Perpetual Futures and Basis Risk:

Evidence from Cryptocurrency

Preliminary

Will Gornall[*] Martin Rinaldi[*] Yizhou Xiao[†]

May 19, 2025

Abstract

We study the design of futures contracts using the cryptocurrency market as a natural laboratory. In these markets, constrained arbitrage capital and volatile speculative demand frequently drive futures prices away from their underlying assets. Perpetual futures emerged in response, using small, frequent funding payments to keep prices aligned. We document that these contracts now dominate trading volume, enhance liquidity, and reduce extreme price dislocations. A tractable model of capital-constrained arbitrage explains these findings. Finally, we argue that perpetual futures may enhance financial stability by improving crisis-time liquidity and substantially reducing the drawdowns of arbitrage strategies.

[*]University of British Columbia, Sauder School of Business.

[†]The Chinese University of Hong Kong (CUHK) Business School.

We are grateful to the faculty at Carnegie Mellon University, CUHK, McGill University, and UBC, and to Yoontae Jeon for their valuable insights, helpful comments, and suggestions. We would also like to acknowledge the support of Aquanow, the MITACS Accelerate program, and the SSHRC who provided funding for this research.

1. Introduction

Cryptocurrency exchanges have evolved into massive and largely unregulated financial markets. Leading venues such as Binance report hundreds of billions of dollars in annual trading volume and are valued in the tens of billions of dollars.[1] Despite their scale, cryptocurrency markets remain highly volatile and fragile. Whereas the US stock market has experienced only four daily drops exceeding 10% over the past century, cryptocurrency markets have recorded more than fifty such declines in just the past decade.[2] The fragility of this asset class is further evidenced by the frequent failures of major exchanges, including Mt. Gox and FTX, and the collapse of cryptocurrencies themselves, such as Terra/LUNA and Mantra.

Despite this volatile environment, speculators eager for leverage have made cryptocurrency futures contracts extremely popular. Most cryptocurrency trading does not take place on a blockchain. Instead, trading volume is concentrated on centralized exchanges like Binance and takes place using *perpetual futures* contracts.

Unlike standard futures, which converge to the spot price via a single transfer at maturity, perpetual futures enforce convergence through frequent, small payments between long and short positions. If the futures price is below the spot, shorts pay longs; if it is above, longs pay shorts. On BitMEX, for example, these funding payments occur every eight hours and are based on the average price difference over that interval. This structure imposes persistent convergence pressure, keeping perpetual futures closely aligned with their underlying, and as a result the actual funding payments are generally small.

We use data on perpetual futures and standard quarterly futures from cryptocurrency markets to establish several empirical facts. First, perpetual futures rapidly gained popu-

[1] For instance, Binance routinely reports daily volumes exceeding $30 billion, raised $2 billion from a UAE sovereign wealth fund in early 2025, and maintains an associated token (BNB) with a market capitalization of $80 billion.

[2] Author's calculations using daily Coinbase data.

larity across many venues. These contracts were introduced by BitMEX in May 2016, as the end result of an evolution from quarterly to monthly to weekly and even daily futures. This contractual form quickly came to dominate BitMEX and has since displaced standard futures as the dominant instrument for cryptocurrency derivatives on every major non-U.S. venue, with U.S. regulators considering allowing this form.

Second, perpetual futures are more liquid than quarterly futures, with spreads that are 49–83% tighter than comparable quarterly futures. Although this liquidity advantage partly reflects their overwhelming popularity, our analysis shows that perpetual futures offered superior liquidity even immediately following their introduction, suggesting inherent structural advantages. Further, we show that standard futures are less liquid farther from maturity.

Third, we document a shift in trader liquidations from quarterly to perpetual futures. Perpetual futures exhibit frequent liquidations throughout, whereas quarterly futures show frequent liquidations prior to the introduction of perpetual futures and fewer afterward. This is consistent with higher-risk traders showing a stronger preference for perpetual futures.

Fourth, we find that perpetual futures exhibit lower risk than quarterly futures during large spot price movements. Quarterly futures prices decline 8–10% more than spot prices during large market events, while perpetual futures decline just 3% more than spot prices. The tendency of quarterly futures to overshoot spot price movements makes them riskier than perpetual futures.

We emphasize that such a divergence between futures and spot prices is particularly hazardous for arbitrageurs engaged in futures-spot-basis trades. These traders buy the spot asset and hedge that by selling futures (or vice versa) using leverage, a strategy that makes them extremely vulnerable to diverging spot and futures prices. Perpetual futures show drawdowns that are about one-third the size of those shown by quarterly futures; moreover, the perpetual futures drawdowns quickly mean revert while the quarterly futures drawdowns persist. These results have financial stability implications as futures-spot-basis arbitrage is a key stabilizer in many markets, including U.S. Treasuries, equity indices, and non-crypto currencies.

Finally, we show that perpetual futures' liquidity advantage is even larger during times of market volatility. Although both types of futures become less liquid around large market movements, quarterly futures spreads increase much more than perpetual futures spreads.

We interpret these findings using a tractable theoretical model in which futures contracts are traded by market makers, liquidity traders, and long-term hedgers. Liquidity traders and market makers are risk-averse to mark-to-market trading losses, whereas long-term hedgers seek to minimize the risk of a timed hedging exposure. Time-varying liquidity trader demand drives futures price away from the spot price. Contracts with shorter maturities or continuous funding payments (like perpetuals) better track the spot price, reducing risk. That risk reduction increases trading volumes and welfare for both the market maker and the liquidity trader.

Perpetual futures can also benefit hedgers. By lowering risk, these contracts reduce the compensation demanded by market makers, lowering the hedger's costs. Standard futures can perfectly hedge exposures that occur at their maturities, but exposures with other timings expose hedgers to the risk that the futures price may differ from the spot price when contracts are sold early (if they expire after the hedging need) or rolled over (if they expire before). We show that if hedging needs are randomly assigned and not bunched to match contract expiration, hedger welfare is better with a perpetual future than any standard futures.

A futures price that closely tracks the spot price is well suited for some markets but poorly suited to others. For example, a perpetual future that tracks the (low) price of natural gas during summer does not allow a utility to hedge winter consumption. This mismatch helps explain why trading in standard futures on physical commodities and non-investible financial assets often spreads across multiple maturities. In contrast, as Fett and Haynes

(2017) document, trading in equity index and U.S. Treasury futures is heavily concentrated on a single maturity. This pattern suggests that market participants are not hedging date-specific exposures and instead seek general market exposure with low collateral demands, operational simplicity, no counterparty risk, and high liquidity. As such, perpetual futures are best suited for tradable financial underlyings.

Our study makes a novel contribution by demonstrating how perpetual futures represent an evolutionary adaptation in contract design that specifically addresses the unique volatility and capital constraints of cryptocurrency markets. This work is the first to comprehensively analyze how this contractual innovation enhances market stability and liquidity by reducing the risk of the futures price diverging from the spot price.

These results link to several lines of prior work. First, two early papers discuss perpetual futures in traditional financial markets. Gehr (1988) documents that these contracts were originally used by the Chinese Gold and Silver Exchange of Hong Kong and Shiller (1993) proposes offering these futures on illiquid assets as a way to better match hedger needs and encourage the production of long-term information. Following these early contributions, the literature remained largely silent on perpetual contracts until their resurgence in cryptocurrency markets. In the cryptocurrency context, the mechanics and pricing of perpetual futures have been discussed by Ackerer et al. (2024) and He et al. (2022), with Deng et al. (2020) additionally constructing optimal hedging strategies in the context of exchange liquidations.

A second group of papers document persistent inefficiencies in cryptocurrency markets. Makarov and Schoar (2020) analyze cross-exchange arbitrage and show that pricing deviations in cryptocurrency markets reflect constraints on arbitrage capital and market fragmentation. Focusing on cryptocurrency derivatives, Christin et al. (2022) articulate an arbitrage strategy using perpetual futures and link its outsized returns to speculator demand. Schmeling et al. (2023) similarly describe futures-spot-basis arbitrage, ascribing its cause to outsized retail demand and a scarcity of arbitrage capital. He et al. (2022) also describe a speculative strategy to exploit demand imbalances. Alexander et al. (2024) document significant—though declining—inefficiencies in cryptocurrency options markets. Ruan and Streltsov (2022) show empirically that perpetual futures' funding settlements increase adverse selection risk in spot markets, leading to higher volumes but wider spreads. De Blasis and Webb (2022) document seasonality and exchange opening effects in perpetual futures liquidity. Alexander et al. (2021) show that perpetual futures improve spot market price discovery and argue that perpetual futures serve an effective hedge against spot BTC price movements.

Our insights on cryptocurrency futures market structure connect to broader themes in financial economics, notably the literature on limits to arbitrage. The noise trader risk channel we focus on builds on work by De Long et al. (1990) and Foucault et al. (2011). We build on Shleifer and Vishny (1997a), extended by Brunnermeier and Pedersen (2009), and applied by Acharya et al. (2013) to commodity futures markets under a relative scarcity of arbitrage capital. Our focus on futures tenor and the risk of price divergence leverages Rutledge (1976) and Castelino and Francis (1982) who report that divergence risk increases with time to maturity. Our paper is perhaps closest to Hazelkorn et al. (2020), who link the difference between spot and futures prices to dealer financing costs in traditional financial markets, and to Figlewski (1984), who examine the pricing gap between equity index futures and their underlyings and link this to limits to arbitrage.

Finally, our paper builds on the literature on contractual innovation and evolution. As Pagano and Roell (1990) and Catalini and Gans (2020) argue, frictions and institutional weaknesses can spur innovation–in this case, the crisis-heavy and arbitrage-capital-light worlds of cryptocurrency futures evolved towards contracts that derisk arbitrage. By offering a new form of risk transfer, these contracts completed the futures market (Duffie and Rahi, 1995; Allen and Gale, 1994). This innovation reduced the need for arbitrage capital (Shen et al., 2014) and addressed unmet trader demands (Tufano, 2003) in this novel market (Miller, 1986). Although completing markets may improve welfare through better risk sharing and reduced hedging costs (Allen and Gale, 1994; Duffie and Rahi, 1995), they can also spur destabilizing speculation or amplify liquidity shocks, particularly when arbitrage capital is limited (Allen and Gorton, 1993; Brunnermeier and Pedersen, 2009; Shleifer and Vishny, 1997b).

2. Futures trading under arbitrage constraints

2.1. The structure of perpetual futures

Traditional futures contracts, which we refer to as *standard futures*, settle via cash or physical delivery at a predetermined maturity date. Prior to maturity, traders are credited (or debited) based on changes in the futures price. The eventual settlement of these futures causes the futures price converges to the price of its underlying. For futures with financial underlyings, arbitrage typically keeps the spot price in line even prior to maturity. If the futures price trades 2% above the spot price, an arbitrageur can buy the spot asset and short the futures contract and earn a profit of 2% at the futures settlement (minus financing costs and dividends, if any).

In normal conditions, this may suffice to keep prices aligned. But when arbitrage capital is scarce or maturities are long, these forces may be too weak to constrain futures prices. Shortening maturities increases arbitrage pressures, but extremely short maturities make it challenging to hold long-term positions.

Perpetual futures eliminate expiry dates entirely and instead, they maintain price alignment through regular *funding payments* between long and short positions. In an idealized form, the long pays the short at a continuous rate:

$$FundingPayment_t = \delta(P_t - S_t) = \delta B_t, \tag{1}$$

where δ is a scaling parameter, P_t is the futures price, S_t is the spot price, and $B_t = P_t - S_t$ is the basis. When the futures price exceeds the spot price, the long side pays the short side, and vice versa. These expected payments create strong price pressure: any divergence triggers an opposing payment stream, pulling futures and spot prices together.

Larger values of δ create stronger convergence incentives, and so by adjusting δ, exchanges can create stronger or weaker convergence pressures. Standard perpetual futures formulas translate a 2% price dislocation into an arbitrage profit of at least 2% per day–a far stronger incentive than 2% paid only at the maturity of a standard future that might be months away for a quarterly future. In practice, the prospect of these payments means futures prices stay close to spot prices, and funding payments are typically small.

Beyond varying δ, exchanges implement variations of Equation (1) to suit their own trading environments. Differences include how P_t and S_t are measured, the timing and averaging of funding calculations, and modifications to reduce noise or prevent manipulation.[3]

2.2. A model of futures contract structure under arbitrage frictions

We model the equilibrium design and pricing of futures contracts in an environment with constrained arbitrage capital and heterogeneous agents.

Our model begins with the standard assumption that the underlying asset price S_t follows a geometric diffusion process:

$$dS_t = rdt + \sigma_S dY_t^S, \tag{2}$$

where $r \geq 0$ denotes the constant risk-free rate, and Y_t^S is a standard Brownian motion.

We define futures contracts using two parameters: an expiration date T and a funding

[3]For example, the Deribit BTC-USD perpetual uses a mark price for P_t, applies an eight-hour funding period, shrinks the resulting payment by up to $0.003125\% \times S_t$, and caps it at $0.0625\% \times S_t$. Kraken uses similar bounds but sets δ to target a one-day payout, bases the payment on the prior clock hour's average difference between a reference future and the spot price, and adjusts for volatility. BitMEX weights the price difference over the previous eight hours, incorporates bounds, adjusts for imputed borrowing costs, and applies a bid-ask-spread-based dampening factor. He et al. (2022) discuss the mechanics of these contracts.

rate δ. Standard futures are characterized by no funding $\delta = 0$ and expiration at a fixed T, whereas perpetual futures have positive funding $\delta > 0$ and no expiration, $T = \infty$. When a contract has a finite maturity, it settles at expiration for the spot asset price S_T.

