Implied Impermanent Loss: A Cross-Sectional Analysis of Decentralized Liquidity Pools*

T. N. Lǐ,† S. Naik,‡ A. Papanicolaou,§ L. Schönleber¶

14th March 2025

Abstract

We derive an option-implied valuation of impermanent loss for liquidity providers on decentralized exchanges and quantify it based on traded option prices. We propose a model that values impermanent loss through the variance of the tokens' relative price. Since the relative price is not the price of a traded asset, we introduce a model for the distribution of the former and a valuation formula induced by a change of numéraire. We show that impermanent loss arises from the tokens' individual risks and their correlation risk. These risks negatively impact pool sizes and explain the cross-sectional returns of liquidity pools.

Keywords: Decentralized Exchanges, Decentralized Finance, Impermanent Loss, Derivatives, Risk-Neutral Pricing, Risk Premium, Staking, Yield Farming.
JEL Classification Codes: G10, G11, G13, G20

*We received helpful comments and suggestions from Eric Budish, Faycal Drissi, Campbell Harvey, Stefan Kassberger, Roman Kozhan, Alfred Lehar, Elisa Luciano, Roberto Marfe, Giovanna Nicodano, Julien Prat, Kirill Shakhnov (discussant), Claudio Tebaldi, Fabio Trojani, Harald Uhlig, Grigory Vilkov, Stefan Voigt, and Anthony Lee Zhang (discussant). We thank Mariia Aksenova for helping us access on-chain data. We also thank the participants of the Decentralised Finance Research Group Oxford-Man Institute (DeFOx), the Operations Research Seminar at North Carolina State University, the LTI Research Seminar at the Collegio Carlo Alberto, the ToDeFi 2024, the Market Microstructure and High Frequency Data Conference, the LMU Oberseminar Finanz- und Versicherungsmathematik, the Tech for Finance: AI and Blockchain - Chaire Fintech Dauphine - PSL, and the UCSB-ECON DeFi Seminar. In addition, we thank the FinTech & Digital Finance Chair of Université Paris Dauphine - PSL in partnership with Mazars and Crédit Agricole CIB, and the Avalanche Foundation for financial support.

†Courant Institute of Mathematical Sciences, New York University, 251 Mercer Street, New York City, New York 10012, United States. *thomli@cims.nyu.edu.*

‡Independent Portfolio Manager, *siddharth@equalto.io.*

§Department of Mathematics, North Carolina State University, 2311 Stinson Drive, Raleigh, North Carolina 27695, United States. *apapani@ncsu.edu.*

¶Collegio Carlo Alberto and University of Turin, Piazza Vincenzo Arbarello 8, 10122 Turin, Italy. *lorenzo.schoenleber@carloalberto.org.*

Implied Impermanent Loss: A Cross-Sectional Analysis of Decentralized Liquidity Pools

14th March 2025

Abstract

We derive an option-implied valuation of impermanent loss for liquidity providers on decentralized exchanges and quantify it based on traded option prices. We propose a model that values impermanent loss through the variance of the tokens' relative price. Since the relative price is not the price of a traded asset, we introduce a model for the distribution of the former and a valuation formula induced by a change of numéraire. We show that impermanent loss arises from the tokens' individual risks and their correlation risk. These risks negatively impact pool sizes and explain the cross-sectional returns of liquidity pools.

Keywords: Decentralized Exchanges, Decentralized Finance, Impermanent Loss, Derivatives, Risk-Neutral Pricing, Risk Premium, Staking, Yield Farming.
JEL: G10, G11, G13, G20

1 Introduction

For tokens representing assets, *decentralized exchanges (DEXs)* have become popular, with trading volume reaching 462 billion USD.[1] DEXs enable participants to buy and sell tokens in bilateral markets through liquidity pools. On a DEX, participants can serve as *liquidity providers (LPs)*, supplying liquidity to both sides of a pool and earning fees from token exchanges. The exchange rate, or relative price, is algorithmically set. However, LPs commonly face a form of adverse selection, the so-called *impermanent loss*, resulting from fluctuations in relative prices. Decentralized liquidity provision has been a major driver of growth in *Decentralized Finance*, with market expansion from 500 million USD to 44 billion USD in 2021 in *total value locked (TVL)*.[2] In this paper, we derive a valuation for the primary risk in decentralized liquidity provision, impermanent loss, and quantify it using traded option prices.

Recent studies argue that well-designed DEXs for token and equity trading could save investors billions since they offer less price impact for liquidity takers (Lehar and Parlour (2023)). On a DEX, LPs supply liquidity and earn fees from token exchanges, with the exchange rate algorithmically determined by an *automated market maker (AMM)*. It is well known that LPs face impermanent loss (Heimbach et al. (2022), Capponi and Jia (2024), Li et al. (2024), Milionis et al. (2024), Harvey et al. (2024)), a risk caused by price changes due to token volatility impacting their positions: Initially, liquidity is provided for a specific ratio of the two token. If the price of one token, for example, depreciates, relative to the other, arbitrageurs will sell the depreciated token to the pool. By doing so, they gradually adjust the pool's ratio to match the new relative price. This allows them to profit at the expense of LPs, who end up with a different ratio than they initially deposited, leaving them worse off than if they had simply held the tokens, not providing liquidity.[3]

In this article, we quantify the risk for decentralized liquidity provision, that is the relative price's variance, using a forward-looking risk measure derived from options traded on a centralized exchange. The primary aim is to develop a model-free valuation of impermanent loss based on these option prices. This valuation is a measurement that filters investors' beliefs about future price movements and allows us to quantify a risk

[1] Source: `https://www.ccn.com/news/crypto/uniswap-dexes-reel-record-450b-volume/`

[2] Source: `https://news.anycoindirect.eu/dex-total-value-locked-at-highest-level-in-2-years`. TVL measures adoption by calculating the USD value of coins and tokens locked in a project's smart contracts.

[3] We provide a concrete numerical example of impermanent loss in Section 3.

premium emerging from the difference of the implied and the realized impermanent loss. Specifically, in this paper, we will study how the (option) *implied impermanent loss (IIL)* is related to future risks and investment opportunities. We first show that IIL comprises three (implied) key determinants—the individual volatilities of two tokens, and (third) their correlation. IIL and its components negatively impact pool sizes, and we further examine the risk-return relationship of liquidity provision by analyzing the explanatory power of the IIL, its components, and their risk premia in both the time series and cross-section of returns for LPs. We now discuss these findings in more detail.

We start with an AMM governing token prices in a liquidity pool using the *constant product rule*. We derive and characterize impermanent loss as the primary risk, showing it is -1/8 times the realized variance of the relative price. For a broad class of stochastic volatility models, it depends on token volatilities and their correlation. We refer to these volatilities and correlations as the "drivers" of impermanent loss. We derive an option-implied valuation of impermanent loss, as done in Carr and Madan (1998) and Carr and Wu (2009) for variance swaps. For equities, a variance swap is typically evaluated by replicating the logarithm of the underlying's price relative to the forward price (this is the so-called *log contract*) plus the underlying's total return. The log contract's value can be replicated using a portfolio of European call and put options with various strike prices and the total return is valued by a zero-coupon bond yield. Two primary challenges arise when the underlying is the tokens' relative price: First, the relative price is not the price of a traded asset, so there is no options market where the relative price is the underlying asset. Instead, options are traded only on individual tokens. Second, since the relative price is not a traded asset, a bond yield is not an appropriate proxy for the valuation of the total return. To address these challenges, we first compute a multivariate distribution of the two tokens derived from traded options, and second, utilize a change of numéraire. In general, our methodology also applies for the simpler case when a stablecoin is the numéraire, it applies for the valuation of impermanent loss in Uniswap V3 where LPs are allowed to allocate tokens within specific price ranges and also applies to other measures of adverse selection such as loss-versus-rebalancing (Milionis et al. (2024)).

We estimate an implied multivariate distribution by minimizing the Hansen and Jagannathan (1991) (HJ) bound, a no-arbitrage approach that ensures market consistency. Then, by choosing one token as the numéraire and constructing a new risk-neutral measure, the relative price becomes a martingale and we apply the Carr and Madan (1998) formula. Our framework enables us to value a spread option, which then permits the computation of an *option-implied correlation (IC)* by inverting Margrabe (1978) for-

mula on *implied volatility (IV)* from the individual token options. [4] Prior studies (e.g., Longin and Solnik (2001), Ang and Chen (2002), and Driessen et al. (2009)) suggest that the average equity market correlation (such as IC) is an informative measurement of systematic risk, diversification benefits, or tail risk. Our measure of (implied) correlation transfers this concept to the market for decentralized liquidity provision.

We concentrate on deriving IIL and its drivers for the Bitcoin-Ethereum token pair from options data, given the liquid cryptocurrency derivatives market for Bitcoin (BTC) and Ethereum (ETH). While the IIL and the IC rely on the implied multivariate distribution, the individual risk of tokens can be inferred from the (model-free) IV. Consistent with our framework, we find that the IV of ETH and the IC between BTC and ETH explains more than 60% (and up to 95%) of the daily levels in the IIL. Intuitively, shocks to one underlying increase the IV, and consequently raise IIL. Conversely, higher IC lowers IIL as positive token comovement does not impact relative prices.

Our new implied measures also allow us to quantify an *impermanent loss risk premia (ILRP)* and a *correlation risk premia (CRP)* which we calculate as ex-ante version (defined as the difference between the respective quantity under the risk-neutral and the real-world probability measures). We document that these risk premia vary over time and are, on average, negative. For example, the average IIL of -0.14 exceeds the *realized impermanent loss (RIL)* of -0.12 (both expressed in volatility terms). LPs fear losing comovement because positive token comovement does not affect relative prices, resulting in a lower IC than the *realized correlation (RC)* and hence in a negative CRP. We show that the ILRP arises from the interplay between the *variance risk premia (VRP)* of BTC and ETH and their CRP. In a last step we study the IIL, its drivers, and the risk premia in distressed markets, noting sharp rises, particularly in the tokens' IVs and VRPs. On days when only one token drops sharply, IC falls significantly while the IIL and the ILRP rise.

Lehar and Parlour (2023) show that equilibrium liquidity pool size balances fee revenue against impermanent loss by adjusting pool size, not prices. If the risk of impermanent loss is high, the equilibrium size of the pool is small because small pools have a higher impact for the same order size. Mechanically, a smaller pool size increases the APR of the pool (ceteris paribus). In line, we document that the high risk implied by the risk measures coincides with smaller pool sizes. In the subsequent stage, we utilize IIL and

[4] Our framework enables IC calculation without IV from a cryptocurrency index. Driessen et al. (2009) and Skinzi and Refenes (2005) introduced the average option-implied correlation, with the later literature (Driessen et al. (2016), Buss et al. (2017), Schönleber (2023), Bondarenko and Bernard (2024)) referring to it as equicorrelation.

its drivers as predictors, revealing that higher risks correspond to higher future APRs. The results also hold out-of-sample.

In the last step, we assess the connection between returns for decentralized liquidity provision and our constructed risk factors. Therefore, we empirically analyze a large cross-section of liquidity pools, revealing an average return for LPs, calculated as an annualized percentage rate (APR), of 15%, with peaks of 50%, alongside significant realized impermanent losses averaging -10% and reaching as low as -40%. Overall, our studies suggest impermanent loss as a key risk: A price change in one token alters the relative price, prompting arbitrageurs to exploit the change, which creates impermanent loss for the LP. However, greater risk is rewarded by a larger average return generated by increased trading activity. Our analyses uncover a positive link between innovation ("shocks") in fundamentals measured by implied quantities and returns for LPs. The Fama and MacBeth (1973) regressions confirm a positive, significant risk-return relationship for IIL and its drivers in the cross-section of liquidity pools. In liquidity provision, while the price of risk for IVs is positive, LPs favor high positive token correlations to mitigate impermanent loss, resulting in a negative price of risk for IC. Our results are consistent in a multivariate cross-sectional framework and when controlling for other pool-specific factors.

