning to the link between movements in carry and futures liquidations in the data.

The risk profile of the futures leg in the cash-and-carry strategy is a critical consideration for traders. Figure 5 illustrates the profit and loss (PnL) trajectories for the futures leg of the strategy prior to maturity. Specifically, the analysis reveals that, assuming a leverage of 10 (which is significantly lower than the maximum leverage offered by most exchanges, as shown in Table A.1), the futures leg of the strategy would have been liquidated in over half of the months in our sample. These findings highlight the substantial risks associated with the futures leg of the cash-and-carry strategy, particularly when it is not protected by cross-margining, despite its high average return.

[Figure 5 about here]

Taking this analysis a step further, as per our previous discussion, increases in carry should in fact predict significant liquidations of existing short positions as traders unwind their carry trades before expiration. To see this, note that the typical cash-and-carry trade is short the futures contract. An increase in carry would thus imply losses on the futures leg of a cash-and-carry trader, and the absence of cross-margining raises the risk that the arbitrageur is forced out of the position. In contrast, since crypto carry is positive on average and carry traders are typically short the futures contract, we would not expect to see decreases in carry to lead to forced liquidations of long positions.

To test this prediction, we regress the amount of buy liquidations and sell liquidations in futures contracts, that is, the closing of open contracts, on the lagged crypto basis, in Table 7. Sell (buy) liquidations are the cumulative amount of short (long) futures positions that were liquidated over a given month (expressed as a percentage of open interest). These measures capture both forced liquidations by the exchange due to missed margin calls and voluntary liquidations by the trader.

[Table 7 about here]

The results of Table 7 are consistent with our conjecture above and show that a higher carry predicts liquidations of short futures positions. The table also shows that crypto carry significantly predicts only sell liquidations: a rise in standardized carry by 10% predicts a 22% increase in total sell liquidations (relative to total open interest) over the next month. Hence, while a high carry predicts overall greater risk (since implied volatility goes up when carry rises as indicated in columns (1) – (2)), it only forecasts contract liquidations of short positions, but not of long positions. The latter fact is consistent with contract liquidations of cash-and-carry trades as these trades involve short futures positions (for most of our sample period).

Difference-in-differences regression around spot BTC ETF introduction. To provide causal evidence on the limits-to-arbitrage explanation for crypto carry, we now run a DiD regression around the spot BTC ETF introduction in 2024. As outlined above, the inability of arbitrageurs to hold spot bitcoin on CME creates a friction and likely prevents them from bringing down the basis. If sophisticated, capital-rich arbitrageurs that come from the regulated fiat world were able to hold spot BTC, that should alleviate (albeit not eliminate) these constraints. The deployment of more capital to the carry trade, in turn, should dampen the basis. To test for that mechanism, we run an event study DiD similar to Eq. (3) around the introduction of the spot BTC ETF in January 2024, which allowed CME investors to hold spot BTC and made it easier to implement the crypto cash-and-carry trade.

Consistent with our conjecture, we find that the introduction of the spot bitcoin ETF decreased the basis on all exchanges significantly but disproportionally more so on the CME. The estimates on the $Post$ and $Treat \times Post$ dummies in column 3 of Table 5 show that the introduction of the ETF decreased the basis of all exchanges by around 3%, and on the CME by additional 5%. These are very large declines of 36% and 97% of the mean basis, respectively (given the mean basis of around 7.5%). Figure 6 illustrates the drop in the basis graphically.

[Figure 6 about here]

3.3.3. Summary: limits to arbitrage

The evidence presented in this subsection challenges the common perception of the crypto carry trade as a risk-free arbitrage opportunity. Instead, this strategy exposes arbitrageurs to significant risks, particularly when small trend-following investors drive up futures prices before maturity, leading to mark-to-market losses for carry traders holding short futures positions. A deeper underlying cause of this risk lies in regulatory barriers, which have kept crypto markets highly segmented from the traditional financial system. The absence of cross-margining further exacerbates the risk of sharp price movements, potentially triggering position liquidations. However, the introduction of the spot bitcoin ETF in 2024, while not directly addressing the cross-margining friction, has facilitated institutional investors' ability to hold spot assets as part of a cash-and-carry strategy. As we demonstrate using causal inference, this new ETF has reduced the segmentation of crypto markets and compressed the basis.

3.4. Robustness checks and other potential drivers of crypto carry

In this section, we report the results of several robustness tests. We also explore the validity of other potential explanations, such as transaction costs and hacking risks, but do not find them to be supported by the evidence.

Forced vs. voluntary liquidations. Our measure of futures positions' liquidations (Table 7) covers both forced and voluntary liquidations since the data provider does not distinguish between those two cases. One concern is that these liquidations might also capture events in which arbitrageurs actively close their short futures positions before the expiration and realize the arbitrage profit. Arbitrageurs do not have to wait until the expiration of the futures contract but can close out their short positions before the expiration, in particular when the basis is smaller compared to the time when the position is entered. To address this concern, we exclude such cases, and re-run our main regressions.

More specifically, we focus on the cases where the basis widens for 70% of the days in a two-week period before expiration and never falls below the initial value at the beginning of the two-week period. This subsample then captures cases where closing the short futures position is not profitable as a trader would face a loss.[18] We then rerun our main liquidation regressions (see Table 7) for this subsample. As shown in columns 3 and 4 of Table 7, the main results also hold for this particular subsample in which cash-and-carry positions are not in profit. The coefficients are also close to the ones for the main sample, which shows that the effects we capture in the longer sample are indeed likely related to liquidations and not to profit-taking before maturity.

Different horizons for returns and accounting for bear markets. For robustness, we assess to what extent different horizons for the past return measures in Table 4 (one day and two weeks) alter our findings. Our results remain broadly similar, with the one-day return being more significant than the two-week return in the multivariate regression for OKEx (column 5) but not for the CME (column 11). These results are excluded from the paper for brevity, but are available on request.

We also verify that the findings of the trend-chasing explanation for crypto carry hold also in a bear market as seen from Table A.3 in the Online Appendix.