Let P_t be the price of the futures contract on that asset and so $B_t = P_t - S_t$ is the contract's basis. The contract is in zero net supply and is traded by three types of agents: a (normalized) unit mass of long-term hedgers (H), mass n^S of liquidity traders (denoted S), and mass n^M of market makers (M). Hedgers seek long-term exposure, market makers provide liquidity but face inventory risk, and liquidity traders (speculators) derive time-varying utility from futures exposure, capturing demand shocks or belief heterogeneity. Each agent i holds x_t^i units of the contract and their wealth, W_i, evolves according to

$$dW_t^i = x_t^i dP_t - x_t^i \delta dt. \tag{3}$$

Liquidity traders derive an additional benefit from futures exposure, drawing on evidence of unbalanced speculative demand (Christin et al., 2022; Schmeling et al., 2023). Specifically, they receive utility flow $x^i G_t^S dt$, where G_t^S is the liquidity trader demand factor. This demand factor could represent time-varying hedging needs, liquidity constraints, or sentiment.[4]

We assume G_t^S follows an Ornstein-Uhlenbeck process, a continuous-time analogue of an AR(1) process:

$$dG_t^S = -\psi G_t^S dt + \sigma_G dY_t^G, \tag{4}$$

where Y_t^G is a standard Brownian motion independent of Y_t^S. The parameter $\psi > 0$ governs the speed at which the liquidity trader demand factor reverts to zero, whereas the parameter σ_G governs its volatility.

[4] An equivalent interpretation of G_t^S is as a belief wedge: liquidity traders hold heterogeneous beliefs about the asset's drift. In particular, they believe the asset follows $dS_t = (G_t^S + r)dt + \sigma_S dY_t^S$. Under that interpretation, G_t^S is the disagreement between market makers and liquidity traders about the asset's drift rate. This dual interpretation is standard in models of speculative trading under heterogeneous beliefs (Harrison and Kreps, 1978).

Liquidity traders and market makers have mean-variance preferences over instantaneous changes in wealth:

$$dW_t^i - rW_t^i + G_t^S x^i \mathbb{I}[i = S]dt - \frac{1}{2}\gamma^i \text{Var}(dW_t^i). \tag{5}$$

The variance penalty reflects either risk aversion or inventory holding costs, consistent with the fragmented nature of cryptocurrency exchanges and constrained arbitrage capital (Garleanu and Pedersen, 2011; Makarov and Schoar, 2020). We use γ^M to denote the risk aversion parameter for market makers and γ^S for liquidity traders.

These preferences and the variation in the liquidity demand factor drive trade between liquidity traders and market makers. Long-term hedgers do not trade and instead have constant hedging demand x^H.

We conjecture that the futures basis B depends on time t and the liquidity demand factor G_t^S. Following standard assumptions in continuous-time analysis, we take $B = B(G_t^S, t)$ to be differentiable in t and twice continuously differentiable in G_t^S.

Given that setup, we define an equilibrium as follows:

Definition 2.1 (Equilibrium)**.** An equilibrium consists of a futures basis function $B(G_t^S, t)$, liquidity trader strategies x_t^S, and market maker strategies x_t^M, such that the following conditions hold:

1. Given $B(G_t^S, t)$, market makers choose their futures contract positions x_t^M to maximize their utilities;

2. Given $B(G_t^S, t)$, liquidity traders choose their futures contract positions x_t^S to maximize their utilities; and

3. The market clears: $x^H + n^M x_t^M + n^S x_t^S = 0$ for $\forall t \in [0, T]$.

The following proposition characterizes the structure of a linear equilibrium in this setting. All proofs are provided in Appendix A.

Proposition 2.2. *There exists an equilibrium in which the futures basis is a linear function of the liquidity trader demand factor G_t^S and long-term hedger demand x^H:*

$$B = k(t)G_t^S + j(t)x^H. \tag{6}$$

In the proof, we derive closed-form expressions for the coefficients $k(t)$ and $j(t)$. The **liquidity trader price coefficient** is given by

$$k(t) = \underbrace{\frac{n^S/\gamma^S}{n^S/\gamma^S + n^M/\gamma^M}}_{\text{Proportional risk-aversion}} \underbrace{\frac{1 - e^{(\psi+\delta)(t-T)}}{\psi + \delta}}_{\text{Liquidity factor risk}} \tag{7}$$

and the **long-term hedger price coefficient** by

$$j(t) = \underbrace{\frac{1}{n^S/\gamma^S + n^M/\gamma^M}}_{\text{Aggregate risk-aversion}} \underbrace{\left(\sigma_G^2 k^2(t) + \sigma_S^2\right) \frac{1 - e^{\delta(t-T)}}{\delta}}_{\text{Hedger demand risk sensitivity}}. \tag{8}$$

The comparative statics of the price coefficients are intuitive. A shorter time to maturity $(T - t)$, a higher funding rate (δ), faster mean reversion in liquidity trader demand (ψ), lower market maker risk aversion (γ^M), or a greater number of market makers (n^M) all attenuate both the liquidity trader and hedger price coefficients. Higher liquidity trader risk aversion (γ^S) or a smaller mass of liquidity traders (n^S) reduces the liquidity trader price coefficient and has an indeterminate effect on the hedger price coefficient: while liquidity traders create price risk for market makers, they also absorb some of the hedger demand. Lower volatility of the underlying asset (σ_S^2) or of liquidity demand shocks (σ_G^2) reduces the hedger price coefficient but does not affect the liquidity trader price coefficient, since both types of trading agents are averse to this volatility. In equilibrium, this volatility affects the trading volume between liquidity traders and market makers but does not affect the price.

In the case of standard futures, $j(t)$ is replaced by its limit as $\delta \to 0$, yielding

$$B_t = \frac{n^S/\gamma^S}{n^S/\gamma^S + n^M/\gamma^M} \frac{1-e^{\psi(t-T)}}{\psi} G_t^S + \frac{1}{n^S/\gamma^S + n^M/\gamma^M} \left(\sigma_G^2 k^2(t) + \sigma_S^2\right)(t-T)x^H. \quad (9)$$

For the standard future, the liquidity trader and long-term hedger price coefficients both monotonically converge to zero as $t \to T$ and the future approaches expiration.

For perpetual futures, we take the limit as $T \to \infty$, yielding

$$B_t = \frac{n^S/\gamma^S}{n^S/\gamma^S + n^M/\gamma^M} \frac{1}{\psi+\delta} G_t^S + \frac{1}{n^S/\gamma^S + n^M/\gamma^M} \left(\sigma_G^2 k^2(t) + \sigma_S^2\right) \frac{1}{\delta} x^H. \quad (10)$$

Because perpetual contracts have no expiration date, this basis does not converge to zero at a fixed time. However, the funding rate δ disciplines deviations from the spot price. As shown in the expression, a sufficiently large δ can shrink the basis arbitrarily close to zero.

For the general contract, the equilibrium liquidity trader positions are given by

$$x_t^S = \frac{1}{n^S} \frac{1}{\gamma^M/n^M + \gamma^S/n^S} \frac{1}{\sigma_G^2 k^2(t) + \sigma_S^2} G_t^S - \frac{1}{n^S} \frac{\gamma^M/n^M}{\gamma^M/n^M + \gamma^S/n^S} x^H \quad (11)$$

and the market maker positions by

$$x_t^M = -\frac{1}{n^M} \frac{1}{\gamma^M/n^M + \gamma^S/n^S} \frac{1}{\sigma_G^2 k^2 + \sigma_S^2} G_t^S - \frac{1}{n^M} \frac{\gamma^S/n^S}{\gamma^M/n^M + \gamma^S/n^S} x^H. \quad (12)$$

The first terms in each expression capture trading between liquidity traders and market makers. It is driven by the incentive to trade (G_t^S), reduced by total risk ($\sigma_G^2 k^2(t) + \sigma_S^2$) and the aggregate risk aversion of the two groups ($\gamma^M/n^M + \gamma^S/n^S$). The second terms reflect the hedging demand from long-term hedgers, which is allocated between liquidity traders and market makers according to their risk tolerance.

Because hedging demand x^H is constant, trading is driven by market makers and liquidity traders and instantaneous trading volume is therefore proportional to

$$\frac{1}{\gamma^M/n^M + \gamma^S/n^S} \; \frac{1}{\sigma_G^2 k^2(t) + \sigma_S^2} \sigma_G^2. \tag{13}$$

Lower basis volatility increases risk-bearing capacity for both market makers and speculators, and so many of the comparative statics of $k(t)$ carry over: trading volume is higher when market makers are less risk averse (γ^M is lower), more numerous (n^M is higher), or when the contract has a shorter time to maturity (T is lower) or a higher funding rate (δ), or when there is faster mean reversion in liquidity trader demand (ψ). Although higher volatility in liquidity trader demand (σ_G) increases price impact, the net effect of increasing that volatility is a rise in trading activity because fluctuations in G_t^S directly drive trades.

The effect of liquidity trader risk bearing is more nuanced. A lower risk aversion (γ^S) or larger mass (n^S) increases demand for trade, but also raises basis volatility, deterring participation. When liquidity traders are sufficiently risk averse (or few in number), the first effect dominates, and reducing γ^S (or increasing n^S) increases trading volume. However, the second effect can also prevail, with increasing liquidity trader risk bearing raising volatility to the point where market makers are deterred from participating–an effect akin to the limits-of-arbitrage mechanism described in Shleifer and Vishny (1997a).

We now present our main result, which is that both liquidity traders and market makers prefer well designed perpetual futures over delivery contracts.

Corollary 2.3. *The utilities of both liquidity traders and market makers are increasing in the funding rate δ. For any standard futures contract with expiration T, there exists a threshold $\overline{\delta}$ such that liquidity traders and market makers prefer the perpetual futures contract if and only if $\delta \geq \overline{\delta}$.*

We conclude our model by analyzing the preferences of long-term hedgers. Hedgers are born at a constant rate x^H/l, live for a fixed duration l, and maintain unit exposure to futures throughout their lifetime. Each hedger has mean-variance preferences over the difference between their wealth and the spot price of the asset at the time of death.

A sufficiently risk-averse hedger prefers a futures contract that matures at exactly their exit date. However, a single futures contract cannot support every hedger's specific needs. We assume that there is a single tradable future at all times. In particular, we assume that for futures contracts with fixed maturities, a futures contract maturing at time $(n+1)T$ is available for trade between nT and $(n+1)T$ for integer n. This stylized structure reflects the fact that the liquidity on CME equity-index and U.S. Treasury futures concentrates in the nearest-to-expiry contract (Fett and Haynes, 2017).

This setup mirrors equity-index and U.S. Treasury futures, where there is only a single, near-expiry contract is liquid (Fett and Haynes, 2017).

This series of contracts allows the hedger to maintain constant exposure by rolling from one contract to the next as older contracts expire, which is common practice on equity-index and U.S. Treasury futures markets. For instance, a hedger born at time $\tau < T$ who lives until time τ with $T < \tau + l < 2T$ will at time t purchase a futures contract maturing at T, at time T roll that future by closing it for S_T and purchase a future maturing at $2T$, and at time $\tau + l$ close that second future. For standard futures, the hedger's terminal wealth is as follows:

$$-P_{T,\tau} + \underbrace{S_T - P_{2T,T}}_{\text{Futures roll}} + \underbrace{P_{2T,\tau+l} - S_{\tau+l}}_{\text{Early winding}} + S_{\tau+l}, \tag{14}$$

where $P_{nT,s}$ denotes the price at time s of a futures contract maturing at nT.

The following corollary shows that appropriately set perpetual futures offer higher aggregate welfare than any such series of delivery futures.

Corollary 2.4. *Consider an equilibrium where a series of standard futures are offered with the contract maturing at time $(n+1)T$ being tradable between nT and $(n+1)T$ for integer n. There exists a single perpetual futures contract that delivers higher aggregate hedger welfare* *than the series of standard futures.*

Standard futures benefit hedgers whose needs align to those futures' maturity dates. However, in aggregate, that benefit is outweighed by the risk of futures rolls and early unwindings borne by hedgers with misaligned dates. Moreover, all hedgers benefit from the lower expected costs that arise because perpetual futures reduce the compensation that market makers and liquidity traders demand for bearing the hedgers' risk.

Combining Corollaries 2.3 and 2.4, we conclude that perpetual futures contracts can lower the magnitude of the futures basis, increase trading volume, and enhance the welfare of all three types of market participant. This result illustrates that well-designed perpetual futures not only discipline prices through funding but also better accommodate the heterogeneity in investor horizons, avoiding the roll frictions created by standard futures.

3. Empirical analysis

3.1. Cryptocurrency trade and orderbook data

We test our model's predictions using a combination of aggregate market data and high-frequency trading records. Appendix Figure B1 summarizes the product offerings of all exchanges that launched cryptocurrency futures prior to 2019. While nearly all surviving cryptocurrency futures exchanges now offer perpetual futures as their primary contract type, we focus on three: BitMEX, Deribit, and Kraken.

These three exchanges are ideal for our analysis for several reasons. First, they were the first to offer perpetual futures contracts. By the time other exchanges introduced these contracts, perpetual futures were already arguably the market standard. Second, all three provide complete historical trading data, which we use from each venue's inception through to a cutoff date of December 31, 2023. Third, and unlike the next set of adopters, these three exchanges have not been credibly accused of manipulating trading data.[5]

BitMEX was the first cryptocurrency exchange to offer cryptocurrency perpetual futures. It provides both trade and top-of-the-book data for essentially its entire existence starting November 11, 2014.