2 Literature Review

Cousaert et al. (2022) study yield farming frameworks, focusing on protocols and tokens used by aggregators, while Heimbach et al. (2022) analyze the risks and returns of LPs in detail. A continuous-time framework of yield farming from the view of the LP is developed in Li et al. (2024). Milionis et al. (2024) have identified loss-versus-rebalancing as the primary risk for LPs. Cartea et al. (2024) introduce a new comprehensive metric of predictable loss for LPs. Augustin et al. (2023) study LP token staking and the return chasing behavior on PancakeSwap. Lehar and Parlour (2023) show that the equilibrium size of a pool balances the fee revenue against the impermanent loss. We contribute to this literature by developing a continuous-time framework replicating impermanent loss under the risk-neutral measure and establishing a risk-return relationship explaining cross-sectional returns in liquidity provision.

The literature on option-implied information for equity is extensive, see Christoffersen et al. (2013) for an overview. In contrast, the literature on IV for cryptocurrencies is still developing. Alexander and Imeraj (2021) construct a term structure of Bitcoin IV

indices covering maturities from one week to three months. Option-implied information predicting future equity returns is discussed in Driessen et al. (2005), Ang et al. (2006), Bollerslev et al. (2009), Conrad et al. (2013), Bollerslev et al. (2015), Buss et al. (2017), and Schönleber (2023). We contribute to the literature by developing a methodology to apply established equity market metrics, like IC, to cryptocurrency markets, empirically quantifying them, and using option-implied information from cryptocurrency options to predict returns and risks in decentralized liquidity provision.

The literature concerning replicating impermanent loss is limited, with few exceptions. Masaaki Fukasawa and Wunsch (2023) investigate the link between constant function markets, variance swaps, and gamma swaps, and replicate the impermanent loss with a weighted variance swap whenever the numéraire of the token pair is a stablecoin. A similar assumption is made by Clark (2020) and Clark (2021) who argues that participating in liquidity provision and encountering impermanent loss essentially involves taking a short volatility position. We contribute to this literature by developing a methodology for replicating impermanent loss, applicable even when both tokens are non-stablecoins, and empirically quantifying it for the BTC-ETH token pair. This work therefore fills the gap in the literature and connects more broadly the literature on option-implied information to decentralized liquidity provision via the (centralized) derivatives market.

3 Mathematical Formulation for Impermanent Loss

In this section, we formulate the mathematical model of an LP. We start by assuming a liquidity pool with a constant product AMM rule[5] and assume that the underlying token prices follow a stochastic volatility model. These assumptions lead to a characterization of the impermanent loss in terms of the volatility of the relative price for tokens. Before diving into the mathematical derivations, we first offer a high-level overview of how a DEX operates and provide a specific example of impermanent loss.

On a DEX, liquidity takers trade at prices set by mathematical rules like the *constant product rule*, ensuring the total liquidity of the pool stays constant. For two tokens, the rule is $L = \sqrt{N_1 N_2}$, where L is the total liquidity, and N_1 and N_2 represent the amounts of tokens 1 and 2. As N_1 increases, N_2 must decrease, and vice versa. If a liquidity taker

[5]For the sake of simplicity, we do not consider pools with *concentrated liquidity (CL)*, a concept that is introduced by Uniswap V3, enabling LPs to consolidate their pool liquidity within a defined range and earn fees when the spot price enters their specified *active zone*. For a detailed exploration of concentrated liquidity, refer to Heimbach et al. (2022).

wants to obtain 1 unit of token 2 from the pool, the AMM will require Δ_1 units of token 1, where Δ_1 satisfies $\sqrt{(N_1 + \Delta_1)(N_2 - 1)} = L$.

Next, we give an example of impermanent loss: Assume the price of token 1 is 100 USD, and token 2 is a stablecoin. Currently, the liquidity pool contains 900 token 1 and 90,000 token 2. An LP adds 10,000 token 2 and 100 token 1 to the pool, with a total contribution value of 20,000 USD. The pool now contains 100,000 token 2 and 1,000 token 1, with a total pool value of 200,000 USD, of which the LP owns 10%. If the price of token 1 increases in the wider market from 100 USD to 150 USD, traders enter the pool, add token 2, and remove token 1 until the new ratio is 150:1. Because of $x/y = 150$, and $x \times y = 100,000,000$, leading to 122,474.48 token 2 and 816.49 token 1 remaining in the pool, because $122,474.48 \times 816.49 = 100,000,000$ and $122,474.48/816.49 = 150$. Withdrawing the LP's 10% share results in 12,247.4 token 2 and 81.6 token 1, totaling a total value of approximately 24,500 USD. Had the LP simply held their assets (HODL), they would have 10,000 USDT and 1,000 token 1 at 150 USD each, for a total value of 25,000 USD. This results in an impermanent loss of approximately 500 USD.

3.1 Tokens Dynamics Under the Constant Product Rule

Let $N_1(t)$ and $N_2(t)$ denote the amounts of two respective tokens that are paired for trading in a liquidity pool, which we simply refer to as *pool* from now on. The constant product rule for this pool is as follows:

$$L = \sqrt{N_1(t) N_2(t)}, \tag{3.1}$$

where $L > 0$ is a constant. From constant product rule (3.1) a relative price of token emerges,

$$R(t) = \frac{N_2(t)}{N_1(t)} = \text{token 2 per token 1}, \tag{3.2}$$

from which we can deduce the amount of each token in the pool,

$$N_1(t) = \frac{L}{\sqrt{R(t)}}, \qquad N_2(t) = L\sqrt{R(t)}. \tag{3.3}$$

We make the following assumption about the market prices and the relative price for tokens in the pool.

Assumption 3.1. *The relative token price in a pool is equal to the ratio of the market prices of tokens,*

$$R(t) = \frac{P_1(t)}{P_2(t)} = \frac{\textit{dollar per token 1}}{\textit{dollar per token 2}} = \textit{token 2 per token 1}, \tag{3.4}$$

where $P_i(t)$ for $i = 1,\ 2$ are the market prices of tokens outside of the pool.

The rationale for equation (3.4) is that if it were not true, then arbitrageurs would enter the pool and exploit the price discrepancy until it corrected itself. Thus, arbitrageurs will ensure that the ratio of $N_1(t)$ to $N_2(t)$ in the pool maintains a relative price equal to the relative price of the greater market external of the pool.

Our model takes token prices to be given by the following stochastic volatility model,

$$dP_i(t) = \mu_i P_i(t)\, dt + \sigma_i(t)\, P_i(t)\, dB_i(t), \quad i = 1,\ 2, \tag{3.5}$$

where $B_i(t)$ are two correlated standard Brownian motions (SBMs) defined on a filtered probability space $\left(\Omega,\ \mathcal{F},\ (\mathcal{F}_t)_{t\geqslant 0},\ \mathbb{P}\right)$, $dB_1(t)\, dB_2(t) = \rho dt$, $\rho \in [-1,\ 1]$ is the correlation between them, and $\sigma_i(t) > 0$ is the $\mathcal{F}_t$-adapted stochastic volatility process. We apply Itô's lemma to the ratio of $P_1(t)$ and $P_2(t)$ to get the stochastic differential equation (SDE) for the relative price $R(t)$,

$$dR(t) = d\left(\frac{P_1(t)}{P_2(t)}\right) = \mu_R(t)\, R(t)\, dt + \sigma_R(t)\, R(t)\, dB_R(t), \tag{3.6}$$

where $\mu_R(t) = \mu_1 - \mu_2 + \sigma_2^2(t) - \rho\sigma_1(t)\,\sigma_2(t)$, $\sigma_R^2(t) = \sigma_1^2(t) - 2\rho\sigma_1(t)\,\sigma_2(t) + \sigma_2^2(t)$, and $B_R(t) = \frac{\sigma_1(t)}{\sigma_R(t)} B_1(t) - \frac{\sigma_2(t)}{\sigma_R(t)} B_2(t)$ is a SBM under the physical probability measure $\mathbb{P}$.

3.2 Impermanent Loss

Let $\mathcal{V}^{\text{staked}}(t,\ s)$ denote the dollar value staked[6] in the pool at time $s \in [t, \infty)$, and let $\mathcal{V}^{\text{held}}(t,\ s)$ denote the value of an un-staked position that has an equal dollar amount as time t; that is $\mathcal{V}^{\text{staked}}(t,\ t) = \mathcal{V}^{\text{held}}(t,\ t)$. For a time increment $\Delta t > 0$, at time $t + \Delta t$ the impermanent loss is defined as follows.

Definition 3.1. *For a time increment $\Delta t = \frac{1}{n}$ where n is a positive integer, the impermanent loss from time t to time $t + \Delta t$ is the staked value minus the held value, divided*

[6]Depositing liquidity into a pool is called *staking*.

by the held value:

$$\Delta IL_n(t) := \frac{\mathcal{V}^{staked}(t,\, t+\Delta t) - \mathcal{V}^{held}(t,\, t+\Delta t)}{\mathcal{V}^{held}(t,\, t+\Delta t)}. \tag{3.7}$$

In discrete time it is straightforward to show that $\Delta IL_n(t) \leqslant 0$, see Heimbach et al. (2022), and in continuous time it is also true that impermanent loss is always nonpositive. In terms of $\Delta IL_n(t)$, the total impermanent loss up to time t is the summation to time t. Total impermanent loss in continuous time, which we denote as $IL(t)$, is obtained by taking the limit as Δt goes to zero.

Proposition 3.1. *In continuous time, impermanent loss is the differential of total impermanent loss $IL(t)$ is equal to negative one-eighth times the variance of the relative price times the length of time increment,*

$$dIL(t) = -\frac{1}{8}\sigma_R^2(t)\, dt, \tag{3.8}$$

where $\sigma_R(t) = \sqrt{\sigma_1^2(t) - 2\rho\sigma_1(t)\,\sigma_2(t) + \sigma_2^2(t)}$ is the volatility of the relative price $R(t)$ seen in equation (3.6)*.*

Proof. See Appendix A.1. □

4 Model-Free Valuation

Within the stochastic volatility framework given in Section 3, the total impermanent loss is negative one-eighth times the realized variance of the relative price $R(t)$ given in equation (3.6). As shown in Demeterfi et al. (1999), for the time window $[0,\, T]$ the realized variance is equal to twice the total returns plus the log contract,

$$\int_0^T \sigma_R^2(t)\, dt = 2\int_0^T \frac{dR(t)}{R(t)} - 2\ln\left(\frac{R(T)}{R(0)}\right), \tag{4.1}$$

where $-2\ln\left(\frac{R(T)}{R(0)}\right)$ is the so-called log contract with the underlying being $R(T)$, see Carr and Lee (2008). From equation (4.1), we can observe that the valuation of impermanent loss is equivalent to the valuation of a variance swap.

4.1 Risk-Neutral Formulation

Let us consider a risk-neutral probability measure $\mathbb{Q}$ that is equivalent to the physical probability measure $\mathbb{P}$ used in Section 3 (see details in Appendix A.2). Under $\mathbb{Q}$, similar to equation (3.5), over some finite time interval $t \in [0,\, T]$, the SDEs for token prices become

$$dP_i(t) = rP_i(t)\,dt + \sigma_i(t)\,P_i(t)\,dB_i^Q(t), \quad i = 1,\, 2, \tag{4.2}$$

where $B_1^Q(t)$ and $B_2^Q(t)$ are SBMs under $\mathbb{Q}$ with correlation $dB_1^Q(t)\,dB_2^Q(t) = \rho dt$, and r is the risk-free interest rate. Analogous to SDE (3.6), by applying Itô's lemma, the relative price $R(t)$ under the same risk-neutral measure has the following SDE:

$$dR(t) = \mu_R^Q(t)\,R(t)\,dt + \sigma_R(t)\,R(t)\,dB_R^Q(t), \tag{4.3}$$

where $\mu_R^Q(t) = \sigma_2^2(t) - \rho\sigma_1(t)\,\sigma_2(t)$, and $B_R^Q(t) = \frac{\sigma_1(t)}{\sigma_R(t)}B_1^Q(t) - \frac{\sigma_2(t)}{\sigma_R(t)}B_2^Q(t)$ is an SBM under $\mathbb{Q}$. Equation (4.3) demonstrates that $R(t)$ is not a martingale under $\mathbb{Q}$. From here forward, we denote by $\mathbb{E}^Q$ the expectation operator under $\mathbb{Q}$.