Hacking risks and storage costs. Another possible factor explaining crypto carry is that it may reflect a compensation for hacking risks or storage costs. However, these factors are unlikely to explain the large variation in carry: Hacking risks would affect the futures and the spot similarly if they were traded on the same exchange. Storage costs, in turn, should be fairly stable and cannot plausibly explain the high volatility and (absolute) size of carry. Moreover, the exchange fixed effects we employ in our regression in Table 3 would capture such factors since variation in hacking risk and storage costs are likely to be exchange-specific. However, Table 3 shows that exchange fixed effects explain less than 0.5% of the variation in carry. In addition, a higher risk of hacking would plausibly correlate with a lower volume of trading on an exchange, since market participants should be reluctant to use such an exchange (Makarov and Schoar (2020)). However, the large trading volumes on OKEx and CME relative to other exchanges (reported in

[18]The initiation of the trade could have taken place prior to the start of a two-week period, of course. However, if the trader has not closed the position before the start of the period, closing it after the start of the period would be suboptimal since the basis is wider and the profit is smaller.

Figure 1) are inconsistent with hacking risk being a major driver of the returns on crypto carry.

Bid-ask spreads and exchange transaction fees. Bid-ask spreads are also unlikely to explain the large returns on crypto carry. Bid-ask spreads on spot bitcoin are much smaller than crypto carry and less than 0.2% for most crypto-native exchanges (Makarov and Schoar (2020)). Average bid-ask spreads for futures are also small relative to the size of crypto carry and are less than 3% for CME futures in our sample period. In addition, note that the carry trader does not incur bid-ask spread costs if she holds the position until maturity but only if she chooses to liquidate the position prematurely. Exchange trading fees are also too small to explain the large magnitude of crypto carry, also because many crypto exchanges do not charge fees on a trade-by-trade basis but based on the trading volume in a given month or week. Makarov and Schoar (2020) estimate these to range from 0.25% of the amount traded to 0.1%, which is much less than the size of crypto carry.

Blockchain transaction fees and mining costs. Another potential driver of crypto carry could be cryptocurrency transfer fees such as blockchain transaction fees or volatile energy costs associated with bitcoin mining. However, such costs are unlikely to explain carry: If agents were to trade the spot and the futures on the same exchange (e.g., OKEx), the cash-and-carry strategy would not involve transactions on the blockchain, and thus would not be affected by cryptocurrency transfer fees. For robustness, we also repeat our main regressions with controls for blockchain transaction fees and the hash rate and find that our main results remain unchanged. These tests are excluded from the paper for brevity but are available upon request.

4. Conclusion

We study the carry of crypto futures for bitcoin and ether, i.e., the difference between futures and spot prices. We show that crypto carry reflects a substantial and volatile inconvenience yield associated with holding spot cryptocurrencies relative to futures. Crypto futures offer a unique opportunity to study convenience yields in a setting largely free from many confounding factors present in traditional asset classes, such as storage costs, exogenous supply shocks, or policy interventions. This clean environment allows us to isolate the effects of demand pressure from smaller, potentially less sophisticated investors and to understand how this pressure interacts with the limits to arbitrage arising from regulatory uncertainty or outright restrictions on investment in a new asset class. The resulting dynamics are especially transparent in crypto markets, where the sharp divide between regulated and unregulated exchanges makes segmentation particularly severe.

Our findings reveal that market segmentation driven by regulatory frictions can lead to persistent discrepancies between futures and spot prices, resulting in significant distortions in asset pricing that may initially appear to present risk-free arbitrage opportunities. However, rather than offering investors a "free lunch", our results suggest that crypto carry reflects the tangible impact of these regulatory barriers and frictions. Such impediments to arbitrage have a particularly pronounced effect in crypto derivatives markets, which are dominated by relatively unsophisticated investors. Specifically, the evidence presented in this paper indicates that past crypto price surges coincide with buying pressure from smaller, trend-following investors seeking leveraged upside exposure, which could create feedback loops that exacerbate mispricings. Viewed through this lens, a greater presence of professional investors employing cash-and-carry strategies could serve as a stabilizing force during such booms, fostering healthier and more efficient market dynamics.

In this context, crypto futures serve as a real-time laboratory to understand how market segmentation influences the functioning of markets, asset pricing and the emergence of convenience yields. More broadly, these findings highlight that, in the realm of financial innovation, establishing clear legal and regulatory frameworks early on can help mitigate distortions and promote more sound market dynamics and practices.

References

Angeris, G., T. Chitra, A. Evans, and M. Lorig (2022). A primer on perpetuals. *Working Paper*.

Brunnermeier, M. and L. H. Pedersen (2009). Market liquidity and funding liquidity. *Review of Financial Studies 22*, 2201–2238.

Christin, N., B. R. Routledge, K. Soska, and A. Zetlin-Jones (2022). The crypto carry trade. *Working Paper*.

Duffie, D. (2020). Still the world's safe haven? *Hutchins Center Working Paper #62*.

Frazzini, A. and L. H. Pedersen (2022). Embedded leverage. *The Review of Asset Pricing Studies 12*(1), 1–52.

Gorton, G. and K. G. Rouwenhorst (2006). Facts and fantasies about commodity futures. *Financial Analysts Journal 62*(2), 47–68.

He, S., A. Manela, O. Ross, and V. von Wachter (2022). Fundamentals of perpetual futures. *Working Paper*.

He, Z., S. Nagel, and Z. Song (2022). Treasury inconvenience yields during the COVID-19 crisis. *Journal of Financial Economics 143*, 57–79.

Heston, S. and K. Todorov (2023). The variance and return premiums across assets. *Working paper*.

Hong, H. and M. Yogo (2012). What does futures market interest tell us about the macroeconomy and asset prices? *Journal of Financial Economics 105*(3), 473–490.

Kogan, S., I. Makarov, M. Niessner, and A. Schoar (2023). Are cryptos different? Evidence from retail trading. *Working Paper*.

Koijen, R. S., T. J. Moskowitz, L. H. Pedersen, and E. B. Vrugt (2018). Carry. *Journal of Financial Economics 127*(2), 197–225.

Liu, Y. and A. Tsyvinski (2021). Risks and returns of cryptocurrency. *Review of Financial Studies 34*, 2689–2727.

Makarov, I. and A. Schoar (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics 135*, 293–319.

Makarov, I. and A. Schoar (2021). Blockchain analysis of the bitcoin market. *Working Paper*.

Plantin, G. and H. Shin (2014). Destabilizing carry trades. TSE Working paper.

Schrimpf, A., HS. Shin, and V. Sushko (2020). Leverage and margin spirals in fixed income markets during the covid-19 crisis. *BIS Bulletin 2*.

Todorov, K. (2021). Launch of the first US bitcoin ETF: mechanics, impact, and risks. *BIS Quarterly Review December 2021*.

Figures and Tables

Figure 1: Size of the crypto futures market.
The figure shows the average daily dollar open interest (number of contracts times futures price) across exchanges. Sample: March 2019 to July 2024. The data on crypto derivatives comes from Skew, Coinmetrics and Bloomberg.