Deribit specializes in cryptocurrency options trading and was the second exchange to adopt perpetual futures. It provides trade, liquidation, and mark-to-market pricing data from its inception in February 2017 onward.

Kraken was the third exchange to adopt perpetual futures and the first regulated venue to do so. It was founded as CryptoFacilities, before being aquired by U.S.-regulated Kraken and rebranded as Kraken Futures. For simplicity, we refer to it as Kraken throughout. Kraken provides trade, liquidation, and mark-to-market pricing data from January 1, 2018 onward.

Spot prices are taken from Coin Metrics and calculated as the median of the last traded prices from Coinbase, Bitfinex, and Bitstamp. These are three of the longest-standing and most liquid exchanges. We use a median to mitigate the effect of outages or anomalies; however, these exchanges are well-aligned in price, unlike venues in jurisdictions with capital controls (Choi et al., 2022).

Aggregate futures volume and product listing data are also drawn from Coin Metrics and cover 20 futures exchanges, accounting for the majority of global cryptocurrency futures activity.

Futures contract structures varied across the sample. In the initial sample period, futures contracts were typically denominated in cryptocurrency because the early exchanges

[5]Our assessment of exchange credibility follows prior research identifying wash trading and fake volume, including Chen et al. (2022), Cong et al. (2022), Bitwise's 2019 SEC report, and the Blockchain Transparency Institute. These studies suggest that other early perpetual futures exchanges, such as HTX and OKX, may have provided unreliable data, making them unsuitable for our main analysis. Manipulation is a well-documented issue in cryptocurrency markets, partly because platforms generally act as exchange, custodian, and broker.

lacked access to the banking system. **Inverse** futures are crypto-denominated contracts referencing the price of USD in that crypto, as opposed to USD-denominated contracts referencing the price of cryptocurrency in USD.[6] Assets other than BTC are often **BTC-quoted**, for example, a future on ETH which was quoted and settled in BTC. Somewhat less commonly, **quanto** futures are based on the ratio of two currencies and paid out in a third currency. For example, an XRP/USD future settled in BTC, whose value (in BTC) is tied to the price of XRP in USD. Later in the sample period, most futures trading occurs using standard futures that pair a cryptocurrency with either U.S. dollars or (more commonly) a "stable coin" cryptocurrency that is pegged to the U.S. dollar. While these structures differ in accounting, they have no bearing on our results. Throughout, we analyze returns in USD and control for these structural differences.

3.2. The rise of perpetual futures

A core prediction of our model is that market participants prefer perpetual futures to standard futures. Figure 1 supports this: across the exchanges that introduced them, perpetuals rapidly became the dominant contract type.

As the figure shows, BitMEX was the first exchange to apply the perpetual futures contract to cryptocurrency and experienced a rapid shift in trading activity toward these contracts.[7] After their success on BitMEX, in 2018 these contracts spread to Deribit and Kraken, who quickly saw the new contracts eclipse their established standard futures offerings. HTX, OKX, and bitFlyer later underwent similar transitions. In all cases, perpetual futures saw

[6] Writing BTC-denominated contracts on USD was pioneered by ICBIT's founder Alexey Bragin to allow futures trading without reliance on fiat currency exchanges, as detailed in Bragin (2015). To illustrate, suppose a trader goes long \$100 (4 BTC) of a futures contract at a price of \$25 per BTC and at settlement the BTC price is \$50. If she was long a conventional BTC/USD future as offered by the CME, she would have a profit of $\$100 = 4 \times (\$50 - \$25)$ because she had a USD-denominated long claim on 4 BTC each of which increased by \$25. If she was long a so-called inverse BTC future, she would have a profit of $2\text{BTC} = \$100 \times (1/25 - 1/50)$ because she had a BTC-denominated short claim on \$100 and the dollar fell from 1/25th of a BTC to 1/50th.

[7] The temporary rebound in standard futures volume in 2017 is driven by a surge in popularity of BTC-JPY standard futures, which did not have an associated perpetual future.

immediate popularity.

Exchanges like Bybit and Binance, which entered the cryptocurrency market later, almost always started with a perpetual futures offering. In many cases, exchanges started with only a perpetual futures offering. For example, the offshore subsidiaries of U.S.-based Coinbase and Gemini initially offered only perpetual futures and did not offer standard futures at all. A similar focus on perpetual futures is seen on emerging decentralized blockchain-based derivatives platforms such as Jupiter, dYdX, and Hyperliquid, which are transparent by nature and have perpetual futures as their only derivative offering.

A notable exception to this global trend is found in the three U.S.-based exchanges: the CME, Coinbase Financial Markets, and ErisX (now owned by CBOE). This divergence is not coincidental. While perpetual futures are not explicitly prohibited in the United States, their regulatory status remains ambiguous and potentially limited by strict rules on swaps. To date, no perpetual futures contracts have been approved by the Commodity Futures Trading Commission (CFTC), the primary U.S. futures regulator. This may change in the future, as the CFTC is exploring the possibility of liberalizing the regulatory framework surrounding perpetual futures.[8]

Figure 2 focuses on the months surrounding the launch of perpetual futures on BitMEX, Deribit, and Kraken. We consider only inverse BTC futures here, because that was the dominant quoting structure for both standard futures and the newly introduced perpetual futures.

For each exchange, we divide BTC futures into three categories based on expiration: (i) quarterly standard futures with calendar-quarter expirations; (ii) short-maturity standard futures, with maturities other than quarter ends; and (iii) perpetual futures. The set of short-maturity contracts varies across exchanges and includes monthly, weekly, 48-hour, and even 24-hour contracts.

[8]See, for example, the CFTC's request for comment on the regulation of perpetual futures.

The top set of subplots shows the normalized futures basis: the difference between the futures price and the spot price, all divided by the spot price. Consistent with our model, the basis of quarterly futures is large, volatile, and converges to zero at maturity. In Bitmex's case, the quarterly futures traded up to 40% above the spot price, while for Deribit and Kraken they traded up to 8% above. Looking back at Equation (6), the larger effects we see on BitMEX could be explained by the relatively scarce arbitrage capital (low η^M) in the early cryptocurrency markets.

Also consistent with the model, the short-maturity standard futures track the spot price more closely thant the quarterly futures. The perpetual futures used by cryptocurrency exchanges track the spot price even more closely due to their high funding rates.

The bottom row of subplots presents the number of trades by contract type.[9] The rapid rise of perpetual futures is clearly visible on all three exchanges. Moreover, the introduction of perpetual futures is associated with an increase in total trading activity, rather than merely a reallocation of trades away from existing contracts. This is consistent with the model prediction that reduced basis risk increases participation.

Also evident is that perpetual futures effectively replaced short-maturity standard futures on all three venues. Because almost all trading occurs using perpetual futures or quarterly futures, we drop short maturity futures from the subsequent analysis.

These patterns are unlikely to be driven by fees or margin requirements. Perpetual futures did not offer fee advantages: they had fee parity with standard futures on BitMEX and Kraken, and were actually subject to higher fees on Deribit (perhaps explaining why perpetual futures are associated with increasing platform volume for Kraken and BitMEX but not Deribit).

Margin requirements are another salient contract feature that could influence trader de-

[9]Traded volume displays qualitatively similar patterns. We focus on trade counts rather than trade volume due to the extreme volatility of BTC prices during the period.

mand. Kraken and Deribit initially launched perpetual futures with the same margin requirements as their standard futures, yet still experienced rapid adoption of perpetual contracts. BitMEX offered lower margin requirements for short-maturity futures and perpetual futures than for quarterly futures; however, the higher margin requirements for longer-maturity futures appear to have been driven by their erratic pricing behavior rather than an attempt to steer demand.

Appendix Figure B2 confirms that elevated basis volatility in quarterly futures persists over time, supporting the model's prediction that perpetual futures offer more stable pricing environments. High and unstable futures bases were present even in recent years, highlighting that the divergence of spot and futures prices is not just a historical artifact.

3.3. Liquidity of quarterly and perpetual futures

A second prediction of the model is that the prices of high-funding-rate perpetual futures are less sensitive to changes in demand, either from liquidity traders or hedgers. We test this prediction using standard measures of market liquidity.

Quoted spread. As we have top-of-the-book data from BitMEX, we construct the quoted spread at time t as

$$QuotedSpread_t = \frac{P_t^{ASK} - P_t^{BID}}{P_t^{MID}}, \tag{15}$$

where P_t^{ASK} and P_t^{BID} denote the best ask and bid quotes at time t, and P_t^{MID} is the midpoint between them. For each day d, we compute the average of the quoted spread measured at the end of each minute.

Roll Spread. We also implement the Roll (1984) estimator of transaction costs, which infers the spread from the absolute value of the serial covariance of price changes. For day

d, we define this as

$$\textit{RollSpread}_d = 2\sqrt{\left|\text{AUTOCOVAR}_{t\in N_d}\left[P_t/P_{t-1}\right]\right|}, \tag{16}$$

where $\left|\text{AUTOCOVAR}_{t\in N_d}[\cdot]\right|$ is computed over all trades on day d. We exclude symbol-day observations with fewer than 10 trades to ensure reliability.

Effective Spread. The effective spread of a trade is twice the difference between its price and a pre-trade reference price, signed based on which side of the trade was the aggressor:

$$\textit{EffectiveSpread}_t = 2 \times \begin{cases} P_t/P_{t-}^{Ref} - 1 & \text{if trade } j \text{ is a buyer-initiated} \\ 1 - P_t/P_{t-}^{Ref} & \text{if trade } j \text{ is seller-initiated,} \end{cases} \tag{17}$$

where P_{t-}^{Ref} is the pre-trade reference price. For BitMEX, we use the midpoint from one second before the trade as the reference price. For Kraken and Deribit, we use the exchanges' reported mark prices, which are smoothed estimates of P_t^{Mid} designed to be robust to manipulation.[10] For Deribit, we lack trade initiator flags, so we classify a trade as buyer-initiated if it occurs above the mark price, and seller-initiated if it occurs below. To minimize bias from imbalanced trade flow, we compute the daily effective spread as the midpoint between the average of spread of buyer-initated trades and the average spread of seller-initiated trades. If the average effective spread on a day is negative, we set it to zero. Finally, we discard this measure for symbol-days with fewer than 10 buys or sells.

[10]Both Kraken and Deribit compute the mark price as $S_t + EWM_{30}[P_s^{ImpactMid} - S_s]$, where $EWM_{30}[\cdot]$ is a 30-second exponentially weighted moving average, $P_s^{ImpactMid}$ is the midpoint between the average price a market sell of a certain quantity would realize and the average price a market buy of that same quantity would realize, and S_s is the exchange-defined spot reference price. Details are provided online by Kraken and Deribit.

Realized spread. The realized spread of a trade is calculated similarly to the effective spread, except that instead of using the mark price prior to the trade, we use the price of the next trade to occur that is at least five minutes from the trade in question. As before, we discard this measure for symbol-days with fewer than 10 buys or sells.

Table 1 reports summary statistics for our market quality metrics. We report these metrics for quarterly futures and perpetual futures on BitMEX, Deribit, and Kraken.

The first set of columns reports our liquidity measures for BTC futures on each venue. Across all three exchanges and all liquidity measures, perpetual futures are much more liquid than quarterly futures. In particular, perpetual futures exhibit 5–8 times as many daily trades and spreads that are 49–83% narrower.

The second set of columns repeats that analysis but includes all contracts. Here, the pattern is mixed: perpetual futures continue to show higher trading volumes, but in some cases exhibit wider average spreads. This reflects the endogenous nature of contract structure: the least liquid coins are listed only as perpetual futures and not as quarterly futures. For example, Deribit offers quarterly futures only on BTC, ETH, and SOL–its most liquid assets–while it lists perpetual futures on 15 different coins. Because less liquid coins have higher spreads, the fact that perpetual futures are listed on less liquid coins means the cross coin average liquidity is lower.

The third set of columns attempts to control for that by considering only coin-day-exchanges where both perpetual futures and quarterly futures are listed and restricting consideration to the most liquid contract on that coin-day-exchange. With this restricted sample, we see a similar pattern to the first set of columns: for a given coin and venue, the perpetual futures are more popular and more liquid, across venues and across measures.

Table 2 explores the determinants of liquidity using a panel regression design. We use the more-liquid sample identified in the third set of columns of Table 1, considering only the most liquid . Letting $y_{c,v,s,t}$ be a liquidity measure taken on day t for coin c using futures contract structure s on venue v:

$$y_{c,v,s,t} \;=\; \beta'\mathbf{X}_{s,t} \;+\; \upsilon_{c,t} \;+\; \gamma_v \;+\; \varepsilon_{c,v,s,t}. \tag{18}$$

The covariate vector X_s contains contract-structure characteristics. In particular, we consider an indicator for the symbol being a perpetual future (as opposed to a standard future) and the log of one plus the number of days until the future's expiration, with perpetual futures set to 0. Coin–day fixed effects $\upsilon_{c,t}$ absorb market conditions, while venue fixed effects γ_v capture time-invariant platform heterogeneity. Standard errors are clustered at the coin-month level.

Specification (1) shows that perpetual futures trade at effective spreads that are 4.6 bp narrower than comparable quarterly futures: comparable to the magnitudes from our matching exercise and large given the 10-14 bp average spreads on these contracts. Specifications (2)–(4) rerun these effective spread tests venue-by-venue and yield qualitatively similar results. Specification (5) introduces $Tenor$, our measure of time to maturity and confirms the model prediction that liquidity is worse at longer maturities. Specifications (6) and (7) repeat the tests in Specification (1) using realized spreads and Roll spreads, respectively; the liquidity advantage of perpetual futures remains robust across these alternative metrics.