Ideally, the $\mathbb{Q}$-measure expected value of equation (4.1) can be computed with a model-free method, giving a purely market-driven prediction of variance. In many pricing problems, the expected total return $\mathbb{E}^Q\left[\int_0^T \frac{dR(t)}{R(t)}\right]$ is equivalent to a short-term zero-coupon bond yield. For the log contract, the formula proposed in Carr and Madan (1998) gives the valuation in terms of a portfolio of European options on $R(T)$; the same valuation can be obtained using the risk-neutral density of $R(T)$ given by the formula proposed in Breeden and Litzenberger (1978). However, there are two main difficulties when performing this model-free valuation for token pairs. The first is that there are not options on the relative prices of tokens. The second is that the relative price is not a martingale under $\mathbb{Q}$, which means that the expected rate of return for $R(t)$ cannot be equated to the short-term interest rate. To manage the first difficulty we compute a joint density from existing options on the separate tokens. To manage the second we employ a change of numéraire.

4.2 Optimal Joint Density

Centralized exchanges offer European options on individual tokens, such as BTC and ETH, which are cash-settled and have margins settled in their respective underlying.

Using the risk-neutral density formula proposed in Breeden and Litzenberger (1978), options with a fixed maturity date and a continuum of strikes provide the risk-neutral distribution (RND) on the future states of the underlying token. For a token pair, the risk-neutral density formula provides the two marginal densities and then a further step must be taken to estimate the dependency structure for their joint RND. This joint distribution is exactly what we would need to value the log contract on $R(T)$.

Let $Call_i(T,\ K)$ denote European call option prices of tokens on their prices $P_i(T)$ with maturity T and strike price of K. The marginal densities under $\mathbb{Q}$ are

$$f_i(T,\ K) = e^{rT}\frac{\partial^2}{\partial K^2}Call_i(T,\ K), \quad \text{for } i = 1,\ 2. \tag{4.4}$$

We can combine the marginal densities $f_1(T,\ K)$ and $f_2(T,\ K)$ to compute a joint RND for $P_1(T)$ and $P_2(T)$ under $\mathbb{Q}$. We assume we are given the joint density under $\mathbb{P}$. Then we estimate the joint pricing kernel (or stochastic discount factor) using an optimization objective based on arbitrage theory. In particular, this estimation must respect the law of one price by preserving market prices, matching marginal densities f_1 and f_2, and avoiding inter-market arbitrage. A financially meaningful objective is to minimize the HJ upper bound, which states that the Sharpe ratio of a portfolio cannot exceed the standard deviation of any pricing kernel. This is expressed as

$$\sup_{\Pi:\mathrm{std}(\Pi)>0}\frac{\mathbb{E}[\Pi]}{\mathrm{std}(\Pi)} \leqslant \inf_{M\in\mathcal{M}(\mu,\nu)}\mathrm{std}(M) \tag{4.5}$$

where Π is the excess return a portfolio, $\mathcal{M}(\mu,\ \nu)$ is the family of pricing kernels for marginal distributions μ and ν, and $\mathrm{std}(\cdot)$ represents the standard deviation. The left-hand side of equation (4.5) represents a portfolio optimization, our primal problem, while the right-hand side represents an optimal pricing problem, our dual problem. If there are Π and M such that $\frac{\mathbb{E}[\Pi]}{\mathrm{std}(\Pi)} = \mathrm{std}(M)$, then there is no duality gap. Equation (4.5) is derived from the Cauchy-Schwartz inequality for the covariance, denoted by $\mathrm{cov}(\cdot)$, between Π and M, $\mathrm{cov}(\Pi,\ M) \geqslant -\mathrm{std}(\Pi)\,\mathrm{std}(M)$. The Cauchy-Schwartz is an equality if and only if Π is a scalar (deterministic) multiple of M. For complete markets, $\mathcal{M}(\mu,\ \nu)$ is a singleton set with a unique M that is replicated by a portfolio of Arrow-Debreu securities, and hence there is no duality gap. For incomplete markets, there is no duality gap if the optimal M is in the span of investible portfolios, which is in general not the case, but when the individual assets each have a complete set of options then the pricing kernel can be replicated, and thus there is no duality gap.

Let $q(x,\ y)$ be a joint density under $\mathbb{Q}$ such that $M(x,\ y) = \frac{q(x,\ y)}{p(x,\ y)}$ is a pricing kernel. Computing the HJ upper bound is a quadratic programming problem:

$$\begin{aligned} \underset{q:q\sim p}{\text{minimize}} \quad & \int\int \left(\frac{q(x,\ y)}{p(x,\ y)}\right)^2 p(x,\ y)\, dxdy, \\ \text{subject to} \quad & \int q(x,\ y)\, dy = f_1(T,\ x), \\ & \int q(x,\ y)\, dx = f_2(T,\ y), \end{aligned} \tag{4.6}$$

where $p(x,\ y)$ is the joint density under $\mathbb{P}$, and $q \sim p$ denotes the set of joint densities such that $q(x,\ y) > 0$ if and only if $p(x,\ y) > 0$ for all $(x,\ y)$. The existence and uniqueness of solutions of a Sharpe-optimal portfolio are proven in Guasoni and Mayerhofer (2020). Trivially, there always exists a solution to equation (4.6) because $\mathcal{M}(\mu,\ \nu)$ always has at least one element, namely the product measure $f_1(T,\ x)\, f_2(T,\ y)$ under which the two tokens are independent. In general, the problem is convex and an approach to solving is to use the method of Lagrange multipliers,

$$\begin{aligned} \mathcal{L} = & \int\int \frac{q^2(x,\ y)}{p(x,\ y)} dxdy - \int \lambda_1(x) \left(\int q(x,\ y)\, dy - f_1(T,\ x)\right) dx \\ & - \int \lambda_2(y) \left(\int q(x,\ y)\, dx - f_2(T,\ y)\right) dy, \end{aligned} \tag{4.7}$$

where $\lambda_1(x)$ and $\lambda_2(y)$ are Lagrange multiplier functions. From equation (4.7), the unique non-negative solution to (4.6) is expressed as $q(x,\ y) = \frac{1}{2}(\lambda_1(x) + \lambda_2(y))\, p(x,\ y)$, and the optimal pricing kernel is

$$M(x,\ y) = \frac{1}{2}(\lambda_1(x) + \lambda_2(y)). \tag{4.8}$$

Assuming that every point $(x,\ y)$ has a positive physical probability of occurring, then the pricing kernel given by equation (4.8) is arbitrage-free if and only if $\lambda_1(x) + \lambda_2(y)$ is strictly positive. The Lagrangian given by equation (4.7) is written without a non-negativity constraint because the avoidance of arbitrage requires a solution to be in the interior, in which case this additional constraint would be inactive. It can also be shown that there is replicability of the Lagrange multiplier functions $\lambda_1(x)$ and $\lambda_2(x)$ using a portfolio of European call and put options, forward contracts, and cash. The formal statements for these results are given in Appendix B.

Using Proposition I.2, we can construct a portfolio that is a (negative) scalar multiple

of the pricing kernel and therefore has a Sharpe ratio equal to the the right-hand side of equation (4.6). That is, with cash, forward contracts, and European call and put options, we construct a portfolio $\mathit{\Pi}$ whose excess returns are

$$\mathit{\Pi}\left(x,\, y\right) = -\frac{1}{2}\left(\lambda_1\left(x\right) + \lambda_2\left(x\right)\right) = -M\left(x,\, y\right).$$

For this $\mathit{\Pi}$ we have $\frac{\mathbb{E}[\mathit{\Pi}]}{\text{std}(\mathit{\Pi})} = \text{std}\left(M\right)$, confirming that there is no duality gap in equation (4.5). For implementation details, Appendix B shows a sparse quadratic program when $f_1\left(T,\, K\right)\Delta K$ and $f_2\left(T,\, K\right)\Delta K$ are the inputs.

After optimizing the HJ bound, the obtained density can price European options on $R\left(T\right) = \frac{P_1(T)}{P_2(T)}$. These prices can be inserted into the formula of Carr and Madan (1998) for valuing the log contract. Hence, we have addressed the issue highlighted before, namely, the absence of options prices on $R\left(T\right)$.

4.3 Change of Numéraire

The relative price $R\left(t\right)$ is a marginal rate of substitution for a pool, and so unless the base token is pegged to the local currency it is not accurate to say that $R\left(t\right)$ grows at the risk-free rate. Specifically, it is difficult to calculate the expectation of the total-return term $\int_0^T \frac{dR(t)}{R(t)}$ of equation (4.1). A change of numéraire resolves this issue and simplifies the computation for the swap valuation. This change of numéraire should be carried out with the base token, which in our case is $P_2\left(t\right)$. For example, we could make a change of numéraire to ETH.

We consider a probability measure $\widetilde{\mathbb{Q}}$ that is equivalent to the risk-neutral probability measure $\mathbb{Q}$ used in equation (4.2), which is defined through a Radon–Nikodym density:

$$\left.\frac{d\widetilde{\mathbb{Q}}}{d\mathbb{Q}}\right|_T := \frac{P_2\left(T\right)}{\mathbb{E}^Q\left[P_2\left(T\right)\right]} = \frac{P_2\left(T\right)}{e^{rT}P_2\left(0\right)}. \tag{4.9}$$

By the change of numéraire and Girsanov theorem explained in Appendix A.2, under this new probability measure $\widetilde{\mathbb{Q}}$, $\widetilde{B}_1^Q\left(t\right) := B_1^Q\left(t\right) - \rho\int_0^t \sigma_2\left(s\right)ds$ and $\widetilde{B}_2^Q\left(t\right) := B_2^Q\left(t\right) - \int_0^t \sigma_2\left(s\right)ds$ are SBMs. Then by implementing Itô's quotient formula, we get the SDE for the relative price under $\widetilde{\mathbb{Q}}$,

$$dR\left(t\right) = R\left(t\right)\left(\sigma_1\left(t\right)d\widetilde{B}_1^Q\left(t\right) - \sigma_2\left(t\right)d\widetilde{B}_2^Q\left(t\right)\right). \tag{4.10}$$

From (4.10) it is obvious that $R(t)$ is a martingale under $\widetilde{\mathbb{Q}}$. From here forward, we denote by $\widetilde{\mathbb{E}}^Q$ the expectation operator under $\widetilde{\mathbb{Q}}$ with $P_2(t)$ as the numéraire.

Proposition 4.1. *Under the probability measure $\widetilde{\mathbb{Q}}$ defined by equation (4.9), the valuation of the variance swap with underlying being the relative price $R(t)$ given by equation (4.3), is equal to the expectation of the log contract on $R(t)$, $\widetilde{\mathbb{E}}^Q\left[\int_0^T \sigma_R^2(t)\,dt\right] = \widetilde{\mathbb{E}}^Q\left[-2\ln\left(\frac{R(T)}{R(0)}\right)\right]$. Consequently, under the same probability measure, the expectation for the impermanent loss given by equation (3.8) is equal to $-\frac{1}{8}$ times the expectation of the log contract:*

$$\widetilde{\mathbb{E}}^Q\left[IL(T)\right] = \frac{1}{4}\widetilde{\mathbb{E}}^Q\left[\ln\left(\frac{R(T)}{R(0)}\right)\right]. \tag{4.11}$$

Proof. See Appendix A.3. □

For the valuation of variance swap on $R(t)$ under $\widetilde{\mathbb{Q}}$, the individual components for the formula proposed in Carr and Madan (1998) or Breeden and Litzenberger (1978) are the following European call and put options on $R(t)$,

$$\begin{aligned}
\widetilde{\mathbb{E}}^Q\left[(R(T)-K)^+\right] &= \frac{e^{-rT}}{P_2(0)}\int\int y\left(\frac{x}{y}-K\right)^+ q(x,\,y)\,dxdy, \\
\widetilde{\mathbb{E}}^Q\left[(K-R(T))^+\right] &= \frac{e^{-rT}}{P_2(0)}\int\int y\left(K-\frac{x}{y}\right)^+ q(x,\,y)\,dxdy.
\end{aligned} \tag{4.12}$$

By using equation (4.12), we can write the expectation of equation (4.11) as follows:

$$\widetilde{\mathbb{E}}^Q\left[IL(T)\right] = -\int_{R(0)}^{\infty}\frac{1}{4K^2}\widetilde{\mathbb{E}}^Q\left[(R(T)-K)^+\right]dK - \int_0^{R(0)}\frac{1}{4K^2}\widetilde{\mathbb{E}}^Q\left[(K-R(T))^+\right]dK.$$

Our methodology also applies to the valuation of impermanent loss if one (or both) of the tokens is a stablecoin, in Uniswap V3 and other adverse selection metrics, such as loss-versus-rebalancing (LVR) (Milionis et al. (2024)). We outline the details below and refer to the Appendix C for more details.