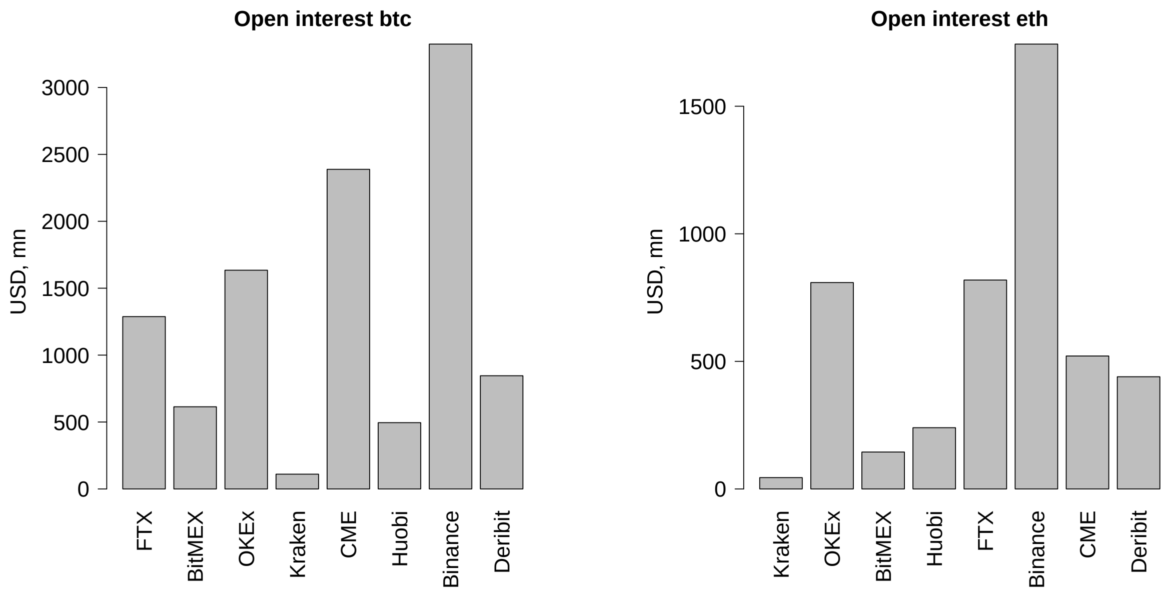

Figure 2: Dynamics of crypto carry.
The figure shows the dynamics of crypto carry for OKEx (top left panel) and CME (top right panel). The lower left-hand panel shows the dynamics of interest rates from Binance, Aave, and 1-month USD LIBOR. The lower right-hand panel depicts fluctuations in the 1-month BTC basis on OKEx and the spot price of bitcoin. Sample: March 2019 to July 2024 for OKEx, August 2020 to July 2024 for CME. The data on crypto derivatives comes from Skew and Coinmetrics.