3.4. Liquidity around the introduction of perpetual futures

Our model suggests that the perpetual futures structure has intrinsically better liquidity in markets with excess speculation. Although Tables 1 and 2 confirm the superior liquidity of perpetual futures, these estimates are hard to interpret in light of perpetual futures' market dominance. The better liquidity of perpetual futures might simply reflect their current dominance, rather than anything intrinsic to their contractual design.

To isolate the effect of contract design, Figure 3 looks back to the launch of the first perpetuals. In these early days, perpetual futures were the entrants rather than the dominant incumbents. For each of our three early adopters, it shows the Roll spreads and effective spreads on perpetual and for front-month quarterly BTC futures.

Perpetual futures exhibit markedly tighter spreads on every exchange and for both liquidity measures, even immediately after their introduction. The economic magnitudes differ. BitMEX shows the largest gains, while Deribit and Kraken record more modest improvements and more noise. These differences may be driven by the level of basis risk during these periods. As Figure 2 shows, BitMEX introduced perpetual futures during a period when the quarterly futures basis was high and volatile–conditions under which the model predicts that the quarterly futures form would inhibit trading. Deribit and Kraken introduced perpetuals when basis volatility was lower, muting the theoretical benefit of perpetual futures.

Table 3 tests how the introduction of perpetual futures impacted market quality using a panel event study design:

$$\begin{aligned} y_{c,v,s,t} &= \sum_i \gamma_i \, \mathbb{I}\big[t \in \text{Quarter } i_{c,v}\big] \, \mathbb{I}\big[\text{Perpetual}_s\big] \\ &\quad + \sum_i \alpha_i \, \mathbb{I}\big[t \in \text{Quarter } i_{c,v}\big] \; + \; \beta' X_{s,t} \; + \; \upsilon_{c,t} \; + \; \gamma_v \; + \; \varepsilon_{c,v,s,t}. \end{aligned} \tag{19}$$

Here, $y_{c,v,s,t}$ is again a liquidity measure, and the variables of interest are indicator variables for quarter and contract type. $\mathbb{I}\big[t \in \text{Quarter } i_{c,v}\big]$ equals one if day t falls in quarter i, where quarter 0 is the 91 days immediately after the introduction of perpetual futures on coin c on venue v, quarter 1 is 91–181 days after, quarter ≥ 2 is more than 182 days after, and quarter ≤ -2 is greater than 91 days prior. We isolate the time-variation in liquidity differentials by interacting these quarter indicators with an indicator variable equal to one if the contract is a perpetual future. Other variables follow Equation 18, including the perpetual indicator, exchange fixed effects, contract controls, and coin-month fixed effects. As before, all standard errors are presented after clustering standard errors at the coin-month level.

Table 3 shows that perpetual futures have significantly lower Roll spreads and significantly more trades than quarterly futures, even in the days following their introduction. For the subsequent quarter and the quarters after that, perpetual futures show large improvements on effective spread, Roll spread, realized spread, and the number of trades.

These interaction terms are not driven by a decline in the liquidity of quarterly futures. quarterly futures liquidity is mostly unchanged, with the Roll spread shrinking and the number of trades growing later in the sample.

3.5. Liquidations

Another notable feature of cryptocurrency markets is the prevalence of large, forced liquidiations. Although data on liquidations in traditional futures markets is lacking, anecdotal evidence suggests liquidations are much more prevalent in cryptocurrency. One driver of this may be that cryptocurrency futures exchanges are non-recourse for both perpetual and quarterly futures, and traders commonly use position-level margin and strategically triggered liquidations as a way to limit downside exposure.[11] That contrasts with traditional futures markets, where futures are recourse and losing traders can not just lose their collateral but be driven into debt.

Figure 4 reports data on the liquidations on Deribit and Kraken, the two venues for which we have complete liquidations information. The top set of plots in each panel report the portion of each venues' trading volume that is accounted for by liquidation trades. On Kraken, we see much higher levels of liquidation and what appears to be a transition of liquidation activity from quarterly futures to perpetual futures. Before the introduction of perpetual futures, liquidations accounted for 2.8–3.6% of quarterly futures volume. After the introduction of perpetual futures, the prevalence of liquidations on quarterly futures

[11] For example, on many venues a trader could enter long position on a BTC future and fund that position with only part of their account value. If BTC's price fell, that trader would be liquidated, but only the part of their account that they committed to the futures contract would be at risk.

plumeted to 0.3-0.7%, whereas the liquidations on perpetual futures stayed high at 1.9-3.0%.

The apparent transition of liquidations from quarterly futures to perpetual futures on Kraken occurs despite both contracts having the same margin requirements. This is consistent with the model prediction that shorter-horizon, financially-constrained traders would prefer perpetuals and migrate toward them once offered.

On Deribit, we see no such pattern and both contracts have similar reported liquidations accounting for 0.4–0.8% of trading volume. The SOL quarterly futures see volatile patterns, but that is associated with the standard future being discontinued due to lack of activity. A potential explanation for this is that Deribit is primarily an options exchange–with the risk seeking traders on Deribit choose options, rather that poorly capitalized futures positions.

The bottom plots in each of Figure 4's panels show the effective spread at which those liquidations occur. Even though these are forced sales, the price impact is limited, with spreads of only a few basis points. Consistent with the rest of the paper, the spreads are higher for standard futures than for perpetual futures. Liquidations (a demand shock) has a smaller effect on the perpetual futures price than on the standard futures price, in the model this is driven by perpetual future's lower liquidity trader price coefficient.

4. Futures contract structure and market stress

4.1. Futures returns during market stress

The cryptocurrency market, with its frequent manias and panics, offers a rich laboratory to study how futures perform in times of market stress. Figure 5 assesses the risk created by basis movements by reporting the returns to spot, futures, and arbitrage positions. We focus on BTC contracts on BitMEX, because it has the longest continuous history of perpetual futures and because it reports quoted spreads.

Our analysis focuses on hourly periods in which the spot BTC price changed by more than 5%. In these moves, the spot BTC prices move by an average of 5.8% within two hours. Such moves indicate severe market stress as throughout our sample, BTC comprised the bulk of total crypto-market capitalization. For context, the worst single-day declines of the 2008 financial crisis were roughly 7–8%. Every futures contract has a long and a short side, and these charts are based on the losses borne by the losing side–i.e., returns to short (long) positions when the price rises (falls).

The left plots reports the returns to four strategies that each maintain a constant notional exposure to BTC: (i) a spot position, (ii) BitMEX's original XBTUSD perpetual future, (iii) the front-quarter (nearest-maturity) standard future, and (iv) the second-nearby (second-nearest-maturity) standard future. All four strategies incur sharp drawdowns, but all the three futures strategies have greater peak losses than the spot strategy. The front quarterly quarterly futures lose 8% (0.5 percentage points) more than the spot strategy and the second-nearby standard future overshoots by 10% (0.6 percentage points). Perpetual futures prices stay closer to the spot prices and so the perpetual futures strategy only overshoots by 3% (0.1 percentage points). Strong arbitrage pressure keeps the perpetual futures price close to the spot price, even during these extreme minutes.

The two quarterly futures contracts suffer similar losses, despite the front quarterly futures contract being far more liquid than the second-nearby standard future. This is inconsistent with liquidity-stories, but entirely consistent with the model. Under the model, the initial overshoot of the futures prices is driven by changes in the liquidity trader demand, whether due to sentiment swings or simply liquidations (e.g., of the optimists). If liquidity trader demand shifts are not persistent, this shock will have a similar effect on both the front and second-nearby quarterly futures.

Remarkably, within 16 hours of the shock the perpetual-futures position has lost less than the spot strategy. This cushioning comes from funding: once the futures–spot gap opened, the losing side of the perpetual futures contract began receiving transfers from the winning side, partially offsetting the losing sides losses. quarterly futures amplify spot shocks; perpetual futures, through funding, dampen them.

The right plot on Figure 5 shows the losses of BTC basis arbitrage trades, where futures positions are hedged using offsetting spot positions. This type of spread trade is popular both in cryptocurrency and in traditional finance, for example, the treasury cash-futures basis trade or the covered interest parity trade. In the long term, these arbitrage strategies converge to a nearly risk-free profit. In practice, they can realize large losses when markets are under stress.

We again plot the returns to the losing side of these basis trades. To reflect the prevalence of leverage in these trades (e.g., Siriwardane et al. (2022)), we have scaled these arbitrage trades by 10 to reflect hypothetical 10-times leverage and assumed 10% annual interest for spot leg. In our data, this basis trade loses 4.3-4.6% in the two hours surrounding the $> 5\%$ spot drawdowns if it is conducted using the quarterly futures contracts. Using the perpetual future reduces that peak drawndown to 1.5%, a roughly two-thirds reduction that is consistent with the reduction in divergence. As before, the losses are not only less severe, the are also shorter lived.

By reducing basis risk, perpetual futures reduce the losses arbitrageurs bear in crises. This might ease the pressure on the financial system, as large arbitrage drawdowns can amplify market turmoil (Edwards, 1999; Glicoes et al., 2024). Because arbitrage capital and market liquidity are tightly linked (Mitchell and Pulvino, 2012; Rösch, 2021), broader use of perpetuals could also bolster liquidity in stressed conditions.

4.2. Futures liquidity during market stress

The limits to arbitrage are particularly acute during crises. Figure 6 illustrates the relationship between liquidity and market stress by showing the median quoted (bid-ask) spread of BitMEX BTC-USD futures around the same set of 5% hourly spot price movements. We plot large price increases (left-plot) and decreases (right-plot) separately.

For perpetual futures, the quoted spread climbs only modestly: in the two hours around a $> 5\%$ spot rally it averages 0.66 bp, just 34% above the 0.49bp baseline observed 1–12 hours earlier. A comparable 25% widening ($0.49 - -0.66$bp) occurs during spot declines. The front future has worse initial liquidity and suffers far larger liquidity deterioration, its spread jumping 161% on the rallies ($1.66 - -4.36$bp) and 249% on the sell-offs ($1.44 - -5.03$bp).

The second-nearby quarterly future begins with spreads roughly twice those of the front standard future (3.19—3.39bp). After the shock, its spread widens less than the front standard future in percentage terms (115—161% versus 161—249%) but more in absolute terms (1.42—2.10bp versus 0.61—1.12bp). This suggests that the better crisis performance of the perpetual future may not be solely driven by its better overall liquidity: the perpetual future has far better performance than the somewhat-less-liquid front quarterly future, while the front quarterly future performs has only slightly better relative performance than the somewhat-less-liquid second quarterly future.

Table 4 formally tests these relationships by comparing the liquidity of perpetual and quarterly futures. We test whether the liquidity advantage of perpetual futures holds during periods of market stress, using minute-level spreads and a panel regression with stress-by-contract interactions. We focus on the BitMEX exchange because it reports quoted spreads and BTC futures because of their superior liquidity. Here, the liquidity measure in minute t in month m using futures contract structure s is

$$\begin{aligned} y_{s,m,t} =& \alpha_P \mathbb{I}[\text{Perpetual}_s] + \gamma_P I[\text{Perpetual}_s] \times \Upsilon_t \\ &+ \alpha_B \mathbb{I}[\text{Back-month}_{s,t}] + \gamma_B I[\text{Back-month}_s] \times \Upsilon_t \\ &+ \gamma_t + \nu_s + \epsilon_{s,m,t}, \end{aligned} \tag{20}$$

where Υ_t is a minute-level measure of market stress defined as either the absolute change in the log spot BTC price for minute (Current volatility) or the average of such minute-level absolute changes over the prior 15 minutes (Past volatility). For clarity, we scale Υ_t to be in percentage terms. Our variable of interest is this measure of market stress interacted with dummies indicating perpetual futures or back-month (non-nearest expiring) quarterly futures. Minute-level fixed effects γ_t absorb variation in Υ_t and symbol-level fixed effects ν_s absorb variation at the contract level, meaning the interaction of market stress and the perpetual (or back-month futures) futures form is identified by the extent to which the futures spread of the contract structure differentially respond to market stress. Other variable definitions mirror Equation 19, with standard errors clustered at the month-level.

In Specifications (1) and (2), we see that perpetual futures spreads shrink relative to front month quarterly futures spreads immediately following periods of high spot BTC price volatility. This holds both when we look at log of quoted spreads plus one basis point (Specification 1) or quoted spreads (Specification 2). These effects are economically meaningful–a one standard deviation increase in past volatility (0.057) translates into a 10 bp (0.23 log) increase in the spread of the front quarterly future relative to the perpetual future. The effects in periods of peak market stress are even more pronounced, with a movement from the median (0.048) to the 99th percentile (0.28) of past volatility translating into 42 bp (0.94 log) increase in the front quarterly futures relative spread. Although the back-month quarterly future is much less liquid than the front-month future, we do not see statistically significant market-stress-liquidity interactions.

In Specifications (3) and (4), we repeat these tests using within-minute absolute price movements. This measure captures the spread at the end of the minutes in which there were large price movements. We see similar effects to the measure of past volatility: in the minutes with the largest market movements, the front-month quarterly futures spread increases significantly and the back-month shows only a limited change. Importantly, the minute-level price changes are measured based on spot BTC, somewhat reducing the concern that our measures are contaminated.

5. Conclusion

Perpetual futures maintain alignment with underlying assets through regular "funding rate" payments exchanged between long and short positions. While standard futures converge to spot prices only at maturity, perpetual futures achieve continuous convergence through their funding payment mechanism. Analysis of cryptocurrency exchange data demonstrates that perpetual futures consistently track spot prices more closely than their quarterly counterparts.