4.4 Implied Impermanent Loss with a Stablecoin as Numéraire

If one of the tokens is a stablecoin then the valuation of impermanent loss can be significantly simplified. The simplest case is when the base token $P_2(t)$ is pegged to a currency that is also the currency unit for risk-neutral prices under $\mathbb{Q}$. In this case, if interest rate r is constant then $P_2(t) = e^{rt}P_2(0)$ with zero volatility, our valuation

of impermanent loss reduces to the standard variance swap formula for a non-dividend-paying-asset, and there is no need to use a change of numéraire. Another case is when the base token is a non-stablecoin (e.g., ETH) and $P_1(t)$ is a stablecoin for a currency other than the currency under $\mathbb{Q}$. In this case, there is not much simplification because $P_1(t)$ is the (risky) exchange rate. Hence, the change of numéraire identified in Section 4.3 is still the best way to value impermanent loss. If both tokens are stablecoins in a currency different from $\mathbb{Q}$, the volatilities $\sigma_1(t)$ and $\sigma_2(t)$ will be small but non-zero. In this case, we again use the change of numéraire from Section 4.3. Hence, our approach has flexibility for the types of token pairs traded. Finally, if interest rates are stochastic—meaning the risk-neutral measure includes a stochastic short-term interest rate process $r(t)$ adapted to $\mathcal{F}_t$—calculating variance swaps becomes significantly more complex.

4.5 Implied Impermanent Loss for V3 and Implied Loss-Versus-Rebalancing

Under the Uniswap V3 protocol, LPs allocate tokens within a specific price band. For a given band C_i, if $R(t) \in C_i$ then the LP earns a reward but is also exposed to impermanent loss. In particular, we show in Appendix C that the pool-wide impermanent loss is a weighted sum of individual corridor variance swaps, with weights proportional to the TVL in each band.

Our framework can also be used to assess loss-versus-rebalancing (LVR), with the advantage of valuing LVR under $\mathbb{Q}$ introduced in Section 4. Using the weighted variance swap framework of Lee (2010b), we show that the implied LVR can be computed using the estimated RND obtained by minimizing the HJ bound, as proposed in Section 4.2.

5 Data Description

This section provides a general description of data sources, the construction of the option-implied variables, and the variables for pools.

5.1 Option Data and Risk-Neutral Moments

We obtain option data through amberdata[7], encompassing hourly implied volatility surfaces for a standardized maturity of 30 days and different moneyness levels. The data is

[7] https://www.amberdata.io/

collected from Deribit, which is the largest exchange for trading cryptocurrency options.[8] On Deribit options are cash-settled and quoted directly in the respective cryptocurrency and not in USD. We consider BTC and ETH as underlying assets. The surface data consists of five moneyness levels (measured by delta) per call and put option.

We interpolate across moneyness to obtain a smooth surface which allows us to extract the RND for each underlying following equation (4.4) proposed in Breeden and Litzenberger (1978).

$$f_i\left(T,\ K;\ \theta\right) = e^{rT}\frac{\partial^2}{\partial K^2} Call_i^{BS}\left(T,\ K,\ \widehat{\sigma}_i\left(T,\ K;\ \theta\right)\right), \quad \text{for } i = 1,\ 2,$$

where $\widehat{\sigma}_i\left(T,\ K;\ \theta\right)$ is a function with parameter θ that is fitted to the IV smile of quoted options, and where $Call_i^{BS}\left(T,\ K,\ \widehat{\sigma}_i\right)$ is the Black-Scholes European call option price on $P_i(T)$ with IV $\widehat{\sigma}_i$. It is shown in Figlewski (2018) that improved estimation of the risk-neutral density is obtained by parametrically fitting the IV smile and then differentiating the option price. In particular, a polynomial of degree four or five is usually sufficient for fitting quoted IVs, and some type of heavy-tailed parametric distribution is used to extrapolate into the tail. Figure 5.1 shows the 4th-order polynomial fitted to the quoted BTC IVs with $T = 30$ days on the 18th of May 2023 at the beginning of the day in UTC; the tails are extrapolated with a log-normal density with volatility parameter taken to be the very last quoted IV. The 4th-order polynomial fit is convenient because the implied density follows from a simple application of the chain rule when computing the second K partial derivative of $Call_i^{BS}\left(T,\ K,\ \widehat{\sigma}_i\left(T,\ K;\ \theta\right)\right)$.

5.1.1 Implied Impermanent Loss

As shown in Section 4.3, we can construct a new measure as $\frac{d\widetilde{\mathbb{Q}}}{d\mathbb{Q}}\Big|_T = \frac{P_2(T)}{\mathbb{E}^Q[P_2(T)]} = \frac{P_2(T)}{e^{rT}P_2(0)}$ to obtain a valuation of the impermanent loss by using equation (4.11),

$$\widetilde{\mathbb{E}}^Q\left[IL\left(T\right)\right] = \frac{1}{4}\widetilde{\mathbb{E}}^Q\left[\ln\left(\frac{R\left(T\right)}{R\left(0\right)}\right)\right] = \frac{1}{4}\int \ln\left(\frac{x}{yR\left(0\right)}\right)\widetilde{q}_T\left(x,\ y\right)dxdy,$$

where $\widetilde{q}_T\left(x,\ y\right)$ is the optimal joint density at time $t = T$ as described in Section 4.2, computed with options of maturity date T. For a discrete mesh $\left\{\left(x_i,\ y_j\right)_{i,j}\right\}$ with increments Δx and Δy between mesh points, upon which we have a mass function $\widetilde{q}_{i,j} = y_j \times q_{i,j}/F_y$ where $q_{i,j} \approx q_T\left(x_i,\ y_j\right)\Delta x \Delta y$ and $F_y = \sum_{i,j} y_j q_{i,j}$ is the forward price on $P_2\left(t\right)$. Then

[8] The options trading volume across different exchanges is displayed in Figure D.1.

(a) BTC – Implied Volatility

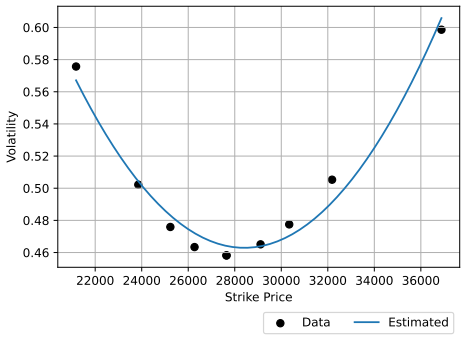

(b) BTC – Risk-Neutral Density

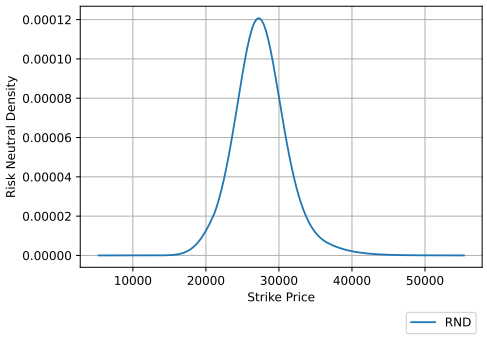

(c) BTC and ETH – Risk-Neutral Joint Density

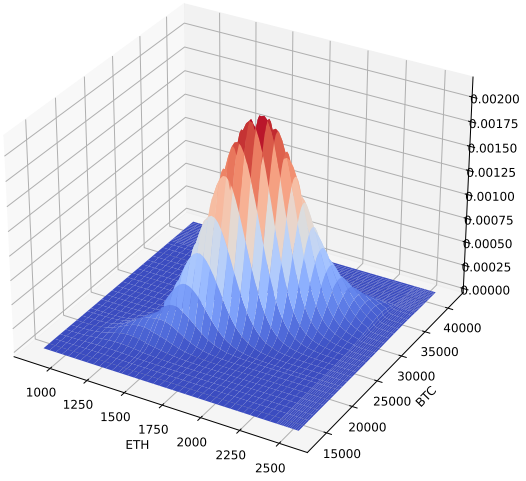

Figure 5.1: BTC Implied Volatility and Risk-Neutral Joint Densities. Panel (a) shows the implied volatility surface for BTC. Panel (b) shows the RND for BTC. The quoted volatilities are fitted with a fourth-order polynomial, and extrapolation of the tail densities uses a log-normal density. Panel (c) shows the risk-neutral joint density for BTC and ETH, which is obtained as described in Section 4.2. The date is the 18th of May 2023.

there is the following estimate of the implied impermanent loss,

$$\widetilde{\mathbb{E}}^{Q}\left[IL\left(T\right)\right] \approx \frac{1}{4}\sum_{i}\sum_{j}\widetilde{q}_{i,j}\ln\left(\frac{x_i}{y_j R\left(0\right)}\right). \tag{5.1}$$

We take the square root of the result calculated by equation (5.1) to express this quantity in terms of volatility and annualize it by multiplying it by $\sqrt{365/30}$. In the empirical analysis, we denote (the absolute value of) the annualized implied impermanent loss of the BTC-ETH token pair in volatility terms as $IIL = \sqrt{\frac{365}{30}\widetilde{\mathbb{E}}^{Q}\left[IL\left(T\right)\right]}$.

5.1.2 Implied Variances

We calculate the implied variances for each underlying for a maturity of $T = 30$ days following Martin (2016) (for the gross return) as

$$IV^2(t) = \frac{2(1+r(t))}{(T-t)\,P^2(t)} \left[\int_{F(t)}^{\infty} Call(t,\, K;\, T)\, dK + \int_0^{F(t)} Put(t,\, K;\, T)\, dK \right].$$

where $P(t)$ is the current token price, $r(t)$ the risk-free rate and $F(t) = \mathbb{E}^Q[P(T)] = e^{rT}P(t)$ is its forward price. In the computations, we discretize the integral via the trapezoidal rule. Finally, we take the square root to get the IV.

5.1.3 Implied Correlation

By using the our optimal joint density $\widetilde{q}_T(x,\, y)$, we calculate the *implied correlation (IC)* between BTC and ETH for a maturity of $T = 30$ days simply as the Pearson correlation coefficient:

$$IC = \frac{\mathrm{Cov}(x, y)}{\sigma_x \sigma_y}$$

5.2 Realized Quantities

We compute the annualized realized volatilities (RV) for BTC and ETH as their annualized standard deviations using hourly return data from CryptoCompare, employing a rolling window approach over 30 calendar days. We repeat the procedure for calculating the realized correlation (RC), as the Pearson correlation coefficient, for BTC and ETH. The calculation of the RIL is outlined in equation (5.3).

5.3 Risk Premia

We calculate the (ex-ante) risk premia defined as the difference between the respective quantity under the risk-neutral (implied) and the real-world (realized) probability measures. Hence the variance risk premium for BTC (VRP_{BTC}) is calculated as IV_{BTC} - RV_{BTC}. The same calculation is made for ETH. Note that we consider the volatilities and not the variance in the calculations, hence the VRP strictly speaking represents a volatility risk premium. The CRP is calculated as the IC-RC. Since we defined both IIL and RIL in positive terms, but they are per definition negative quantities, the impermanent loss risk premia ($ILRP$) is defined as $ILRP = -IIL - (-RIL)$.

5.4 Liquidity Pool Data

We analyze daily data from Uniswap V3, the third version of the Uniswap DEX. We source data using The Graph, a protocol that enables access to Ethereum blockchain data via a query language called GraphQL. We only consider pools where at least one of the tokens is either BTC or ETH. We filter out pools with an average TVL lower than 1,000,000 and fewer than 100 days of daily observations, leaving us with 167 pools. We focus on the prices of both tokens in the pool, the accumulated fees collected from the liquidity takers earned by LPs, and the TVL in the pool. We then calculate the impermanent loss and historical annual percentage rates (APRs) for each pool from May 2021 to February 2025.