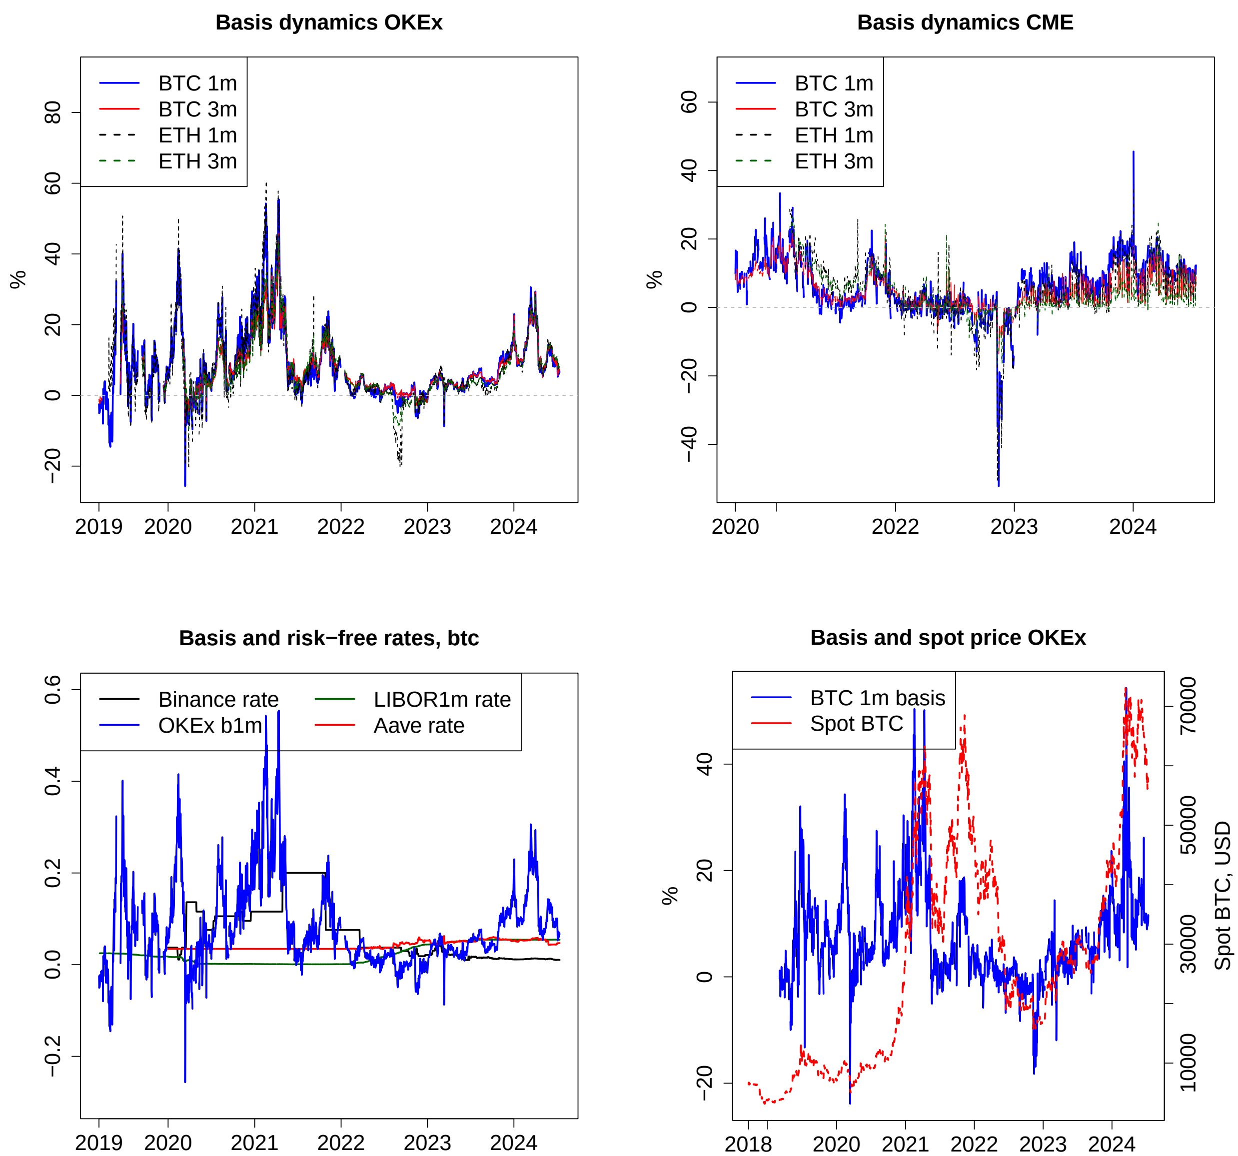

Figure 3: CME positions and regressions.

The left panel shows the dynamics of positions in BTC futures from different types of investors. The right panel shows the estimates of the following regression at the weekly frequency: basis 1-month$_t = \alpha + \beta_D \Delta X_{D,t} + \beta_I \Delta X_{I,t} + \beta_L \Delta X_{L,t} + \beta_N \Delta X_{N,t} + \epsilon_t$, where basis 1-month$_t$ is 1-month BTC futures basis, $\Delta X_{D,t}, \Delta X_{I,t}, \Delta X_{L,t}, \Delta X_{N,t}$ are the changes in positions of dealer intermediaries, institutional investors, leveraged funds, and nonreportables, respectively. All β-s except the one on dealer intermediaries are significant at the 10% level. Reference category is other reportables. The estimate of β_N ("Nonreportable") for instance shows that an increase in the weekly change of nonreportable positions by 1,000 contracts is correlated with 0.2% increase in basis. Sample: March 2019 to July 2024. The data on crypto derivatives comes from Skew and Coinmetrics, the data on investor positions from the CFTC.

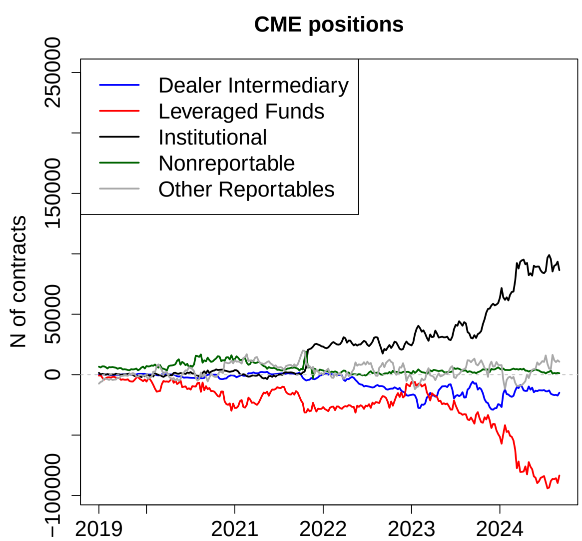

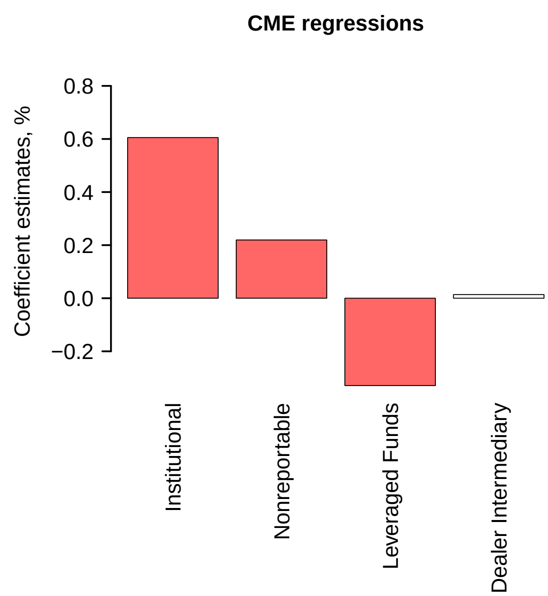

Figure 4: Basis dynamics around the introduction of micro bitcoin futures on the CME. The figure shows the dynamics of 1-month BTC basis on CME (red) and the average basis on OKEx, Huobi, Deribit and Kraken (blue) around the introduction of micro bitcoin futures on CME (marked with dashed vertical line).

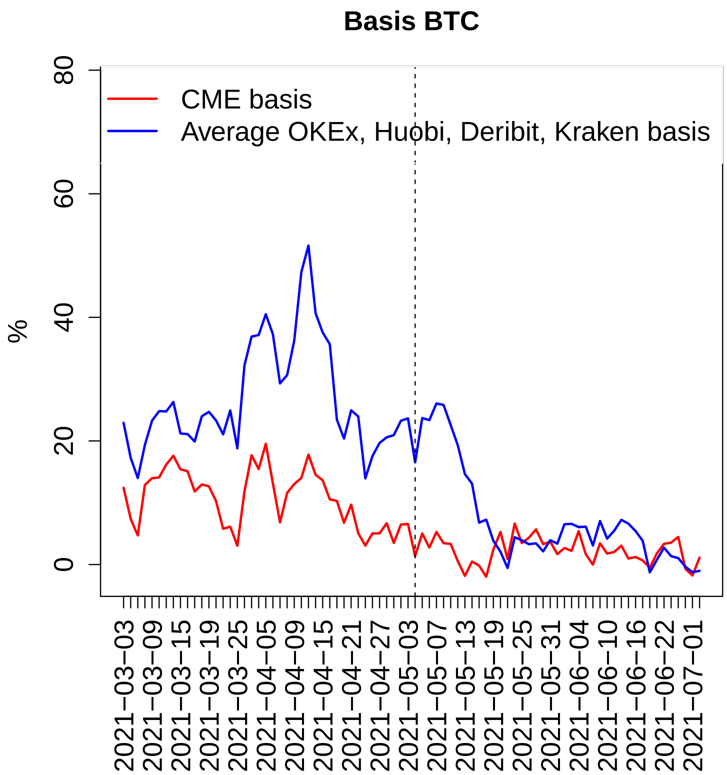

Figure 5: Trading strategy
The figure shows the cumulative returns on the strategy of selling (buying) CME BTC futures when basis is positive (negative) 19 trading days before expiration. 19 trading days capture most of the months in our sample since 28 calendar days before expiration corresponds to 19–22 trading days. The lines correspond to different paths, the dashed horizontal lines illustrate several levels of leverage (L). E.g., with a leverage of 5, the position is bankrupt if the cumulative return before maturity is below -0.2. Sample: March 2019 to July 2024.