These contracts now dominate trading on every major non-U.S. venue and, even after accounting for their popularity, exhibit superior market quality. In particular, perpetual futures remain more liquid during periods of market instability.

Tighter alignment between spot and futures prices reduces risk. Following large spot price movements, quarterly futures exhibit losses 8–10% greater than the spot decline, while perpetual futures move nearly one-for-one with spot. This price divergence is particularly hazardous for arbitragers conducting basis trades because quarterly futures strategies impose crisis drawdowns that are three times as large as perpetual futures strategies.

We build a continuous-time model of financially constrained arbitrageurs to rationalize these facts. Funding rate payments increase arbitrage pressure and cause the futures price to closely track the spot price. Reducing the risk of price divergence reduces the compensation arbitrageurs demand, leading to more trade and higher welfare. Hedgers benefit from these reduced costs and can also benefit from better hedging. Hedger risk is minimized by standard futures whose expiration perfectly coincides with their hedging needs, however, those needs may not coincide with the maturity of the standard futures. Because of these misaligned hedgers, perpetual futures can improve aggregate hedger welfare by reducing the risk of rolling futures or closing them prior to maturity.

Our results suggest that perpetual futures emerged as a contractual innovation tailored to the volatile, arbitrage-capital-constrained environment of cryptocurrency markets. The success of this contractual form in this challenging setting highlights its potential as a financial innovation. This has important policy implications, particularly as the CFTC considers whether to permit perpetual futures in U.S. markets.

References

Acharya, V. V., Lochstoer, L. A., and Ramadorai, T. (2013). Limits to arbitrage and hedging: Evidence from commodity markets. *Journal of Financial Economics*, 109:441–465.

Ackerer, D., Hugonnier, J., and Jermann, U. (2024). Perpetual futures pricing. Technical report, National Bureau of Economic Research.

Akey, P., Robertson, A., and Simutin, M. (2022). Noisy factors. *Rotman School of Management Working Paper Forthcoming*.

Alexander, C., Chen, X., Deng, J., and Wang, T. (2024). Arbitrage opportunities and efficiency tests in crypto derivatives. *Journal of Financial Markets*, 71:100930.

Alexander, C., Deng, J., and Zou, B. (2021). Hedging with bitcoin futures: The effect of liquidation loss aversion and aggressive trading. *arXiv preprint arXiv:2101.01261*.

Allen, F. and Gale, D. (1994). *Financial innovation and risk sharing*. MIT press.

Allen, F. and Gorton, G. (1993). Churning bubbles. *The Review of Economic Studies*, 60(4):813–836.

Bragin, A. (2015). Inverse futures in bitcoin economy. *Available at SSRN 2713755*.

Brunnermeier, M. K. and Pedersen, L. H. (2009). Market liquidity and funding liquidity. *The review of financial studies*, 22(6):2201–2238.

Castelino, M. G. and Francis, J. C. (1982). Basis speculation in commodity futures: The maturity effect. *The Journal of Futures Markets (pre-1986)*, 2(2):195.

Catalini, C. and Gans, J. S. (2020). Some simple economics of the blockchain. *Communications of the ACM*, 63(7):80–90.

Chen, J., Lin, D., and Wu, J. (2022). Do cryptocurrency exchanges fake trading volumes? an empirical analysis of wash trading based on data mining. *Physica A: Statistical Mechanics and its Applications*, 586:126405.

Choi, K. J., Lehar, A., and Stauffer, R. (2022). Bitcoin microstructure and the kimchi premium. *Available at SSRN 3189051*.

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

Cong, L. W., Li, X., Tang, K., and Yang, Y. (2022). Crypto wash trading. Technical report, National Bureau of Economic Research.

De Blasis, R. and Webb, A. (2022). Arbitrage, contract design, and market structure in bitcoin futures markets. *Journal of Futures Markets*, 42(3):492–524.

De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of political Economy*, 98(4):703–738.

Deng, J., Pan, H., Zhang, S., and Zou, B. (2020). Minimum-variance hedging of bitcoin inverse futures. *Applied Economics*, 52(58):6320–6337.

Duffie, D. and Rahi, R. (1995). Financial market innovation and security design: An introduction. *Journal of Economic Theory*, 65(1):1–42.

Edwards, F. R. (1999). Hedge funds and the collapse of long-term capital management. *Journal of Economic Perspectives*, 13(2):189–210.

Fett, N. and Haynes, R. (2017). Liquidity in select futures markets. *White paper. Office of the Chief Economist, United States Commodity Futures Trading Commission*.

Figlewski, S. (1984). Hedging performance and basis risk in stock index futures. *The Journal of Finance*, 39:657–669.

Foucault, T., Sraer, D., and Thesmar, D. J. (2011). Individual investors and volatility. *The Journal of Finance*, 66(4):1369–1406.

Garleanu, N. and Pedersen, L. H. (2011). Margin-based asset pricing and deviations from the law of one price. *The Review of Financial Studies*, 24(6):1980–2022.

Gehr, A. K. (1988). Undated futures markets. *Journal of Futures Markets*, 8(1):89–97.

Glicoes, J., Iorio, B., Monin, P., and Petrasek, L. (2024). Quantifying treasury cash-futures basis trades. *FEDS Notes*, 3(2024):8–3.

Harrison, J. M. and Kreps, D. M. (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, 92(2):323–336.

Hazelkorn, T. M., Moskowitz, T. J., and Vasudevan, K. (2020). Beyond basis basics: Liquidity demand and deviations from the law of one price. Technical report, National Bureau of Economic Research.

He, S., Manela, A., Ross, O., and von Wachter, V. (2022). Fundamentals of perpetual futures. *arXiv preprint arXiv:2212.06888*.

Makarov, I. and Schoar, A. (2020). Trading and arbitrage in cryptocurrency markets. *Journal of Financial Economics*, 135(2):293–319.

Miller, S. M. (1986). Financial innovation, depository-institution deregulation, and the demand for money. *Journal of Macroeconomics*, 8(3):279–296.

Mitchell, M. and Pulvino, T. (2012). Arbitrage crashes and the speed of capital. *Journal of Financial Economics*, 104(3):469–490.

Pagano, M. and Roell, A. (1990). Trading systems in european stock exchanges: current performance and policy options. *Economic policy*, 5(10):63–116.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of finance*, 39(4):1127–1139.

Rösch, D. (2021). The impact of arbitrage on market liquidity. *Journal of Financial Economics*, 142(1):195–213.

Ruan, Q. and Streltsov, A. (2022). Perpetual futures contracts and cryptocurrency market microstructure. *Available at SSRN 4218907*.

Rutledge, D. J. (1976). A note on the variability of futures prices. *The Review of Economics and Statistics*, pages 118–120.

Schmeling, M., Schrimpf, A., and Todorov, K. (2023). Crypto carry. *Available at SSRN 4268371*.

Shen, J., Yan, H., and Zhang, J. (2014). Collateral-motivated financial innovation. *The Review of Financial Studies*, 27(10):2961–2997.

Shiller, R. J. (1993). Measuring asset values for cash settlement in derivative markets: hedonic repeated measures indices and perpetual futures. *The Journal of Finance*, 48(3):911–931.

Shleifer, A. and Vishny, R. W. (1997a). The limits of arbitrage. *Journal of Finance*.

Shleifer, A. and Vishny, R. W. (1997b). The limits of arbitrage. *The Journal of finance*, 52(1):35–55.

Siriwardane, E., Sunderam, A., and Wallen, J. L. (2022). Segmented arbitrage. Technical report, National Bureau of Economic Research.

Tufano, P. (2003). Financial innovation. *Handbook of the Economics of Finance*, 1:307–335.

Figure 1: Rise of perpetual futures across exchanges.
This figure shows the share of futures volume accounted for by perpetual futures on the twenty cryptocurrency exchanges covered by Coin Metrics. These exchanges represent the bulk of global cryptocurrency futures volume. For BitMEX, Kraken, and Deribit, we use our own scraped data; for all other exchanges, we rely on Coin Metrics.

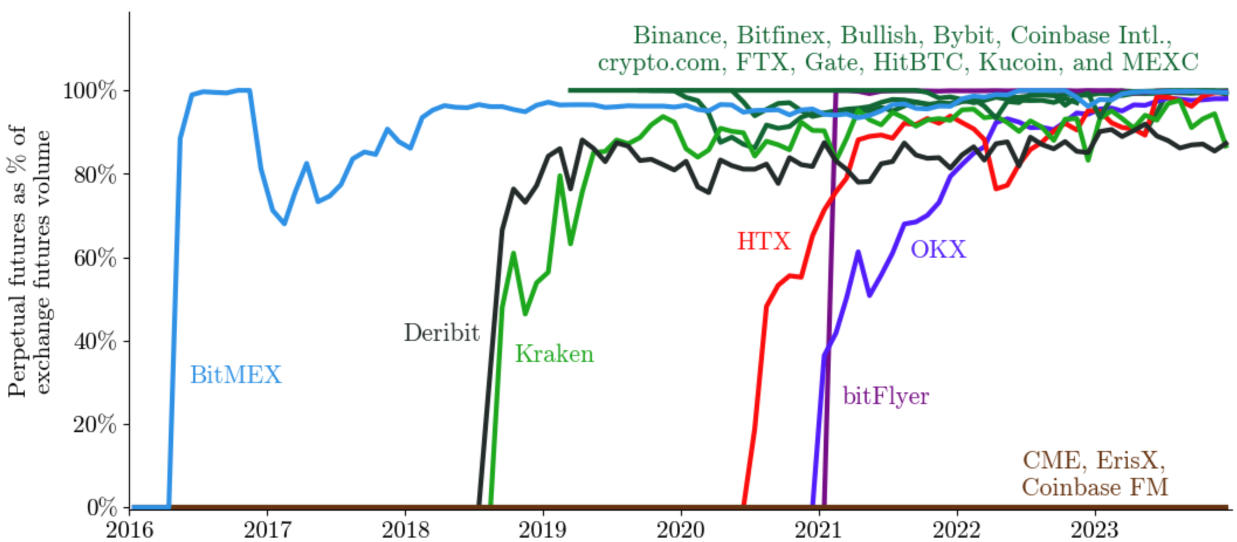

Figure 2: Futures markets around the introduction of perpetual futures.
This figure reports key features of Bitcoin (BTC) futures markets on BitMEX, Deribit, and Kraken around the time each introduced perpetual futures contracts. We group contracts into three categories: (i) quarterly futures, which have standard calendar-quarter expirations; (ii) short-maturity standard futures, including contracts with monthly, weekly, 48-hour, or daily expirations; and (iii) perpetual futures. The top set of panels show the weekly average basis for each contract type normalized by the spot price, measured as the ratio of the futures price to its spot underlying minus one. The bottom plots show the share of daily trades accounted for by each contract type.

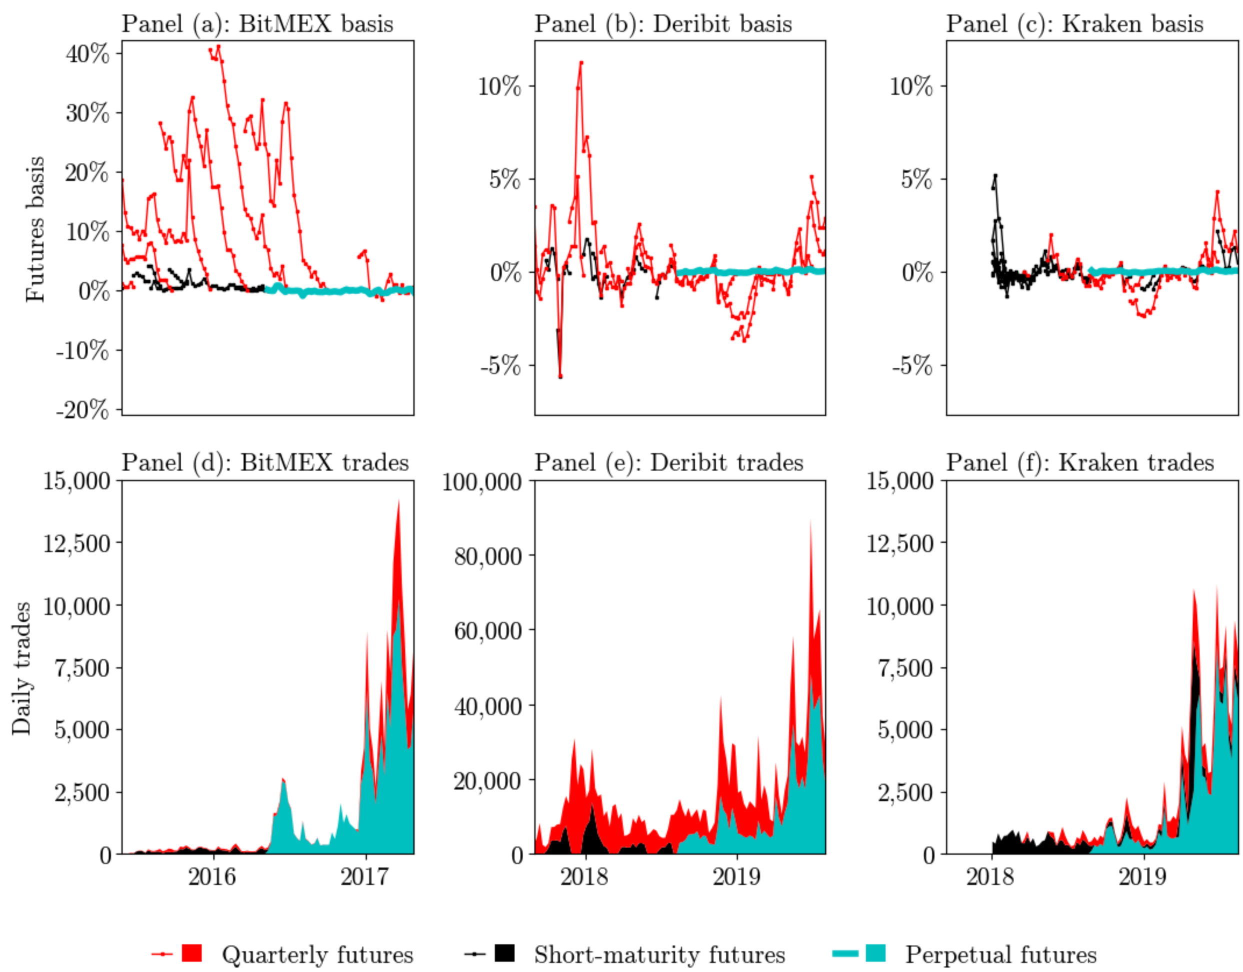

Figure 3: Futures market quality following perpetual future introduction.
This figure reports the liquidity of Bitcoin (BTC) futures markets on BitMEX, Deribit, and Kraken around the time each introduced perpetual futures contracts. We group contracts into three categories: (i) quarterly futures, which have standard calendar-quarter expirations; (ii) short-maturity standard futures, including contracts with monthly, weekly, 48-hour, or daily expirations; and (iii) perpetual futures. The top set of panels show the Roll spread in basis points, computed as twice the square root of the autocorrelation of relative price changes. The bottom set of panels show the effective spread in basis points, which is the average amount that liquidity takers overpay relative to a fair pre-trade price. For this chart only, we set the effective spread equal to 1/2 of a basis point for days where the effective spread is negative, to allow for the use of logorithmic scales.