For the calculation of the daily realized impermanent loss we follow the equation given in Heimbach et al. (2022):

$$IL\left(t\right) \coloneqq \frac{V_{Pool}\left(t\right) - V_{Hold}\left(t\right)}{V_{Hold}\left(t\right)} = \frac{2\sqrt{\frac{R(t)}{R(t-\Delta t)}}}{1+\frac{R(t)}{R(t-\Delta t)}} - 1, \tag{5.2}$$

where $V_{Hold}\left(t\right)$ is the value of the LP's assets when the relative price is $R\left(t-\Delta t\right)$ at the initial time $t-\Delta t$, and $V_{Pool}\left(t\right)$ is the value of the LP's assets after the relative price changes to $R\left(t\right)$ at the later time t. Equation (5.2) is equivalent to the definition of impermanent loss given in equation (3.7), see Li et al. (2024).

We annualize the daily realized impermanent loss by raising $\left(1+IL\left(t\right)\right)$ to the power of 365. Given that this quantity is subject to noise, we then apply a 30-day moving average. To compare with the implied impermanent loss, we take the square root of the annualized realized impermanent loss to express it as volatility. For the BTC-ETH token pair, we denote this annualized realized impermanent loss by RIL:

$$RIL\left(t\right) = \sqrt{\left[1+IL\left(t\right)\right]^{365}}. \tag{5.3}$$

The estimated daily returns for liquidity provision are quoted in terms of APRs, which denote the annualized fraction of fees collected from the liquidity takers ($fees\left(t\right)$) over some period divided by the total volume of the pool, with the latter being measured in TVL. Hence, we calculate the APR as follows:

$$APR\left(t\right) = \frac{fees\left(t\right)}{TVL\left(t\right)}. \tag{5.4}$$

6 Empirical Analysis

In this section, we first compare the IIL and its components for the BTC-ETH token pair using options data, with realized counterparts from historical data. Next, we analyze the IIL and its components, finding that higher IVs increase IIL, while higher IC decreases it due to the reduced impact of positive token comovement on relative prices. We find that the $ILRP$ becomes more negative as the CRP decreases and the VRPs increase. Using a large cross-section of liquidity pools where either Bitcoin or Ethereum is one of the tokens of the respective token pair, we show that our measures are negatively related to the pool sizes. Hence a high level of risk corresponds to a low pool size. In line, we show that our implied variables positively predict future aggregated APRs in the time series. We investigate liquidity provision returns in the cross-section, emphasizing impermanent loss as a critical risk factor shaping the risk-return relationship. Token price changes lead to increased trading activity resulting in higher APR for LPs but also trigger impermanent loss. The Fama and MacBeth (1973) cross-sectional regressions confirm the significance of the IIL and its drivers with consistent results in a multivariate setting. LPs favor high positive correlations among tokens to mitigate impermanent loss, resulting in a negative price of risk for IC.

6.1 Implied Impermanent Loss and Its Drivers

Figure 6.1 displays the (absolute) implied variables together with their realized counterparts. From Panel (a) it is visible that the IIL co-moves with RIL, which is a noisy time series if not smoothed.[9] IV and RV for BTC in Panel (b) and ETH in Panel (c) display a strong comovement. Panel (d) depicts the dynamics of IC together with the RC. IC, being a bounded variable, displays (relatively) smaller fluctuations than the IVs.

Table 6.4 reports summary statistics for the implied variables. The IIL is about 0.135. The second and third columns represent IV_{BTC} and IV_{ETH} with mean values of 0.648 and 0.756 and standard deviations of 0.161 and 0.223, respectively. The average value of 0.818 for the IC, suggests a strong average correlation between the price movements of BTC and ETH as perceived by the derivatives market. For the realized variables the RIL is approximately 0.123. In terms of realized volatility, the mean values for RV_{BTC} and RV_{ETH} in the fifth and sixth columns are 0.555 and 0.692. Notably, RC exhibits

[9]The RIL smoothed with only 7 or 2 days is displayed in Figure D.2. As visible, the RIL can be substantially larger than the IIL.

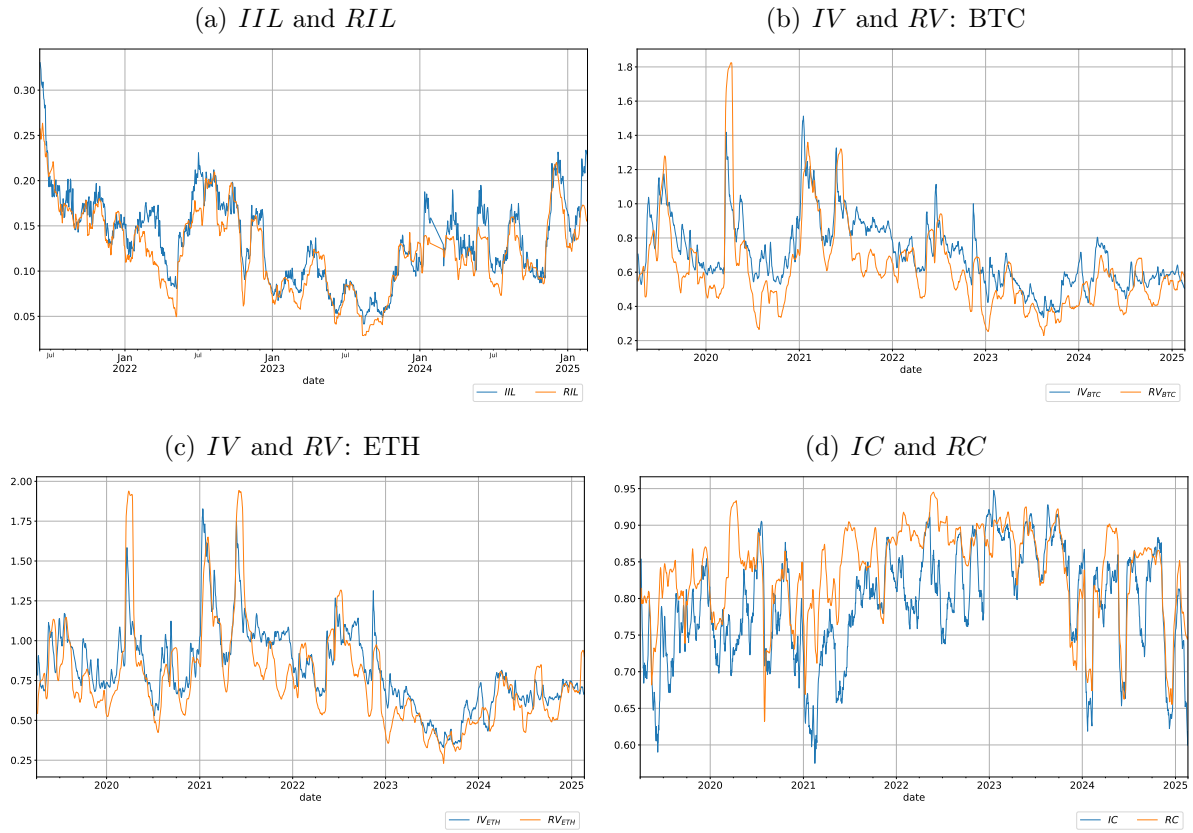

Figure 6.1: Implied and Realized Variables – BTC and ETH. The figure shows the IIL and RIL for BTC and ETH in Panel (a), the IV and RV for BTC and ETH in Panel (b) and Panel (c), and the IC and RC for BTC and ETH in Panel (d). The implied quantities are calculated by the equations in Section 5.1. The realized quantities are calculated as discussed in Section 5.2. The data is sampled daily, and the sample period is from 2021-05 to 2025-02.

an average of 0.854 for BTC and ETH, as measured from realized data. In summary, there are differences in the mean values derived from options data or realized data for the impermanent loss, volatilities, and correlation.

Table 6.2 presents the time-series correlation of the various variables in levels. The correlation between the IIL and RIL is about 0.913. Calculating the RIL using a shorter moving average, such as 7 or 2 days, reduces the correlation to 0.615 or 0.402, respectively. Naturally, there exists a strong comovement between the IVs and the RVs (the correlation ranges from $0.75 - 0.80$). The correlation between IC and RC is 0.655. Interestingly, the correlations between the IC and the IVs are negative; for example, the correlation with IV_{BTC} is -0.172 and similar for IV_{ETH} -0.212; the correlation with IIL is -0.772. The correlation between the IVs and the IIL is 0.665 (0.751) for BTC (ETH).

	IIL	IV_{BTC}	IV_{ETH}	IC	RIL	RV_{BTC}	RV_{ETH}	RC
Mean	0.135	0.648	0.756	0.818	0.123	0.555	0.692	0.854
Std	0.048	0.161	0.223	0.069	0.044	0.174	0.261	0.061
Per 10	0.075	0.456	0.465	0.729	0.065	0.361	0.403	0.762
Median	0.136	0.608	0.712	0.828	0.122	0.545	0.657	0.873
Per 90	0.194	0.879	1.069	0.899	0.176	0.727	0.984	0.905

Table 6.1: Summary Statistics – Implied and Realized Variables. The table reports the summary of the statistics time-series mean, median, standard deviation, and the 10% and 90% percentiles of the measures. The implied quantities are calculated by the equations in Section 5.1. The realized quantities are calculated as discussed in Section 5.2. The sample period is from 2021-05 to 2025-02.

The pronounced interplay of positive and negative relationships among the components, with the IIL, highlights the importance of factoring in the univariate and multivariate distribution of individual tokens during the estimation process.

	IIL	IV_{BTC}	IV_{ETH}	IC	RIL	RV_{BTC}	RV_{ETH}	RC
IIL	1.000	0.665	0.751	-0.772	0.913	0.789	0.852	-0.238
IV_{BTC}	0.665	1.000	0.938	-0.172	0.687	0.773	0.761	0.227
IV_{ETH}	0.751	0.938	1.000	-0.212	0.758	0.749	0.802	0.201
IC	-0.772	-0.172	-0.212	1.000	-0.661	-0.448	-0.480	0.641
RIL	0.913	0.687	0.758	-0.661	1.000	0.735	0.793	-0.238
RV_{BTC}	0.789	0.773	0.749	-0.448	0.735	1.000	0.958	0.262
RV_{ETH}	0.852	0.761	0.802	-0.480	0.793	0.958	1.000	0.214
RC	-0.238	0.227	0.201	0.641	-0.238	0.262	0.214	1.000

Table 6.2: Contemporaneous Correlations. The table reports the correlation coefficients for the measures. The implied quantities are calculated by the equations in Section 5.1. The realized quantities are calculated as discussed in Section 5.2. The sample period is from 2021-05 to 2025-02.

In our empirical evaluation of the impact of the drivers on IIL, we employ straightforward regressions using daily levels of IIL and corresponding levels on its drivers. To underscore the individual contributions of each component, we initially isolate the most crucial variable, followed by the second and third most influential variable, and the most influential pairs of variables. Given the strong correlation between IV_{ETH} and IV_{BTC}, we refrain from including these variables together. The outcomes of these analyses are presented in Table 6.3: The majority of IIL is accounted for by IV_{ETH}, exhibiting R^2 values of nearly 70% while IV_{BTC} only explains 59%. When examining the outcomes

with two regressors, the inclusion of IC increases the R^2 by approximately 25% (30%) compared to the univariate regression results for IV_{ETH} (IV_{BTC}). Following our stylized framework, IV_{ETH} and the IC have opposing effects on the IIL.[10]

	(1)	(2)	(3)	(4)	(5)
Intercept	-0.018	-0.021	0.706	0.414	0.440
	0.000	0.000	0.000	0.000	0.000
IV_{BTC}	–	0.254	–	–	0.171
	–	0.000	–	–	0.000
IV_{ETH}	0.215	–	–	0.151	–
	0.000	–	–	0.000	–
IC	–	–	-0.692	-0.478	-0.507
	–	–	0.000	0.000	0.000
R^2	68.743	59.962	67.524	94.865	89.835

Table 6.3: Drivers of the Implied Impermanent Loss. The table reports the results of slope, p-value, and the R^2 of the described regression procedure. We report the specifications with one ((1)-(3)) and two ((4)-(5)) independent variables. The implied quantities are calculated by the equations in Section 5.1. The data is sampled daily from 2021-05 to 2025-02.