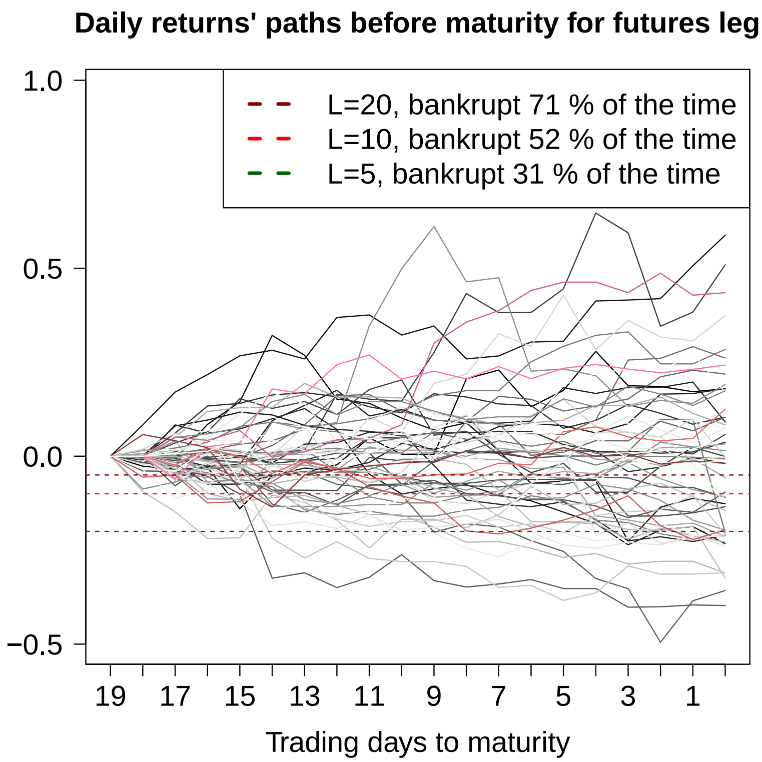

Figure 6: Basis dynamics around the introduction of spot BTC ETF.
The figure shows the dynamics of 1-month BTC basis on CME (red) and the average basis on OKEx, Huobi, Deribit and Kraken (blue) around the introduction of the spot BTC ETF (marked with dashed vertical line).

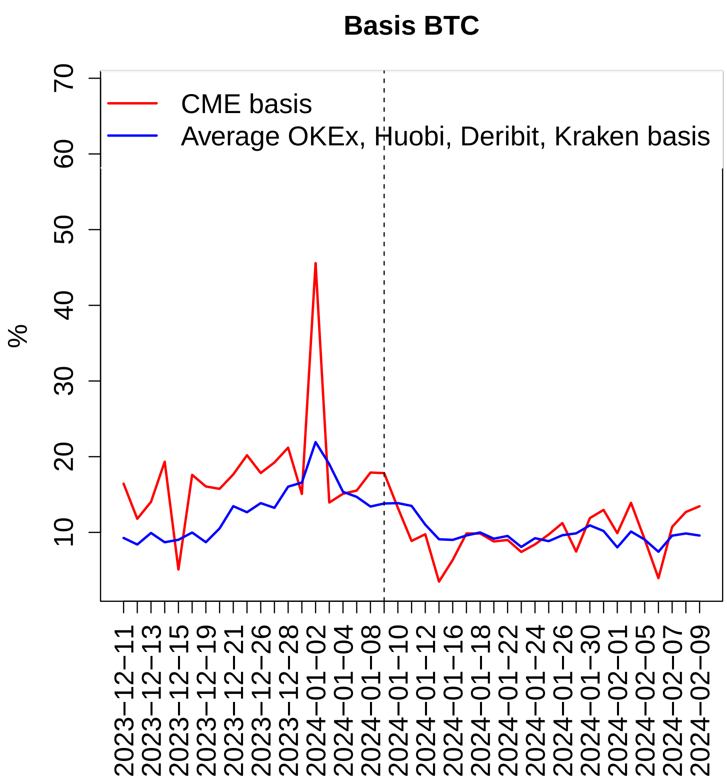

Table 1: Summary statistics for crypto carry.

The table shows the average basis (crypto carry), constant maturity of 1 month or 3 months and in annualized %. For BTC, the sample starts in March 2019 for OKEx and in August 2020 for CME. For ETH, it starts in April 2019 for OKEx and in February 2021 for CME. The final date is July 2024 for all time series. Frequency is daily. The data on crypto derivatives comes from Skew and Coinmetrics.

Panel A: BTC crypto carry				
	OKEx	CME	OKEx	CME
	1 month		3 months	
Mean	8.21	6.41	7.86	5.12
Median	5.62	6.62	5.66	3.76
Std	9.48	8.15	7.30	5.21
Skewness	1.35	-0.95	1.37	0.49
Min	-25.65	-52.2	-7.87	-13.61
Max	55.42	45.57	45.93	22.33
AR(1)	1	1	0.98	0.86
Observations	1866	980	1615	994
Panel B: ETH crypto carry				
	OKEx	CME	OKEx	CME
	1 month		3 months	
Mean	8.21	4.91	7.49	3.67
Median	5.59	4.86	5.06	1.97
Std	10.65	8.77	8.11	5.64
Skewness	1.24	-1.12	1.15	1.06
Min	-20.23	-50.48	-11.21	-16.36
Max	60.69	34.4	44.98	26.42
AR(1)	0.94	1	0.98	1
Observations	1828	869	1604	877

Table 2: Correlations of crypto carry.
Panel A of the table shows daily correlations of the basis within exchange. Panel B reports correlations across exchanges for the 1-month BTC basis. The data on crypto derivatives comes from Skew and Coinmetrics.