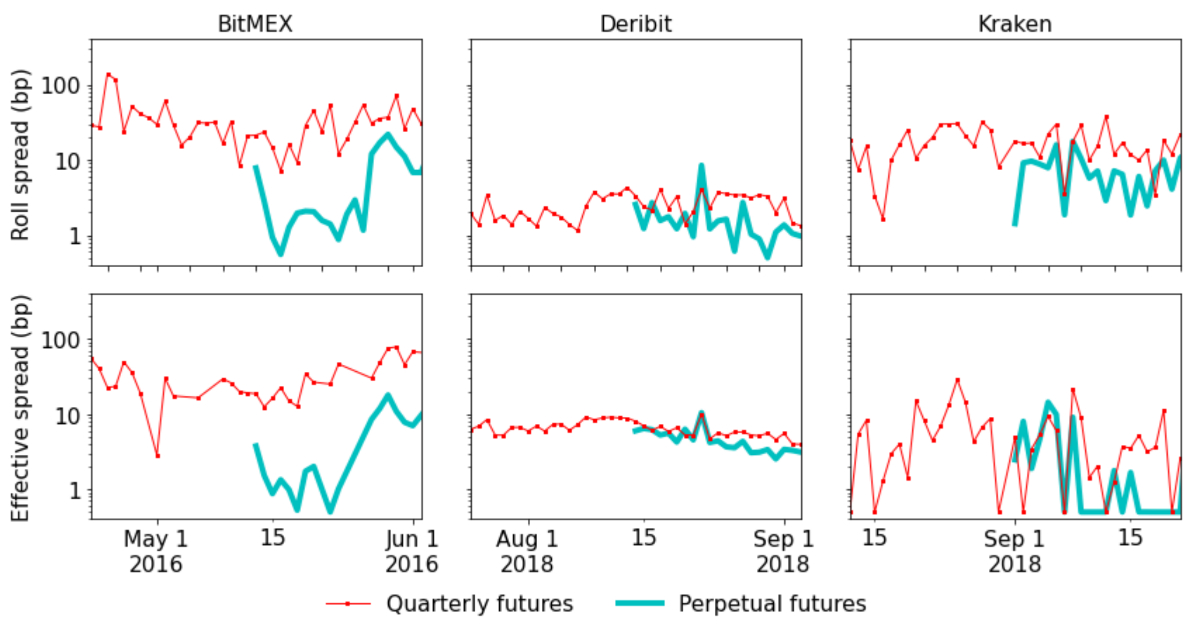

Figure 4: Liquidations by product type.
This figure shows the portion of weekly volume that is liquidations using Deribit (panel a) and Kraken (panel b) data. The top panel reports dollar-weighted monthly averages. Each line marks the percentage of trades in a given month that are liquidations for a different type of contract. The bottom panel shows the effective spread on liquidation trades. Monthly futures includes both monthly and quarterly futures. The dotted vertical line marks the introduction of perpetual futures.

Panel a: Deribit

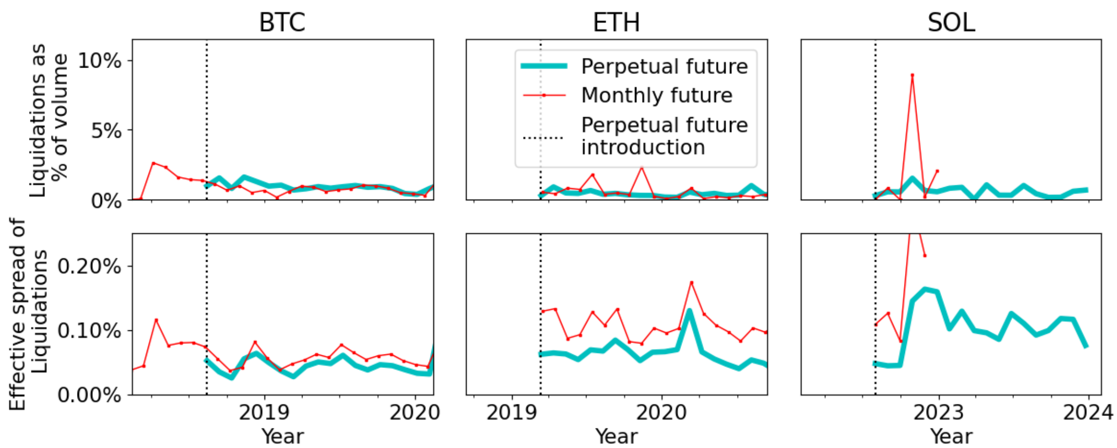

Panel b: Kraken

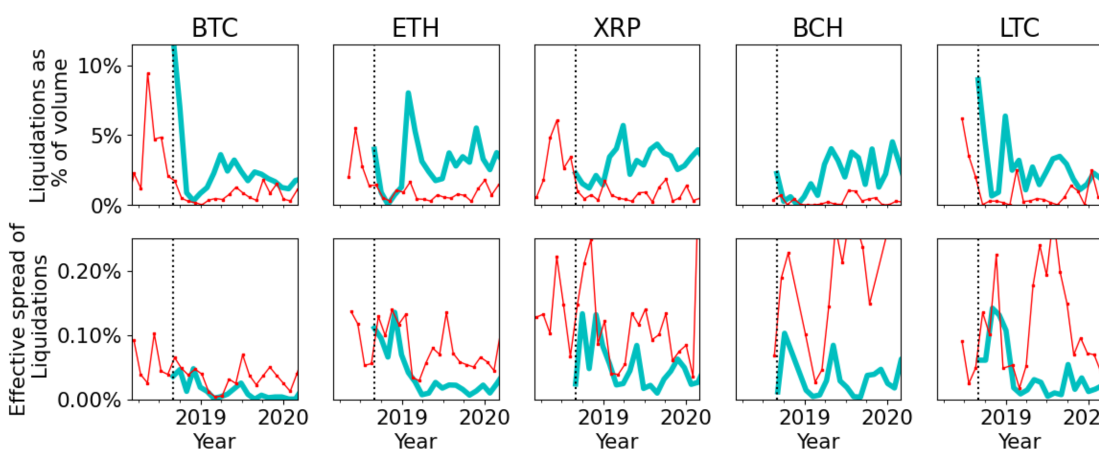

Figure 5: Crisis returns of futures and hedged arbitrage strategies.
This figure shows the average cumulative returns of various futures strategies around the time of > 5% absolute spot price movements. We consider the BTC-USD set of BitMEX contracts for events in the 2015-2023 period where both perpetual futures and two quarterly futures were available. If multiple > 5% absolute changes occur, we consider the first such change in a 36-hour period. The left plot shows the cumulative returns of a strategy that holds \$100 spot BTC (using the average of Coinbase, Bitstamp, and Bitfinex prices) and the cumulative returns to portfolios with the same BTC exposure constructed using BitMEX XBTUSD perpetual futures, front month (closest-maturing) futures, and second nearby (second closest-maturing) quarterly futures. The right panel shows the returns to arbitrage strategy that buys spot BTC and shorts the respective future with 10X leverage. Each line presents the cumulative cash-flows to the losing side of each trade, e.g., a long BTC trade around price decreases and a short BTC trade around price increases. We plot 15-minute averages of 1-minute cross-scenario medians.

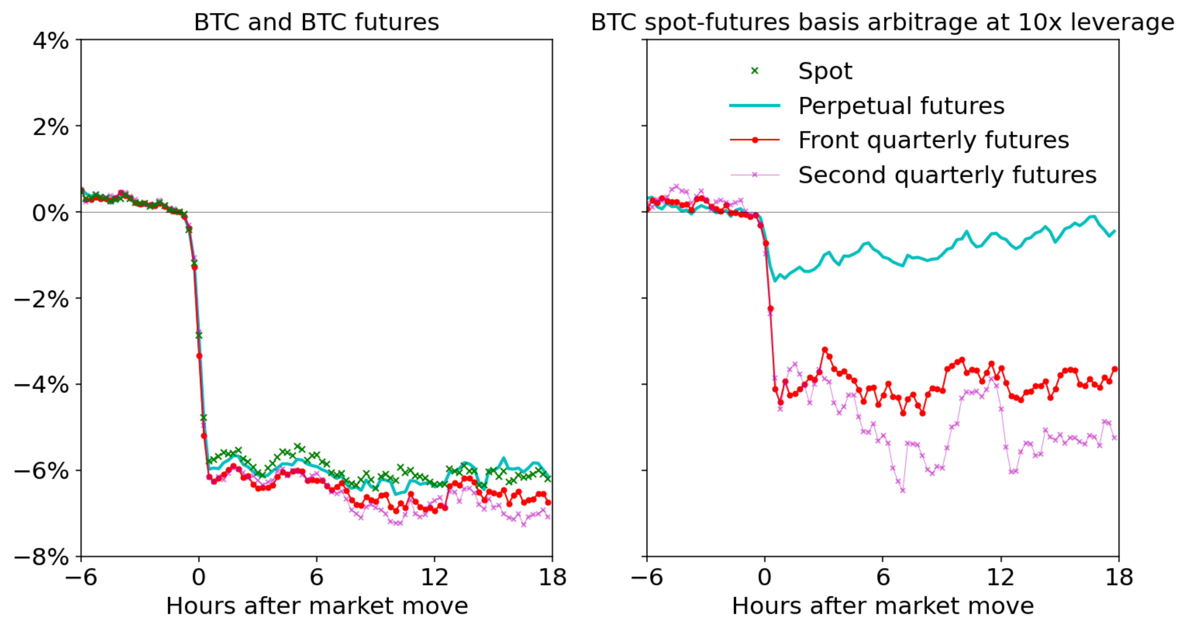

Figure 6: Quoted spreads of futures around the time of price jumps.
This figure shows the median quoted spread of various types of futures around the time of > 5% hourly spot price increases (left plot) and < −5% spot price decreases (right plot). We consider BTC-USD BitMEX contracts for market movements in the 2015-2023 period where both perpetual futures and two quarterly futures were available. If multiple > 5% changes occur, we consider the first such change in a 36-hour period. Each line presents the quoted (bid-ask) spread for a type of contract. We plot rolling 5-minute averages of 1-minute cross-event medians.

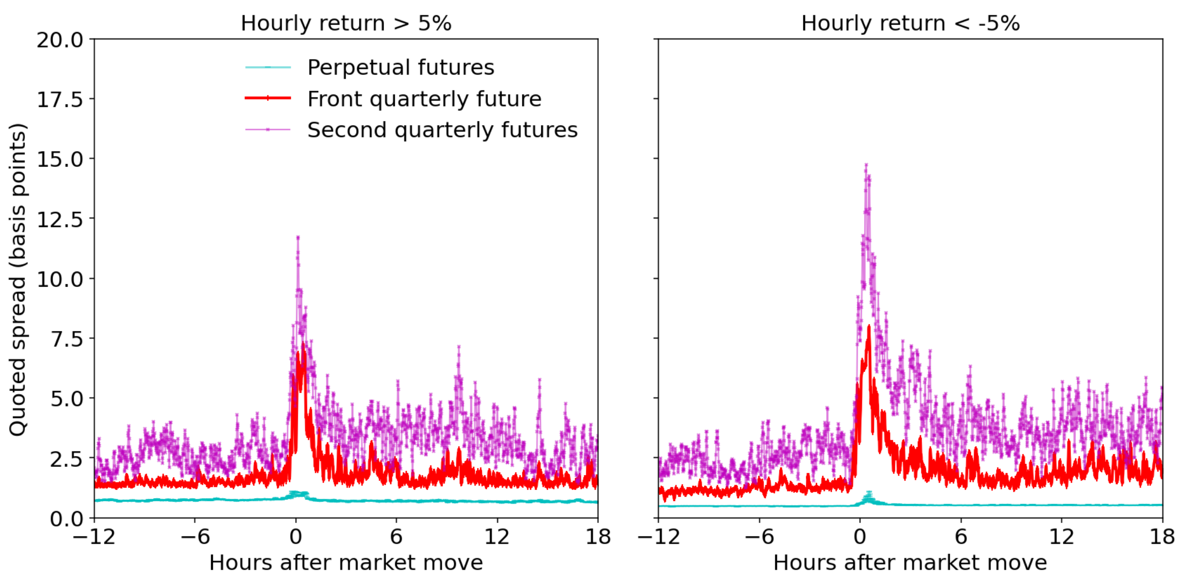

Table 1: Summary of market quality measures.
This table reports averages of liquidity metrics for standard quarterly futures and perpetual futures across BitMEX, Deribit, and Kraken. The first set of columns summarizes the entire sample and the second set of columns summarizes matched pairs of standard and perpetual futures. This matching pairs each quarterly future symbol-day with the perpetual future which has a number of trades closest to the quarterly future's number of trades, considering only matches within venue-coin-month and provided there exists a matching future with 70-120% the number of trades. Differences are tested using standard errors clustered at the coin-month level. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively.