6.2 Risk Premia

We investigate the risk premia of our measures, which we calculate ex-ante as the implied minus the realized quantity following Section 5.3. Table 6.4 reports summary statistics and Figure 6.2 displays the risk premia over time. The $ILRP$ is on average -0.012 with a 10% quantile value of -0.04 and an absolute minima of -0.08.[11] The average VRP_{BTC} (VRP_{ETH}) is 0.09 (0.06), with a low 10% quantile value of -0.052 (-0.111). The absolute minima of both VRP_{BTC} and VRP_{ETH} can reach up to -0.80. The CRP is on average -0.037 with a 10% quantile value of -0.120. The LP favors a strong token correlation and is concerned about its decline, contrary to an equity investor. Consequently, the CRP tends to be negative on average, as the IC derived from option prices is lower than RC, signaling greater risk. As visible from Figure 6.2, the sharp declines in the $ILRP$ and VRPs often occur simultaneously. However, there

[10] We have carried out the same exercise for changes in the IIL and its components. Doing so leads to comparable results but with worse explanatory power.

[11] When not applying a 30-day rolling window on the RIL the ILRP decreases to -0.042 and a 10% percentile of -0.14. When calculating the risk premia as the difference between implied variance and realized variance, and then expressing this difference in volatility terms, the risk premia increase. In this case, the ILRP is nearly 0.06, while the VRP for BTC and ETH is approximately 0.3.

are instances where drops in the $ILRP$ do not align with those in the VRPs, and vice versa. In addition, the $ILRP$ and the CRP comove, especially in the second half of the sample.[12]

	$ILRP$	VRP_{BTC}	VRP_{ETH}	CRP
Mean	-0.012	0.090	0.060	-0.037
Std	0.019	0.115	0.159	0.057
Per 10	-0.038	-0.052	-0.111	-0.120
Median	-0.010	0.091	0.065	-0.025
Per 90	0.011	0.238	0.249	0.030

Table 6.4: Summary Statistics – Risk Premia. The table reports the summary of the statistics time-series mean, median, standard deviation, and the 10% and 90% percentiles of the risk premia. The implied quantities are calculated by the equations in Section 5.1. The realized quantities are calculated as discussed in Section 5.2. The sample period is from 2021-05 to 2025-02.

Next, we investigate which drivers' risk premia influence $ILRP$. To do so, we repeat the method conducted in Section 6.1 by running simple regressions using daily $ILRP$ values and the corresponding risk premia of the drivers ($VRPs$ and CRP). To underscore the individual contributions of each component, we initially isolate the most crucial variable, followed by the second and third most influential variable. The outcomes of these analyses are presented in Table 6.5: The majority of $ILRP$ is accounted for by the CRP, exhibiting R^2 values of nearly 20% while VRP_{BTC} and VRP_{BTC} only explain 6%. Focusing on the multivariate regressions (4) and (5): The coefficient for the CRP loads positively on the $ILRP$. Hence, an increase in RC (for a fixed IC) decreases the CRP and coincides with a decrease in the $ILRP$ (a more negative IIL as opposed to the RIL). The coefficients for the VRPs are negative: A larger VRP (more risk) decreases the $ILRP$.

6.3 Bear and Bull Markets and Implied Variables

Next, we investigate the dynamics of our measures in bear and bull markets. We start with the bear markets, the two worst days for BTC and ETH over our sample were on the 14th of July 2022 and the 10th of November 2022 when BTC (ETH) experienced a

[12]For equities, see Carr and Wu (2009) and Bollerslev et al. (2009) for evidence on the VRP, and refer to Driessen et al. (2009) and Buraschi et al. (2014), Bondarenko and Bernard (2024) for evidence on the CRP.

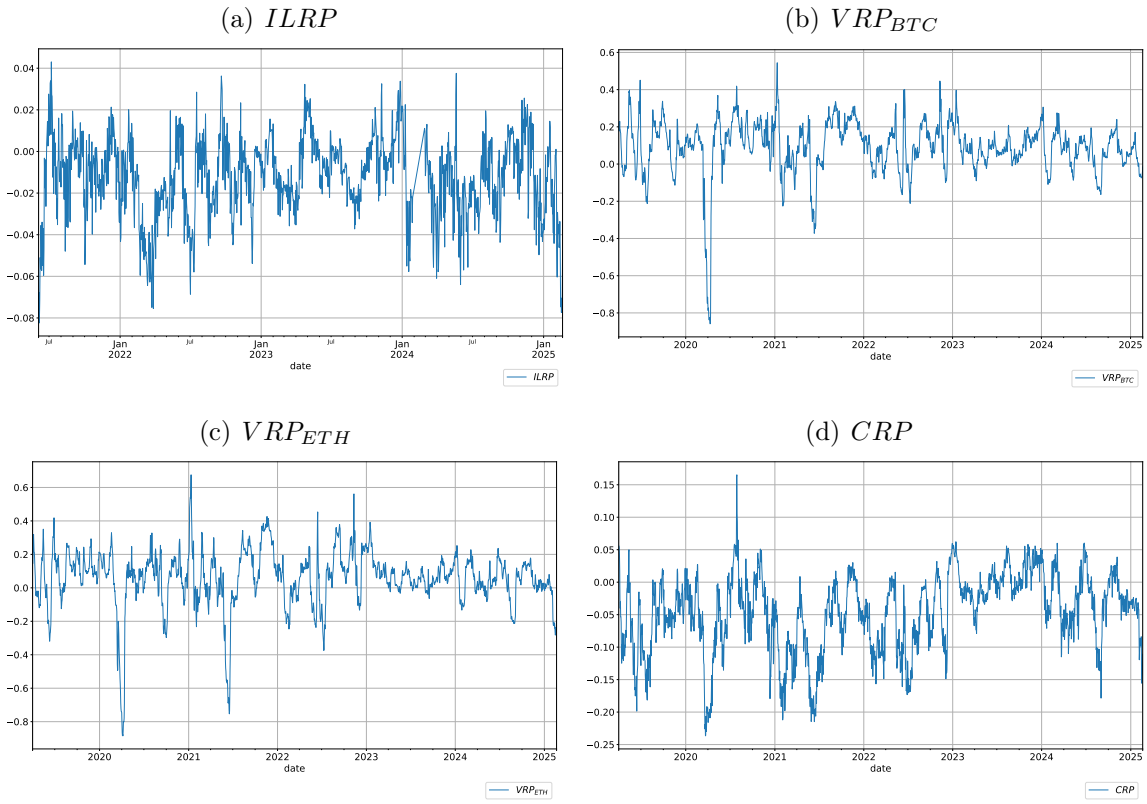

Figure 6.2: Ex-Ante Risk Premia. The figure shows the ex-ante risk premia for the impermanent loss ($ILRP$), the VRPs for BTC and ETH, and the CRP. The implied quantities are calculated by the equations in Section 5.1. The realized quantities are calculated as discussed in Section 5.2. The data is sampled daily, and the sample period is from 2021-05 to 2025-02.

	(1)	(2)	(3)	(4)	(5)
Intercept	-0.007	-0.015	-0.017	-0.006	-0.007
	0.000	0.000	0.000	0.000	0.000
VRP_{BTC}	–	–	0.048	–	-0.006
	–	–	0.000	–	0.323
VRP_{ETH}	–	0.037	–	-0.009	–
	–	0.000	–	0.070	–
CRP	0.162	–	–	0.176	0.168
	0.000	–	–	0.000	0.000
R^2	19.912	6.784	6.514	20.059	19.903

Table 6.5: Drivers of the Implied Impermanent Loss Risk Premia. The table reports the results of slope, p-value, and the R^2 of the described regression procedure. We report the specifications with one ((1)-(3)) and two ((4)-(5)) independent variables. The risk premia are calculated following Section 5.3. The data is sampled daily from 2021-10 to 2025-02.

negative return of -0.16 (-0.18) and -0.14 (-0.17). Figure 6.3 displays the measures around these days. Comparing Panel (a) and Panel (e) it turns out that the IIL steeply increased on the second event while it did react delayed on the first event. For the IV_{BTC} a rapid increase is visible in both Panel (b) and Panel (f) and the same is true and even more pronounced for IV_{ETH}, as depicted in Panel (c) and Panel (g). Lastly, as shown in Panel (d) and Panel (h), the strong negative return in both tokens increases the IC.

We then investigate the dynamics of our measures for days when one token experiences a large decline while the other does not. For example, BTC experienced a large drawdown on the 5th of December 2021 with a negative return of -0.08 while ETH only experienced a small drawdown of -0.02. Adding to that, on the 16th of September 2022, ETH experienced a large negative drawdown of -0.10 while BTC only experienced -0.02. Panel (a)-(d) of Figure 6.4 displays the IIL and the IC on these two days for BTC and ETH: As visible, the IIL increases before and during the drawdown. The effect is stronger for the second event (Panel (b)). In line with the intuition, the IC drops sharply in both events (Panel (c) and Panel (d)).

Whenever there is a shift in the relative price between the assets, the LP incurs impermanent loss, even if both underlying assets generate positive returns. And hence. even though both underlying assets may experience positive returns, this does not shield the LP from impermanent loss. We examine a day when both underlying assets rose in price, with one increasing more than the other. For example, on the 19th of July 2022, BTC realized a gain of almost 0.04 and ETH of over 0.11. The results are presented in Panel (e)-(h) of Figure 6.4: As shown, while IV_{ETH} increases (leading to higher IIL), IC also rises (reducing IIL), but this still results in a net increase in IIL around the event.

6.4 Risk Premia

Next, we analyze the behavior of the risk premia for some of the previously mentioned events. The dynamics of the risk premia at the worst day for BTC and ETH are displayed in Figure 6.5 (Panel (a) - Panel(d)). Both tokens decreased, which led to a smaller change in the RIL (as compared to the IIL) and hence the $ILRP$ overall decreased to -0.04. As both tokens experience a drawdown, the VRPs first increase due to the increase in IV and decrease after the drawdown (due to the increase in RV). The CRP was negative before the event, indicating that the IC was lower than the RC. However, as the drawdown approached, IC increased.