Panel A: OKEx correlations			
	BTC, 1 month	BTC, 3 months	ETH, 1 month
BTC, 3 months	0.97		
ETH, 1 month	0.94	0.93	
ETH, 3 months	0.94	0.97	0.97

Panel B: BTC, 1 month correlations				
	Deribit	Kraken	OKEx	Huobi
Kraken	0.93			
OKEx	0.97	0.93		
Huobi	0.95	0.90	0.97	
CME	0.67	0.75	0.66	0.60

Table 3: Crypto carry and interest rate differentials.
The table shows pooled regressions of 1-month BTC basis for all exchanges on exchange fixed effects, time fixed effects, the change in VIX (proxy for risk constraints), and interest rate spreads. The data on crypto derivatives comes from Skew and Coinmetrics. Standard errors in columns 3–5 are double-clustered by exchange and time. Here and in all subsequent regression tables *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels, respectively.

	Basis, 1 month				
	(1)	(2)	(3)	(4)	(5)
ΔVIX			0.002		
			(0.002)		
Binance-LIBOR spread				-0.14***	
				(0.02)	
Aave-LIBOR spread					-0.56**
					(0.08)
Exchange FEs	Yes	No	No	No	No
Time FEs	No	Yes	No	No	No
Observations	7,891	7,891	5,436	5,196	5,181
Adjusted R^2	0.00	0.88	0.001	0.02	0.01

Table 4: Crypto carry, attention by smaller investors, and trend-chasing.
The table shows daily regressions of 1-month BTC basis on measures of smaller investors' attention and trend-chasing. All independent variables in columns (1)–(12) are standardized. The first six columns are for OKEx, the last six for CME. "Google" is an index for Google searches of "BTC", "bitcoin", "bitcoin futures", "bitcoin price", "bitcoin leverage" from Google trends. "$r^{spot}_{t-1w,t}$" is one-week spot BTC return, "$r^{spot}_{t-1m,t}$" is 1-month spot BTC return. In columns (4) – (6), (10) – (12), we add variables capturing price pressure. "Signed volume" refers to futures trading volume multiplied with the sign of the contemporaneous return (-1 for negative return, +1 for positive), "OI" is open interest in the futures market. The data on crypto derivatives comes from Skew and Coinmetrics. Here and in all subsequent regression tables, standard errors are computed using the Newey-West method with automatic bandwidth selection.

	Basis, 1 month$_t$											
	OKEx						CME					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Google$_t$	3.21***			1.85***	1.35**	0.84**	0.82			0.56	0.06	-0.01
	(0.77)			(0.55)	(0.50)	(0.40)	(0.60)			(0.53)	(0.40)	(0.11)
$r^{spot}_{t-1w,t}$		2.28***			0.55**	0.58**		2.44***			0.49	0.80***
		(0.53)			(0.21)	(0.23)		(0.36)			(0.36)	(0.20)
$r^{spot}_{t-1m,t}$			4.13***		3.22***	0.21			4.32***		4.73***	0.85***
			(0.66)		(0.30)	(0.14)			(0.54)		(1.08)	(0.28)
Signed volume$_t$				2.53***	0.35	-0.02				2.35***	-0.71	-0.19
				(0.64)	(0.27)	(0.08)				(0.43)	(0.66)	(0.18)
OI$_t$				2.62***	2.52***	0.25***				2.67***	2.55***	0.56***
				(0.45)	(0.21)	(0.08)				(0.44)	(0.41)	(0.12)
Basis, 1 month$_{t-1}$						0.92***						0.76***
						(0.02)						(0.03)
Intercept	7.89***	8.13***	7.96***	7.94***	8.10***	0.64***	6.02***	6.37***	6.39***	5.50***	5.78***	1.35***
	(0.65)	(0.88)	(0.69)	(0.64)	(0.19)	(0.10)	(0.81)	(0.45)	(0.58)	(0.79)	(0.66)	(0.27)
Observations	1,864	1,864	1,864	1,720	1,717	1,716	980	980	979	980	979	978
Adjusted R^2	0.12	0.06	0.19	0.27	0.35	0.92	0.01	0.08	0.27	0.25	0.37	0.80

Table 5: Difference-in-differences regressions. The table shows DiD regressions around the introduction of micro BTC futures on CME (columns 1 and 2) and the spot BTC ETF introduction (column 3) using a two-months window around the respective event. *Treat* is a dummy equal to 1 for CME, and *Post* is a dummy equal to 1 for the period after the event date, $r_{ETH,t}$ is the daily return on ETH.

Basis, 1 month$_t$			
	Micro BTC futures		Spot ETF
	(1)	(2)	(3)
Treat	-0.18^{***}	-0.18^{***}	0.05^{***}
	(0.01)	(0.01)	(0.001)
Post	-0.18^{***}	-0.18^{***}	-0.03^{***}
	(0.02)	(0.02)	(0.01)
Treat × Post	0.11^{***}	0.11^{***}	-0.05^{***}
	(0.01)	(0.01)	(0.01)
$r_{ETH,t}$		-0.01	
		(0.10)	
Intercept	0.28^{***}	0.28^{***}	0.13^{***}
	(0.02)	(0.02)	(0.01)
Observations	292	292	291
Adjusted R^2	0.53	0.53	0.27

Table 6: Strategy characteristics.

The table shows characteristics of the strategy that sells (buys) CME BTC futures and buys (sells) spot BTC when basis is positive (negative) 28 days before expiration. The first column shows the futures leg of the strategy for January 2018–July 2024 since this is the longest available sample for CME BTC futures. The second column is March 2019–July 2024 (our main sample). The third column shows the spot leg of the strategy for March 2019–July 2024. The returns are monthly. Max drawdown is the maximum value drop after the peak of the variable (e.g., for spot 2019–2024, the return drops from 62% to -40%).

	Futures 2018–2024 (1)	Futures 2019–2024 (2)	Spot 2019–2024 (3)
Mean (%)	2.78	2.00	-0.29
Std (%)	17.81	17.16	17.41
Skewness	0.11	0.34	-0.24
Max drawdown	0.95	0.87	1.02
Observations	80	66	66

Table 7: Crypto carry, contract liquidations and volatility.
The table shows predictive regressions of crypto carry for cumulative buy or sell liquidations, implied volatility and realized volatility for OKEx. The predictive horizon is 1 month. The first two columns present the results for liquidations using the entire sample. In contrast, columns 3 and 4 display the results for a subsample where the basis widens for at least 70% of the days within a 2-week period before expiration and never falls below the initial value at the start of the carry trade in that same period. The construction of this subsample seeks to rule out instances where traders liquidate short positions due to profits made from the widening basis. The data on crypto derivatives comes from Skew and Coinmentrics. All independent variables are standardized. "Buy liq" are cumulative buy liquidations of futures contracts, scaled by open interest, whereas "Sell liq" are sell liquidations, also scaled by open interest. "Ivol" a measure of expected 1-month implied volatility in BTC based on variance swaps, "Rvol" is 1-month realized volatility, "Skew re" is realized 1-month skew.