	BTC			All coins			Matched coin-days		
	Qtly.	Perp.	Diff.	Qtly.	Perp.	Diff.	Qtly.	Perp.	Diff.
BitMEX									
Observations	3,069	2,788	-281	12,335	17,557	5,222	7,284	7,962	678
Distinct coins	1	1	0	12	31	19	8	8	0
Distinct symbols	54	3	-51	199	64	-135	132	17	-115
Daily trades	3,716	35,193	31,477***	1,905	8,181	6,276***	2,022	17,424	15,402***
Effective spread (bp)	14.4	2.1	-12.3***	19.9	12.2	-7.7***	9.7	3.6	-6.1***
Roll spread (bp)	28.3	3.8	-24.5***	31.8	26.4	-5.3*	22.3	9.7	-12.6***
Realized spread (bp)	19.2	2.8	-16.4***	24.5	21.7	-2.8	16.9	7.5	-9.5***
Kraken									
Observations	2,045	1,938	-107	9,063	29,053	19,990	8,657	9,289	632
Distinct coins	1	1	0	5	54	49	5	5	0
Distinct symbols	27	2	-25	134	60	-74	129	11	-118
Daily trades	602	8,454	7,852***	310	1,205	895***	318	3,178	2,860***
Effective spread (bp)	9.2	2.7	-6.5***	14.1	11.8	-2.3***	14.0	6.9	-7.1***
Roll spread (bp)	12.4	4.7	-7.7***	29.0	36.5	7.5***	28.5	11.1	-17.4***
Realized spread (bp)	7.6	3.2	-4.4***	16.7	19.2	2.6***	15.7	7.2	-8.5***
Deribit									
Observations	2,537	1,966	-571	4,420	9,985	5,565	3,849	3,862	13
Distinct coins	1	1	0	3	15	12	3	3	0
Distinct symbols	30	2	-28	53	18	-35	47	6	-41
Daily trades	2,378	16,165	13,787***	1,823	6,026	4,203***	1,707	15,065	13,358***
Effective spread (bp)	8.3	3.9	-4.5***	10.0	16.3	6.3***	10.3	4.5	-5.8***
Roll spread (bp)	7.6	2.1	-5.6***	9.6	15.7	6.1***	9.1	2.6	-6.5***
Realized spread (bp)	8.8	1.3	-7.5***	11.9	9.6	-2.4***	11.6	1.9	-9.6***

Table 2: Market quality of perpetual and delivery futures.
This table reports regressions of market quality metrics on various features for standard quarterly futures and perpetual futures on BitMEX, Deribit, and Kraken. A data point is a symbol-exchange-day, considering only six of the most liquid cryptocurrencies (BTC, ETH, SOL, XRP, BCH, LTC). The dependent variable is the effective spread in basis points for Specifications (1) to (5), the realized spread in basis points for (6), and the Roll spread in basis points for (7). *Perpetual* denotes perpetual futures. *Tenor* is the log of one plus the contract's maturity in days, where the maturity of perpetual futures is assumed to be 0. *Inverse*, *BTC-quoted*, and *Quanto* are contract features. Exchange refers to the futures trading venue. Standard errors are clustered at the coin-month level. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

	Effective spread					Realized spread	Roll spread
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Perpetual	−4.49***	−2.88***	−5.39***	−6.81***	0.55	−5.47***	−5.02***
	(0.30)	(0.81)	(0.68)	(0.89)	(0.91)	(0.32)	(0.34)
Tenor					1.15***		
					(0.20)		
Inverse	4.19***	−8.00	0.74	−1.88	4.03***	0.40	2.40***
	(0.65)	(6.48)	(1.14)	(1.78)	(0.65)	(0.56)	(0.72)
BTC-quoted	4.50***	−1.40		3.10	4.90***	−1.49***	0.94
	(0.83)	(2.09)		(2.84)	(0.83)	(0.50)	(0.65)
Quanto	3.44***	−1.59			3.47***	−0.95	1.23
	(0.94)	(2.22)			(0.94)	(0.72)	(0.99)
Kraken	6.47***			16.43***	6.69***	0.23	4.85***
	(0.66)			(1.42)	(0.66)	(0.52)	(0.55)
Deribit	8.32***		14.77***		8.16***	−0.45	0.25
	(0.83)		(0.98)		(0.82)	(0.45)	(0.41)
Coin-day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Exchange	All	BitMEX	Deribit	Kraken	All	All	All
Observations	54,926	21,454	11,004	22,468	54,926	54,926	54,926
R^2	18.99%	13.91%	51.88%	26.00%	19.36%	11.98%	9.98%

Table 3: Perpetual futures introductions and market quality.
This table reports a panel event study of market quality metrics for standard quarterly futures and perpetual futures on BitMEX, Deribit, and Kraken. A data point is a symbol-exchange-day, considering only six of the most liquid cryptocurrencies (BTC, ETH, SOL, XRP, BCH, LTC). The dependent variables are measures of spread (columns 1-3) and market activity (column 4). The variables of interest are indicators for the months surrounding the introduction of perpetual futures for that coin on that venue, with Month t being the month of introduction. The month prior to the introduction ($t-1$) is the omitted variable. The indicators for the months following the introduction of perpetual futures are interacted with Perpetual, which is a dummy variable equal to one for perpetual futures contracts. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively, for standard errors based on coin-month clustering.

	Effective spread (1)	Roll Spread (2)	Realized spread (3)	# Trades (4)
Perpetual × Quarter 0	−2.98	−2.21*	−4.42	0.59***
	(2.54)	(1.23)	(2.89)	(0.16)
Perpetual × Quarter 1	−3.36***	−2.88**	−3.51**	0.62***
	(0.67)	(1.39)	(1.41)	(0.13)
Perpetual × Quarter ≥ 2	−4.61***	−5.16***	−5.44***	2.00***
	(0.29)	(0.35)	(0.33)	(0.06)
Quarter ≤ −2	−3.40	−2.15	−1.92	0.59
	(5.38)	(2.79)	(2.89)	(0.44)
Quarter 0	0.43	2.17	0.21	−0.22
	(3.54)	(2.33)	(3.80)	(0.29)
Quarter 1	−2.40	−0.13	−4.39*	0.09
	(2.40)	(2.21)	(2.43)	(0.30)
Quarter ≥ 2	1.42	−0.69	−7.08***	1.15***
	(2.19)	(1.53)	(2.07)	(0.22)
Inverse	4.26***	2.22***	0.04	−1.00***
	(0.67)	(0.73)	(0.54)	(0.12)
BTC-quoted	4.66***	0.76	−1.99***	−1.61***
	(0.87)	(0.66)	(0.50)	(0.11)
Quanto	3.77***	1.01	−1.79**	−0.95***
	(0.98)	(1.02)	(0.75)	(0.16)
Coin-day FE	Yes	Yes	Yes	Yes
Exchange FE	Yes	Yes	Yes	Yes
Observations	54,926	54,926	54,926	54,926
R^2	19.19%	10.20%	13.32%	66.25%

Table 4: Perpetual futures and market stress using high-frequency data.
This table reports the association between contractual form and market liquidity during periods of market stress. A data point is a minute-level observation of a BTC-USD inverse future on the BitMEX cryptocurrency exchange. The dependent variables are the log of the end of minute quoted spread plus 1 basis point (columns 1 and 3) and the end of minute quoted spread in basis points (columns 2 and 4). The variables of interest are the interaction between market stress measures and indicators for futures contract forms. Past volatility is the absolute log change in spot BTC prices averaged over the preceding 15 minutes. Current volatility is the absolute log spot BTC price change in the current minute. Perpetual is an indicator variable equal to one for perpetual futures contracts and Back-month is an indicator variable equal to one for standard quarterly futures that are not the nearest to expiration future. ***, **, and * denote statistical significance at the 1%, 5%, and 10% levels, respectively, for standard errors based on two-way clustering at the symbol and month level.

	Log quoted spread (1)	Quoted spread bp (2)	Log quoted spread (3)	Quoted spread bp (4)
Perpetual × Past volatility	−4.00***	−179.51***		
	(0.91)	(27.76)		
Back-month × Past volatility	−0.14	154.01		
	(0.50)	(123.91)		
Perpetual × Current volatility			−1.56***	−63.88***
			(0.37)	(9.63)
Back-month × Current volatility			−0.10	57.46
			(0.20)	(44.99)
Back-month	0.53***	6.46*	0.53***	13.09**
	(0.16)	(3.73)	(0.15)	(6.25)
Minute FE	Yes	Yes	Yes	Yes
Symbol FE	Yes	Yes	Yes	Yes
Observations	12,250,302	12,250,314	12,250,302	12,250,314
R^2	11.31%	1.80%	9.97%	1.15%

Appendix A Proofs

A.1. Proof of Proposition 2.2

Proof. For an agent of type $i \in \{M, S\}$ and her time t position of future contract x_t^i, her wealth evolves as:

$$\begin{aligned} dW_t^i =& x_t^i dS_t + x_t^i dB\left(G_t^S, t\right) - x_t^i \delta B\left(G_t^S, t\right) dt \\ =& x_t^i[(G_t^i + r)dt + \sigma_S dY_t^S + B_t(.)dt - \psi G_t^S B_G(.)dt + \frac{1}{2}\sigma_G^2 B_{GG}(.)dt \\ &+ \sigma_G B_G(.)dY_t^G - \delta B(.)dt], \end{aligned} \tag{21}$$

where $B_t(.)$ is the first order derivative with respect to t, $B_G(.)$ is the first order derivative with respect to G_t^S, and $B_{GG}(.)$ is the second order derivative with respect to G_t^S.

For an agent of type i, she chooses the optimal x_t^i to maximize her expected preference:

$$E[dW_t^i - rW_t^i + G_t^S x^i \mathbb{I}[i = S]dt - \frac{\gamma^i}{2} Var(dW_t^i)]. \tag{22}$$

Substituting equation (21), one can solve the optimal x_t^i, for $i \in \{M, S\}$:

$$x_t^i = \frac{G_t^S \mathbb{I}[i = S] + B_t(.) - \psi G_t^S B_G(.) + \frac{1}{2}\sigma_G^2 B_{GG}(.) - \delta B(.)}{\gamma^i \left(\sigma_G^2 B_G^2(.) + \sigma_S^2\right)}. \tag{23}$$

The market clearing condition $x^H + n^S x_t^S + n^M x_t^M = 0$ suggests that

$$x^H \left(\sigma_G^2 B_G^2(.) + \sigma_S^2\right) + (\frac{n^M}{\gamma^M} + \frac{n^S}{\gamma^S})[B_t(.) - \psi G_t^S B_G(.) + \frac{1}{2}\sigma_G^2 B_{GG}(.) - \delta B(.)] + \frac{n^S}{\gamma^S} G_t^S = 0. \tag{24}$$

We conjecture $B\left(G_t^S, t\right) = k(t)G_t^S + j(t)x^H$. Then we have $B_t(.) = k'(t)G_t^S + j'(t)x^H$; $B_G(.) =$

$k(t)$; and $B_{GG}(.) = 0$. Substituting those and canceling out terms, we have

$$\begin{aligned} 0 =& [n^S\gamma^M + (n^S\gamma^M + n^M\gamma^S)(k' - \psi k - \delta k)]G^S \\ &+ [(j' - \delta j)(n^S\gamma^M + n^M\gamma^S) + \gamma^M\gamma^S(\sigma_G^2 k^2 + \sigma_S^2)]x^H. \end{aligned} \tag{25}$$

Because in equilibrium the market clear condition always holds for any G_t^S and x^H. We can solve $k(t)$ and $j(t)$ separately. Start with $k(t)$, given the boundary condition $k(T) = 0$ (because basis converges to 0 as t approaches T), and $n^S\gamma^M + (n^S\gamma^M + n^M\gamma^S)(k' - \psi k - \delta k) = 0$, one can solve the ODE:

$$k(t) = \frac{n^S\gamma^M}{n^S\gamma^M + n^M\gamma^S}\frac{1}{\psi + \delta}\left(1 - e^{(\psi+\delta)(t-T)}\right). \tag{26}$$

Similarly, we can solve $j(t)$ given the the boundary condition $j(T) = 0$ and $(j' - \delta j)(n^S\gamma^M + n^M\gamma^S) + \gamma^M\gamma^S(\sigma_G^2 k^2 + \sigma_S^2) = 0$:

$$j(t) = \frac{\gamma^M\gamma^S(\sigma_G^2 k^2 + \sigma_S^2)}{(n^S\gamma^M + n^M\gamma^S)\delta}(1 - e^{\delta(t-T)}). \tag{27}$$