For a more comprehensive view of risk premia dynamics, Figure 6.5 (Panel (e) -

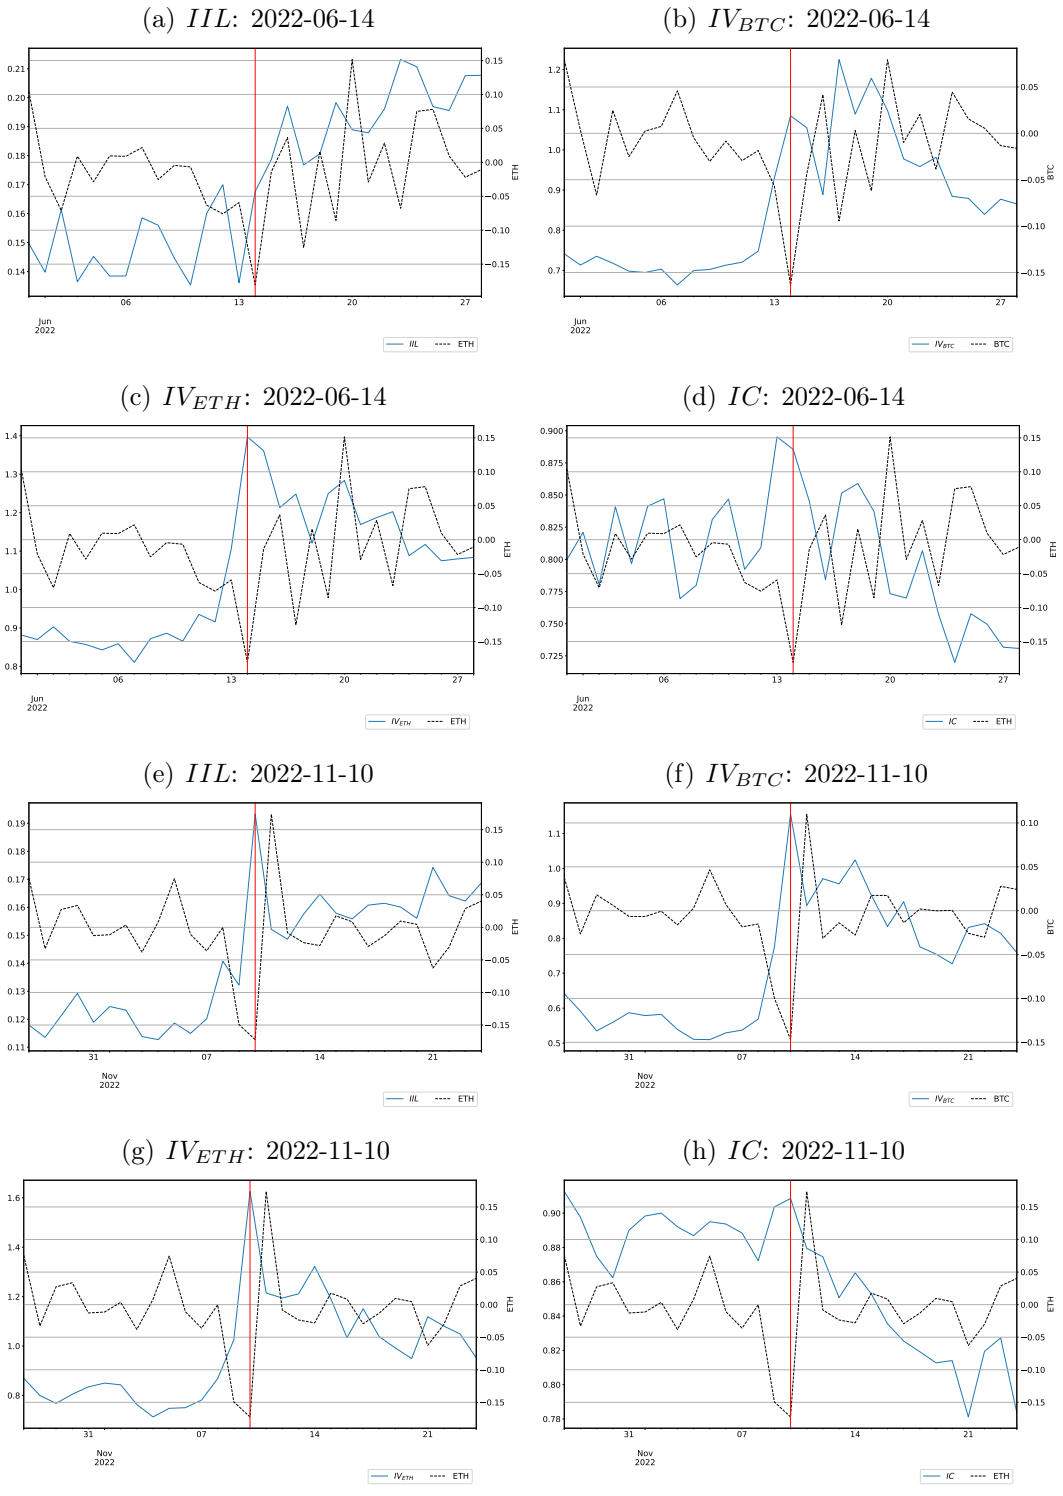

Figure 6.3: Implied Variables and Worst Days of both Underlyings. The figure displays the implied variables around the worst days of the underlying in our sample, which was 2022-06-14 (Panel (a)-(d)) and 2022-11-10 (Panel (e)-(h)) where BTC experienced a drawdown of -0.16 and -0.14, and ETH -0.18 and -0.17 respectively. The return of the respective underlying is plotted on the right y-axis. The implied quantities are calculated by the equations in Section 5.1.

Panel(h)) illustrates the behavior of risk premia on September 16, 2022, when ETH underwent a significant negative drawdown of -0.10, whereas BTC experienced only -0.02. As visible, the $ILRP$ does not decrease but rather increases, meaning that the RIL increased stronger than the IIL. Since ETH experienced a large drawdown, VRP_{ETH} decreases, falling from a high level. The CRP decreases, meaning that RC increased stronger than IC.

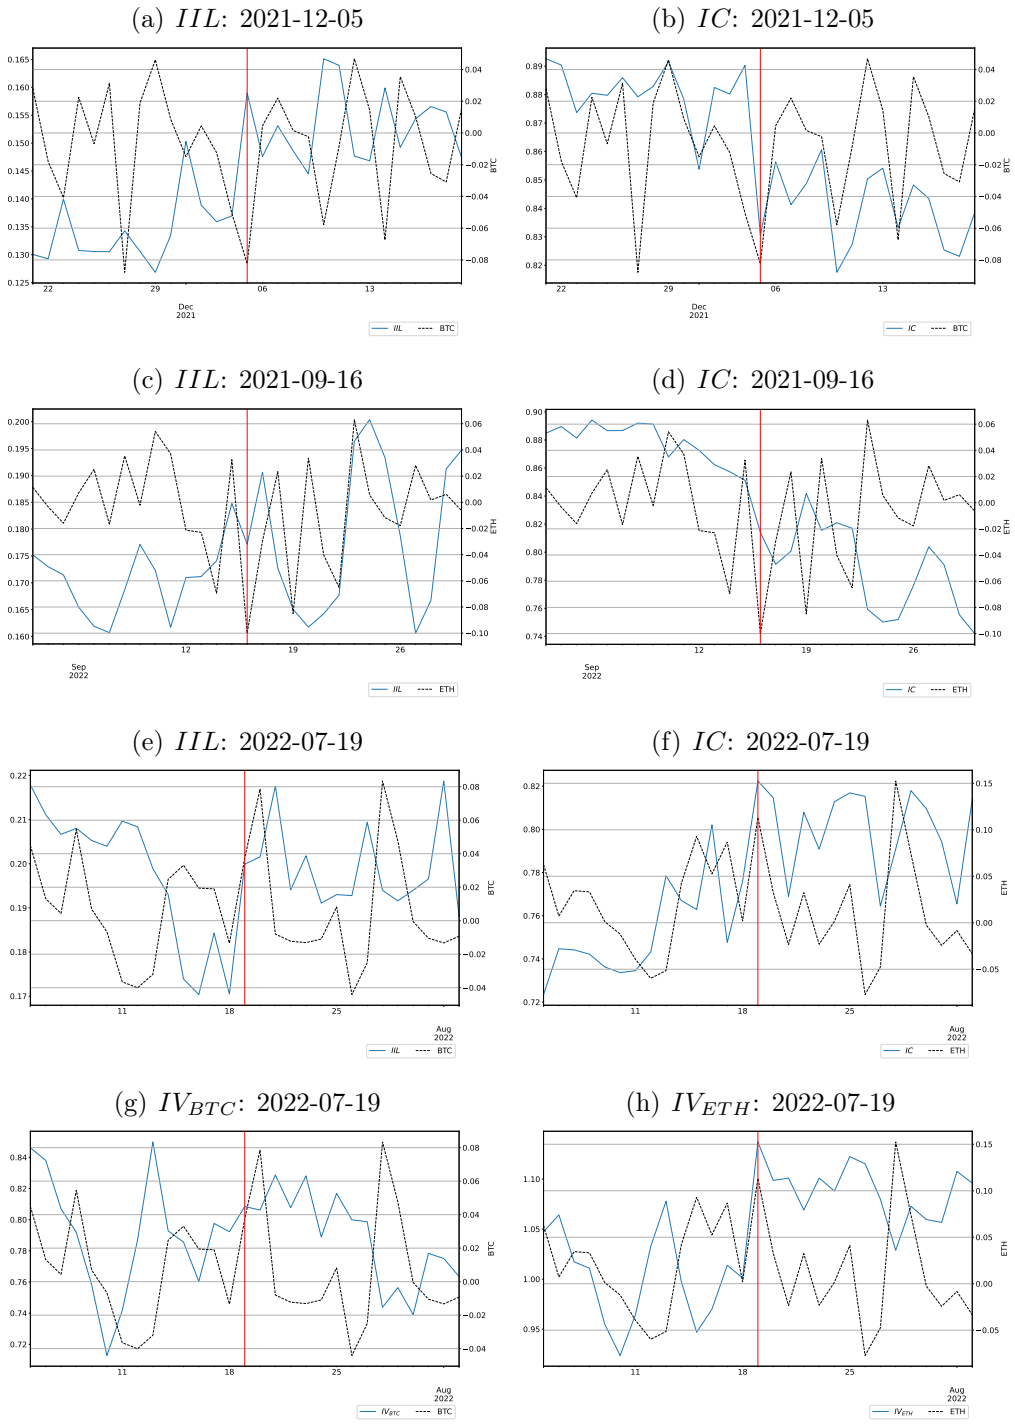

Figure 6.4: Implied Variables and Bad and Good Days for each Underlying . Panel (a) - Panel (d) of the figure displays the IIL and IC around the worst days of each underlying while the other underlying did not experience its worst day, which was 2021-12-05 for BTC (-0.08) and ETH (-0.02) and 2021-09-16 for ETH (-0.10) and BTC (-0.02). Panel (e) - Panel (h) displays the implied variables around one of the best days of ETH, which was 2022-07-19 where ETH experienced an increase of 0.11, and BTC 0.04 respectively. The return of the respective underlying is plotted on the right y-axis. The implied quantities are calculated by the equations in Section 5.1.