	Buy liq$_{t+1m}$	Sell liq$_{t+1m}$	Buy liq$_{t+1m}$	Sell liq$_{t+1m}$	Ivol$_{t+1m}$	Rvol$_{t+1m}$
	(1)	(2)	(3)	(4)	(5)	(6)
Basis, 1m$_t$	0.41	2.22***	0.74	2.02***	7.08***	6.27***
	(0.45)	(0.79)	(0.58)	(0.72)	(1.71)	(1.94)
Ivol$_{t+1m}$	2.53***	3.69***	1.75**	3.74***		
	(0.53)	(0.86)	(0.86)	(1.25)		
Skew re$_t$	1.09***	−1.90***	1.04***	−1.57***		
	(0.39)	(0.51)	(0.38)	(0.44)		
OI$_{t+1m}$	0.25	−0.15	−0.31	1.12		
	(0.54)	(0.72)	(0.68)	(0.74)		
Intercept	7.83***	10.84***	8.88***	9.33***	75.38***	50.78***
	(0.45)	(0.61)	(0.53)	(0.52)	(1.67)	(2.86)
Observations	1,619	1,619	211	211	1,712	1,845
Adjusted R^2	0.23	0.29	0.17	0.52	0.11	0.08

Online Appendix for

"Crypto Carry"

A.1. Additional figures and tables

In this Online Appendix, we provide additional tables and figures: summary statistics for carry across more exchanges, correlations of crypto and commodities carry, and a robustness test using a bear market subsample. We also provide some of the main regressions for ETH.

Table A.1: Summary statistics for basis across exchanges.
The table shows crypto basis with constant maturity of 1 month (Panels A and C) or 3 months (B and D), annualized and in %. The table also shows the max leverage for exchanges. Huobi had a max leverage of 125 before June 2021 but then changed it to 5 reportedly. Similarly, Binance and FTX also reduced their leverage in July 2021. Source: Skew, coinmetrics, exchanges websites.

Panel A: BTC, 1 month

	Deribit	Kraken	OKEx	Huobi	CME
Mean	7.49	7.40	8.21	7.49	6.41
Median	5.52	5.30	5.62	5.12	6.62
Std	8.73	9.07	9.48	9.50	8.15
Skewness	1.64	1.16	1.35	1.47	-0.95
Min	-28.72	-23.87	-25.65	-33.22	-52.2
Max	59.89	54.31	55.42	54.28	45.57
AR(1)	1	0.92	1	1	1
Observations	1484	1813	1866	1739	980
Max leverage	100	50	125	5-125	2

Panel B: BTC, 3 months

	BitMEX	Deribit	OKEx	FTX	Huobi	Binance	CME
Mean	6.05	7.3	7.86	8.58	6.21	8.95	5.12
Median	4.47	6.07	5.66	7.07	3.47	7.24	3.76
Std	6.95	6.76	7.3	7.47	7.66	7.56	5.21
Skewness	1.47	1.24	1.37	1.46	1.6	1.46	0.49
Min	-11.34	-6.55	-7.87	-6.88	-9.39	-2.15	-13.61
Max	45.6	45.63	45.93	45.59	44.85	46.76	22.33
AR(1)	1	0.98	0.98	1	0.98	0.98	0.86
Observations	1952	1950	1615	1157	1435	1456	994
Max leverage	285	100	125	20-101	5-125	50-125	2

Panel C: ETH, 1 month

	Kraken	OKEx	Deribit	Huobi	CME
Mean	7.72	8.21	6.41	6.83	4.91
Median	5.47	5.59	4.26	4.48	4.86
Std	10.94	10.65	9.36	10.88	8.77
Skewness	1.08	1.24	1.46	1.29	-1.12
Min	-22.34	-20.23	-22.72	-19.57	-50.48
Max	58.81	60.69	64.98	57.98	34.4
AR(1)	0.92	0.94	1	1	1
Observations	1668	1828	1444	1739	869

Panel D: ETH, 3 months

	Deribit	BitMEX	OKEx	FTX	Huobi	Binance	CME
Mean	7.09	8.42	7.49	8.68	5.85	8.29	3.67
Median	5.35	6.65	5.06	6.58	2.94	5.99	1.97
Std	7.43	10.6	8.11	10.38	8.32	8.35	5.64
Skewness	1.18	1.12	1.15	1.11	1.5	1.12	1.06
Min	-10.66	-15.01	-11.21	-7.88	-7.39	-8.74	-16.36
Max	44.7	49.22	44.98	50.19	44.82	46.97	26.42
AR(1)	1	0.92	0.98	0.98	0.97	0.98	1
Observations	1902	128	1604	770	1435	1421	877

Table A.2: Comparison of crypto carry with carry on commodities, equities, and U.S. Treasuries.

The table shows summary statistics of futures carry for the S&P 500 index, 10-year U.S. Treasuries, and commodities. The last row illustrates the daily correlations of CME one-month BTC carry with that of the S&P 500 index, 10-year U.S. Treasuries, and commodities. All correlations are insignificant at the 5% level. Source: Bloomberg, CME. Sample is from February 2018 to July 2024.

	S&P 500	U.S. Treasuries	Oil	Gold
Mean carry (%)	0.71	0.58	-0.46	0.29
Median carry (%)	0.46	0.19	0.00	-0.04
Std carry(%)	2.47	1.12	4.12	4.27
Skewness	1.51	0.82	-8.49	1.52
Min (%)	-13.95	-1.88	-81.36	-23.81
Max (%)	18.72	4.15	13.52	44.52
AR(1)	0.38	0.98	0.41	0.33
Observations	1616	1616	1564	1596
Correlation	0.12	-0.05	-0.02	0.05

Table A.3: Crypto carry, attention by smaller investors, and trend-chasing in a bear market.
The table shows daily regressions of 1-month BTC basis on measures of smaller investors' attention and trend-chasing. All independent variables in columns (1) – (12) are standardized. The first six columns are for OKEx, the last six for CME. "Google" is an index for Google searches of "BTC", "bitcoin", "bitcoin futures", "bitcoin price", "bitcoin leverage" from Google trends. "$r^{spot}_{t-1w,t}$" is one-week spot BTC return, "$r^{spot}_{t-1m,t}$" is 1-month spot BTC return. In columns (4) – (6), (10) – (12), we add variables capturing price pressure. "Signed volume" refers to futures trading volume multiplied with the sign of the contemporaneous return (-1 for negative return, +1 for positive), "OI" is open interest in the futures market. Standard errors are computed using the Newey-West method with automatic bandwidth selection. The bear market sample is November 2021-January 2023.