Thus,

$$P_t = S_t + \frac{n^S\gamma^M}{n^S\gamma^M + n^M\gamma^S}\frac{1}{\psi + \delta}\left(1 - e^{(\psi+\delta)(t-T)}\right)G_t^S + \frac{\gamma^M\gamma^S(\sigma_G^2 k^2 + \sigma_S^2)}{(n^S\gamma^M + n^M\gamma^S)\delta}\left(1 - e^{\delta(t-T)}\right)x^H. \tag{28}$$

To show that the absolute basis is decreasing in both ψ and δ, we only need to show that both $k(t)$ and $j(t)$ are (weakly) decreasing in both ψ and δ. That is equivalent to show that $\frac{1}{x}(1 - e^{-x(T-t)})$ is decreasing in x, for $\forall x > 0$. Taking the first order derivative with respect to x, we obtain:

$$\frac{d}{dx}\frac{1}{x}(1 - e^{-x(T-t)}) = \frac{1}{x^2}[(1 + x(T - t))e^{-x(T-t)} - 1]. \tag{29}$$

We then want to show that $(1 + x(T-t))e^{-x(T-t)} - 1 < 0$. To achieve this, notice that $(1+x(T-t))e^{-x(T-t)} - 1 = 0$ when $x = 0$, and taking the first order derivative with respect to x again, we have $(T-t)e^{-x(T-t)} - (T-t)[1 + x(T-t)]e^{-x(T-t)} < 0$, for $\forall x > 0$. □

A.2. Proof of Corollary 2.3

Proof. Given the equilibrium price equation 28 and market clear condition, we can further obtain x_t^i:

$$x_t^S = -\frac{1}{n^S\gamma^M + n^M\gamma^S}(\gamma^M x^H - \frac{n^M G_t^S}{\sigma_G^2 k^2 + \sigma_S^2}), \tag{30}$$

and

$$x_t^M = -\frac{1}{n^S\gamma^M + n^M\gamma^S}(\gamma^S x^H + \frac{n^S G_t^S}{\sigma_G^2 k^2 + \sigma_S^2}). \tag{31}$$

We then show that the trading volume is increasing in ψ and δ. Because x^H is a constant, the instantaneous trading volume can be characterized as:

$$\begin{aligned} |\frac{dn^S x_t^S}{dt}| =& \frac{n^S n^M}{(n^S\gamma^M + n^M\gamma^S)(\sigma_G^2 k^2 + \sigma_S^2)}|\frac{dG_t^S}{dt}| \\ =& \frac{n^S n^M}{(n^S\gamma^M + n^M\gamma^S)(\sigma_G^2 k^2 + \sigma_S^2)}| - \psi G_t^S dt + \sigma_G dY_t^G|. \end{aligned} \tag{32}$$

Since $k(t)$ is decreasing in both ψ and δ, the instantaneous trading volume increasing in both ψ and δ.

From the proof of Proposition 2.2, the subjective expected utility of speculators S is

$$\begin{aligned} E[\frac{\gamma^S\left(\sigma_G^2 B_G^2(.) + \sigma_S^2\right)}{2}(x_t^S)^2] =& QE[(\gamma^M x^H - \frac{n^M G_t^S}{\sigma_G^2 k^2 + \sigma_S^2})^2] \\ =& Q[(\gamma^M x^H)^2 + (\frac{n^M}{\sigma_G^2 k^2 + \sigma_S^2})^2 \frac{\sigma_G^2}{2\psi}], \end{aligned} \tag{33}$$

where $Q \equiv \frac{\gamma^S\left(\sigma_G^2 B_G^2(.)+\sigma_S^2\right)}{2(n^S\gamma^M + n^M\gamma^S)}$ is a constant. Because $k(t) > 0$ is decreasing in δ, specualtor S's subjective expected utility is increasing in δ. Similarly, one can prove the same result for market maker M.

To prove the last result, first notice that we only need to show that for any standard delivery contract, there exists $\overline{\delta}$ such that the all perpetual future contracts with a refunding rate higher than $\overline{\delta}$ will have a lower $k(t)$. It is straightforward to see that when $\delta = \overline{\delta} \equiv (\frac{1}{1-e^{\psi(t-T)}} - 1)\psi$, the perpetual future contract and the standard delivery future share the same basis. Then the perpetual future contracts have a lower absolute basis and higher open interest if and only if $\delta \geq \overline{\delta}$. $\square$

A.3. Proof of Corollary 2.4

Proof. Let τ be the random birth time. We reuse $k(t)$ and $j(t)$ to denote $k(t \bmod T)$ and $j(t \bmod T)$. Let H^τ denote G shifted by the hedger's random birth time, $H^\tau_s = G_{s+\tau}$. An agent born at time τ almost surely has total cash flow

$$\begin{aligned}\pi(\tau, \{H^\tau\}) =& H^\tau_l k(\tau + l) + j(\tau + l)h - H^\tau_0 k(\tau) - j(\tau)h \\ &- \sum_{n\in\mathcal{Z}:0<Tn-\tau<l} \left(H^\tau_{Tn-\tau}k(0) + j(0)h\right) - \int_\tau^{\tau+l} \delta\left(k(\tau+t)H^\tau_t + j(\tau+t)h\right)dt.\end{aligned}$$

Use $\bar{}$ to denote the realized means of variables between 0 and T. Then

$$\begin{aligned}E[\pi(\tau, \{H^\tau\})|\{H^\tau\}] =& \frac{1}{T}\int_0^T \pi(\tau, \{H^\tau\})dt \\ =& \left(H^\tau_l - H^\tau_0\right)\bar{k} - l\left(\frac{1}{T}\left(\bar{H}k(0) + j(0)h\right) + \delta\left(\bar{k}\bar{H} + \bar{j}h\right)\right).\end{aligned}$$

Which allows us to calculate unconditional moments

$$E[\pi] = E[E[\pi_\tau|\{H\}]] = -lh\left(\frac{1}{T}j(0) + \delta\bar{j}\right),$$

$$
\begin{aligned}
Var[\pi] &= Var[E[\pi_\tau|\{H\}]] + E[Var[\pi_\tau|\{H\}]] \\
&\geq Var[E[\pi_\tau|\{H\}]] \\
&= Var\left[\left(H_l^\tau - H_0^\tau\right)\bar{k} + \bar{H}\frac{l}{T}\left(k(0) + \delta\bar{k}\right)\right],
\end{aligned}
$$

with that inequality binding for perpetual futures.

Fix T for the delivery future and consider a perpetual future with δ' such that

$$
\frac{1}{\psi}\left(1 - \frac{1 - e^{-\psi T}}{\psi T}\right) = \frac{1}{\psi + \delta'},
$$

letting $'$ denote this future's properties. By construction, we have $k(0) > \bar{k} \geq \bar{k}'$ and so $\frac{1}{T}j(0) > \delta j'$. Thus, this perpetual futures contract has a higher expected cash flow without higher variance. $\square$

Appendix B Additional tables and figures

Figure B1: Contractual evolution of BTC futures.
This chart summarizes the futures products offerings of exchanges offering cryptocurrency futures prior to Dec. 2018. Product offering data are sourced from the Wayback Machine - Internet Archive, following a procedure akin to Akey et al. (2022) and by scraping the longest-living BTC forum, BitcoinTalk.org. This chart considers only futures contracts and does not consider exchanges offering margin loans for spot trading (e.g., the early days of Bitfinex) or bilateral bets (e.g., over-the-counter contracts for difference or TeraExchange). This chart excludes ErisX which offered futures trading starting 2019 and was acquired by the CBOE in 2021.

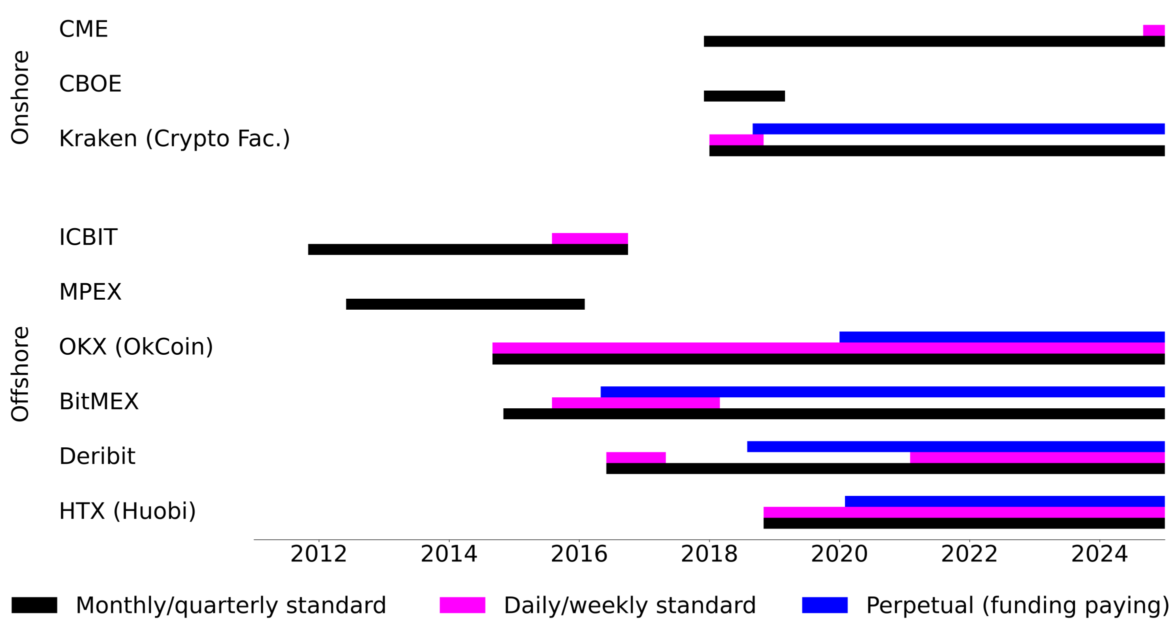

Figure B2: Cryptocurrency futures prices relative to spot prices.
This figure shows the daily average futures basis for the nearest two expiring quarterly BTC-USD futures and the perpetual BTC-USD future on the cryptocurrency exchange BitMEX. The basis for each trade is calculated as the ratio of the futures price to the price of the underlying spot asset, minus one. The spot price is the simple average of the last traded price on Bitfinex, Coinbase, and Coinstamp. The futures price is the mid-point at the end of the minute. The plotted basis is the average of the minutes in two-week periods.

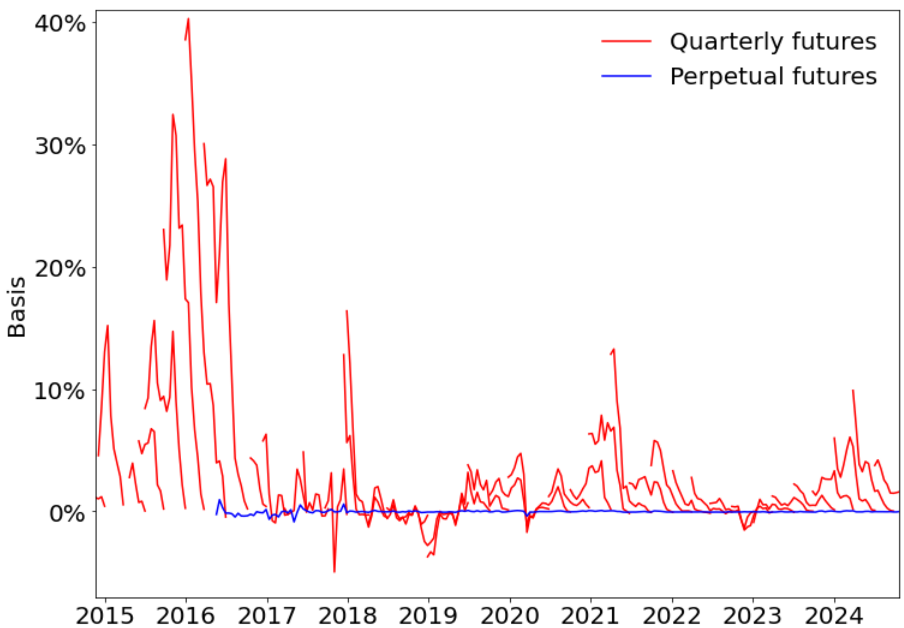

Appendix C Additional Theoretical Results

Our last result concerns the change of the speculator population. Given the high volatility in the cryptocurrency market, liquidity traders are often wiped out of the market, resulting in a large drop of n after a dramatic price move. The following corollary studies the effect of massive speculative position liquidation on future market trading.

Corollary C.1. *The absolute speculative basis decreases in* $\frac{n^M}{n^S}$.

When the market power of liquidity traders drops, the speculative demand in the future market becomes less volatile and the absolute basis shrinks. Consequently, market makers can afford a larger inventory capacity and are willing to trade with a bigger open interest in the future market.

C.1. Proof of Corollary C.1

Proof. From the proof of corollary 2.2, $k(t)$ can be rewritten as

$$\begin{aligned} k(t) =& \frac{n^S\gamma^M}{n^S\gamma^M + n^M\gamma^S}\frac{1}{\psi+\delta}\left(1 - e^{(\psi+\delta)(t-T)}\right) \\ =& \frac{1}{1 + \frac{n^M}{n^S}\frac{\gamma^S}{\gamma^M}}\frac{1}{\psi+\delta}\left(1 - e^{(\psi+\delta)(t-T)}\right). \end{aligned} \tag{34}$$

Then $k(t)$ is decreasing in $\frac{n^M}{n^S}$. □