	Basis, 1 month$_t$											
	OKEx						CME					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Google$_t$	0.97*			0.95*	1.46**	0.34*	1.22*			0.59*	1.20**	0.51
	(0.55)			(0.54)	(0.63)	(0.18)	(0.66)			(0.32)	(0.56)	(0.37)
$r^{spot}_{t-1w,t}$		0.71**			0.89**	0.41*		1.04**			0.72*	0.15
		(0.29)			(0.37)	(0.22)		(0.49)			(0.41)	(0.22)
$r^{spot}_{t-1m,t}$			0.90*		2.28***	1.37*			1.60**		2.11*	1.70*
			(0.53)		(0.64)	(0.76)			(0.66)		(1.18)	(0.79)
Signed volume$_t$				-0.18	-1.85***	-0.55				0.55*	-1.18**	-1.08**
				(0.51)	(0.69)	(0.50)				(0.31)	(0.59)	(0.50)
OI$_t$				0.59	-0.32	0.21				3.70***	3.15***	2.87***
				(1.02)	(0.91)	(0.85)				(0.43)	(0.46)	(0.45)
Basis, 1 month$_{t-1}$						0.89***						0.85***
						(0.03)						(0.02)
Intercept	2.24***	2.25***	2.25***	2.20***	2.31***	0.71***	-1.14	-0.97	-0.94	-1.04*	-1.16	0.67**
	(0.49)	(0.39)	(0.81)	(0.66)	(0.62)	(0.09)	(0.92)	(0.96)	(0.75)	(0.60)	(0.80)	(0.31)
Observations	419	419	419	419	419	418	302	302	302	302	302	301
Adjusted R^2	0.04	0.02	0.04	0.05	0.22	0.89	0.02	0.02	0.04	0.24	0.29	0.87

Table A.4: Crypto carry, attention by smaller investors, and trend-chasing for ETH.
The table shows daily regressions of one-month ETH basis on measures of smaller investors' attention and trend-chasing. All independent variables in columns (1) – (12) are standardized. The first six columns are for OKEx, the last six for CME. "Google" is an index for Google searches of "ETH", "ethereum", "ethereum futures", "ethereum price", "ethereum leverage" from Google trends. "$r^{spot}_{t-1w,t}$" is one-week spot ETH return, "$r^{spot}_{t-1m,t}$" is one-month spot ETH return. In columns (4) – (6), (10) – (12), we add variables capturing price pressure. "Signed volume" refers to futures trading volume multiplied with the sign of the contemporaneous return (-1 for negative return, +1 for positive), "OI" is open interest in the futures market. The data on crypto derivatives comes from Skew and Coinmetrics. Standard errors are computed using the Newey-West method with automatic bandwidth selection.

	Basis, 1 month$_t$											
	OKEx						CME					
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
Google$_t$	3.30***			5.20***	4.30***	0.41***	2.67***			3.26***	3.09***	0.90***
	(0.89)			(0.62)	(0.68)	(0.14)	(0.54)			(0.57)	(0.54)	(0.19)
$r^{spot}_{t-1w,t}$		3.34***			0.95**	0.54***		1.68***			0.67	0.62**
		(0.52)			(0.46)	(0.10)		(0.51)			(0.43)	(0.28)
$r^{spot}_{t-1m,t}$			5.70***		3.74***	0.25**			3.97***		2.31***	0.47**
			(0.72)		(0.69)	(0.12)			(0.90)		(0.69)	(0.22)
Signed volume$_t$				-7.67***	-5.70***	-0.62***				2.35***	-0.71	-1.19***
				(0.96)	(1.12)	(0.16)				(0.43)	(0.66)	(0.32)
OI$_t$				4.70***	3.61***	0.43***				3.73***	3.21***	0.89***
				(0.81)	(0.10)	(0.08)				(0.72)	(0.63)	(0.23)
Basis, 1 month$_{t-1}$						0.90***						0.73***
						(0.02)						(0.03)
Intercept	7.89***	8.13***	8.02***	8.85***	8.37***	0.84***	3.63***	4.97***	5.22***	6.76***	6.55***	1.69***
	(0.65)	(0.95)	(0.77)	(0.64)	(0.63)	(0.14)	(0.84)	(0.71)	(0.86)	(1.01)	(0.93)	(0.34)
Observations	1,826	1,826	1,826	1,720	1,717	1,704	869	869	869	869	869	688
Adjusted R^2	0.12	0.06	0.19	0.27	0.43	0.90	0.09	0.03	0.14	0.30	0.38	0.73

Previous volumes in this series

1086 April 2023	CBDC policies in open economies	Michael Kumhof, Marco Pinchetti, Phurichai Rungcharoenkitkul and Andrej Sokol
1085 March 2023	Supervisory policy stimulus: evidence from the Euro area dividend recommendation	Ernest Dautović, Leonardo Gambacorta and Alessio Reghezza
1084 March 2023	BigTech credit and monetary policy transmission: micro-level evidence from China	Yiping Huang, Xiang Li, Han Qiu and Changhua Yu
1083 March 2023	Commodity prices and the US Dollar	Daniel M Rees
1082 March 2023	Public debt and household inflation expectations	Francesco Grigoli and Damiano Sandri
1081 March 2023	What happens to emerging market economies when US yields go up?	Julián Caballero and Christian Upper
1080 March 2023	Did interest rate guidance in emerging markets work?	Julián Caballero and Blaise Gadanecz
1079 March 2023	Volume dynamics around FOMC announcements	Xingyu Sonya Zhu
1078 March 2023	Greenhouse gas emissions and bank lending	Koji Takahashi and Junnosuke Shino
1077 February 2023	Understanding post-COVID inflation dynamics	Martín Harding, Jesper Lindé and Mathias Trabandt
1076 February 2023	The shape of business cycles: a cross-country analysis of Friedman's plucking theory	Emanuel Kohlscheen, Richhild Moessner and Daniel M Rees
1075 February 2023	Overcoming original sin: insights from a new dataset	Mert Onen, Hyun Song Shin and Goetz von Peter
1074 February 2023	Non-bank lending during crises	Iñaki Aldasoro, Sebastian Doerr and Haonan Zhou
1073 February 2023	Constrained liquidity provision in currency markets	Wenqian Huang, Angelo Ranaldo, Andreas Schrimpf, and Fabricius Somogyi

All volumes are available on our website www.bis.org.