(a) $ILRP$: 2022-06-14

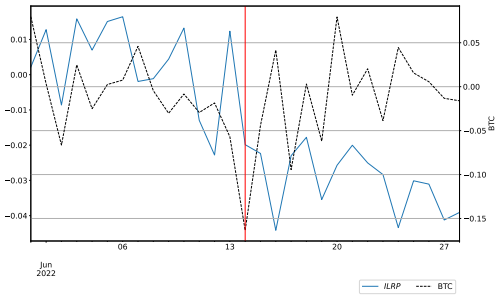

(b) VRP_{BTC}: 2022-06-14

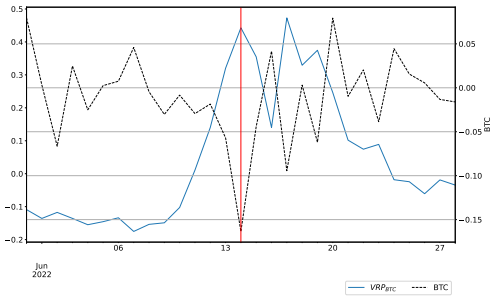

(c) VRP_{ETH}: 2022-06-14

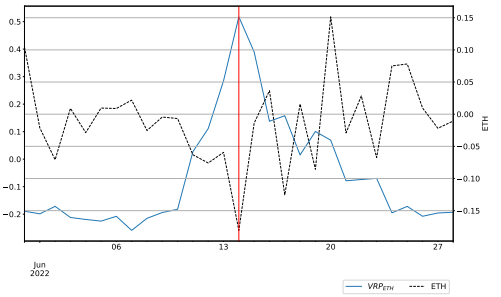

(d) CRP: 2022-06-14

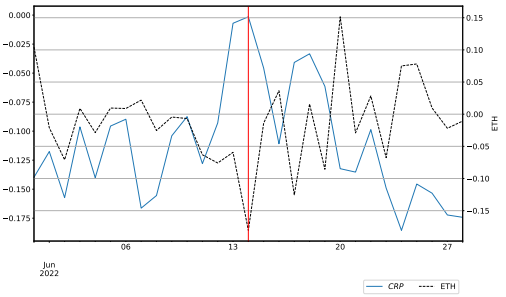

(e) $ILRP$: 2022-09-16

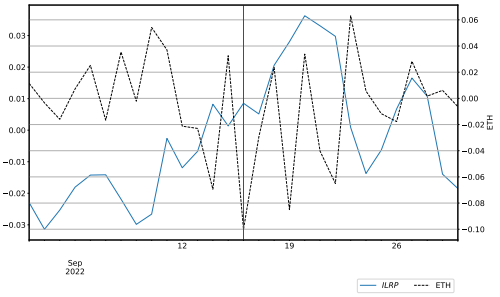

(f) VRP_{BTC}: 2022-09-16

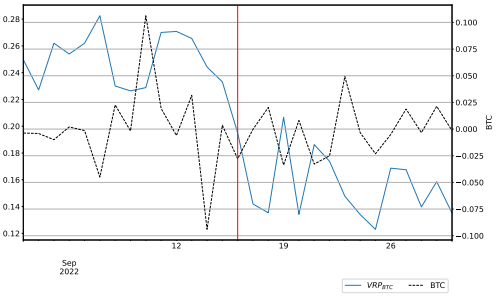

(g) VRP_{ETH}: 2022-09-16

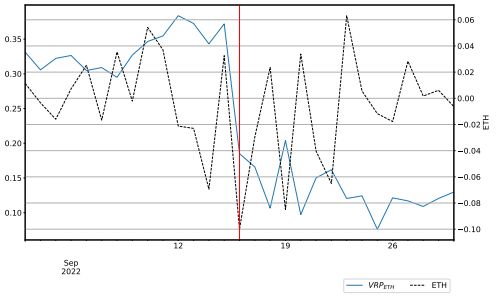

(h) CRP: 2022-09-16

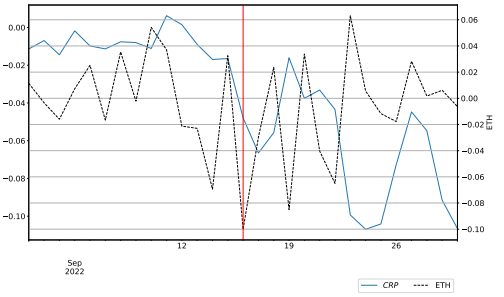

Figure 6.5: Risk Premia and Good Days of both Underlyings. The figure displays the risk premia around the worst day of the underlying in our sample (Panel (a) - Panel (d)), which was 2022-06-14 where BTC experienced a drawdown of -0.16, and ETH -0.18 respectively. Panel (e) - Panel (h) displays the risk premia around the worst day of ETH, while BTC did not experience its worst day, which was 2021-09-16 for ETH (-0.10) and BTC (-0.02). The return of the respective underlying is plotted on the right y-axis. The risk premia are calculated by the equations in Section 5.3.

6.5 Implied Impermanent Loss and Liquidity Pool Size

Lehar and Parlour (2023) demonstrate that equilibrium pool size balances fee revenue against impermanent loss. This equilibrium is reached by adjusting pool size rather than prices, as larger pools have a smaller price impact. High impermanent loss risk results in smaller equilibrium pool sizes due to the greater price impact on small pools. We test this hypothesis by examining the contemporaneous correlation between the implied (Panel (a)) and realized (Panel (b)) measures and the aggregated pool size measured as TVL. As shown in Figure 6.6 the correlation is negative (in the first half of the sample and the last year of the sample) for IIL and both IVs since an increase in risk translates into a lower aggregated pool size. Consistently, for IC the correlation is positive since a decrease in risk (stronger comovement) is accompanied by a larger pool size.

(a) Implied Measures and Aggregated TVL

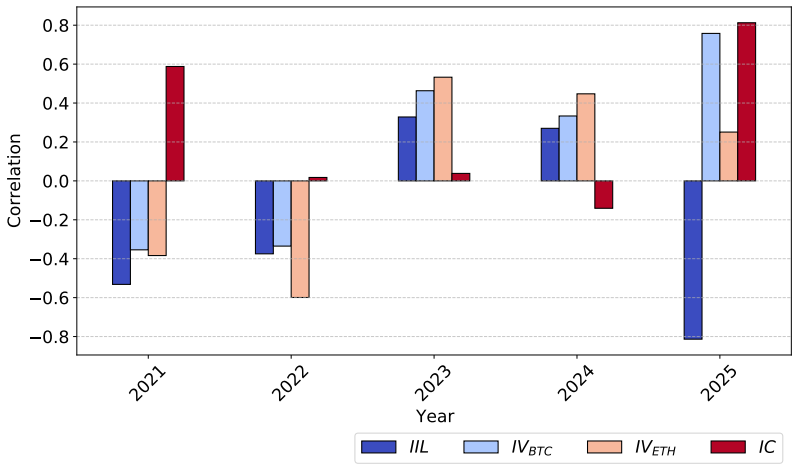

Figure 6.6: Contemporaneous Correlations – Implied Measures and Aggregated Pool Size (TVL). The plots report the correlation coefficients calculated for each year separately between the implied quantities, and the TVL which is calculated as the aggregate TVL of all pools. The implied quantities are calculated by the equations in Section 5.1. The sample period is from 2021-05 to 2025-02.

6.6 Uniswap - Risks and Rewards

To obtain an overview, some summary statistics are presented first. The realized impermanent loss given by equation (5.3) (but without taking the square root) represented as a time-series average for a cross-section of pools, is visually presented in Figure 6.7. The annualized realized impermanent loss across the pools is on average at around -8.89%,

with extreme values up to −40%.

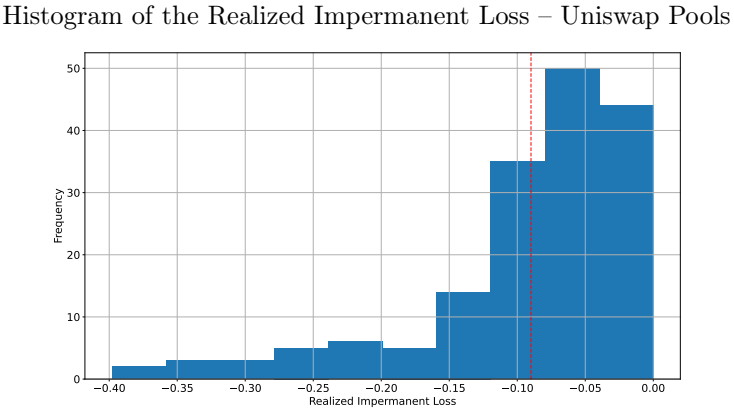

Figure 6.7: Cross-Section: Realized Impermanent Loss – Uniswap Pools. The figure shows the histogram of the (average) realized impermanent loss (annualized) in variance terms of the Uniswap pools. The dotted line represents the average impermanent loss. The data is sampled daily, and the sample period ranges from 2021-05 to 2025-02. The data is winzorized at the 5% quantile.

Next, we present the average APRs calculated by equation (5.4) for the cross-section of pools. Figure 6.8 showcases the histogram of the time-series average gross APRs in Panel (a) and net APRs in Panel (b); net APR is calculated as gross APR minus individual impermanent loss. The average gross APR is approximately 15%, however, due to substantial impermanent loss, the average net APR decreases to only 3%, and there are instances of notably negative net APRs, extending to as low as −37%.

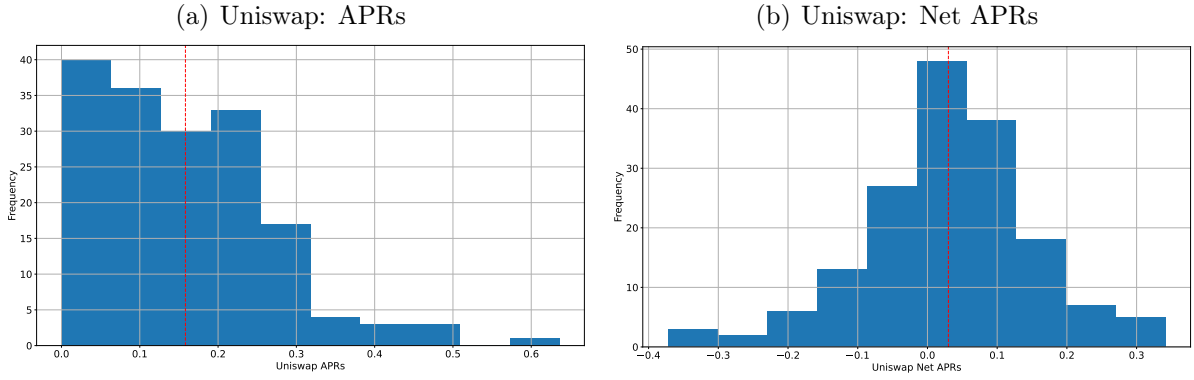

Figure 6.8: Histogram: Uniswap Pool APRs. The figure shows the histogram of the (average) APRs of the Uniswap pools. In Panel (a) the gross APRs are displayed, and in Panel (b), the net APRs. The dotted line represents the average. The data is sampled daily, and the sample period is from 2021-05 to 2025-02. The data is winzorized at the 5% quantile.

In Figure 6.9, the gross and net APRs for the Uniswap pools, calculated as the cross-

sectional average across individual pools, are depicted over time. The graph illustrates the volatility of both gross and net APRs. Notably, the gross APR has exhibited a consistent decline, whereas the net APR has remained relatively stable, hovering around 5-10% since 2023.

(a) Uniswap: APRs

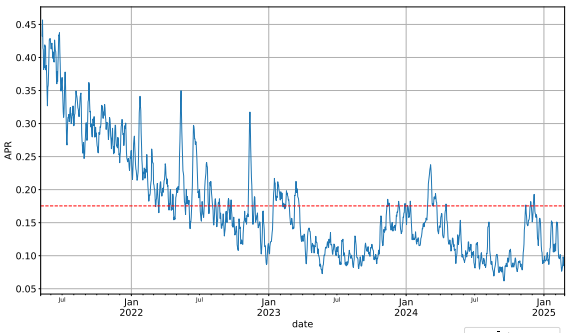

(b) Uniswap: Net APRs

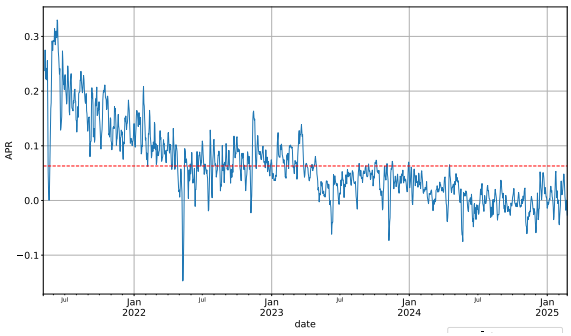

Figure 6.9: Uniswap Pool APRs. The figure shows the cross-sectional average of the APRs of the Uniswap pools ($\bar{\text{U}}$niswap). Panel (a) displays the gross APRs, and Panel (b) displays the net APRs. The dotted line represents the average. The data is sampled daily, and the sample period is from 2021-05 to 2025-02. The data is winzorized at the 5% quantile. In the plots, the five-day moving average is depicted.

6.7 APR Predictability

For the next empirical analysis, we leverage the information content of the IIL and its drivers as forward-looking measures to illustrate their predictive relationship with future APRs by estimating the following time-series regression:

$$APR_{t\to t+\tau} = \alpha + \beta_0 IIL\,(t) + \beta_1 IIL\,(t-1) + \varepsilon_t \tag{6.1}$$

where $APR_{t\to t+\tau}$ denotes the cumulative aggregated APR for all pools where at least one of the tokens of the token pair is either BTC or ETH from the period from t to $t+\tau$. The values for time horizon τ are set to be $1, 2, 3, 5, 7, 10, 14$, and 30 days. Standard errors are corrected for autocorrelation due to overlapping APR observations, following the method outlined in Newey and West (1987).

6.7.1 In-Sample APR Predictability

We first perform an in-sample analysis using equation (6.1). The results are displayed in Figure 6.10 for the IIL (Panel (a) and Panel (b)) which displays a positive significant

sign for both coefficients over the full predictive time horizon of 30 days. We repeat the regression analysis for the drivers of the IIL in Panel (c) - Panel (h). For the IVs (Panel (c) - Panel (f)), the betas are positive. For IV_{ETH}, the betas of the lagged values become significant for a time horizon of ten days and larger. For IC (Panel (g)) the betas are as expected negative but not significant. Increased correlation (less risk for the LP) lowers the future return. Overall, a larger level of risk corresponds to a larger future return. Again, realized quantities barely predict future APRs, as discussed in the Robustness Section 7.

6.7.2 Out-of-Sample APR Predictability

We document the out-of-sample (OOS) performance of the implied measures for the APRs. As in most studies, the forecasting performance of a specific model s is compared to that of a model based on the historical mean of the respective return ($s = 0$). Therefore, the OOS R^2 is calculated as

$$R^2_{s,\tau} = 1 - \frac{MSE_{s,\tau}}{MSE_{0,\tau}}, \tag{6.2}$$

where $MSE_{s,\tau} = \frac{1}{N}\sum^{N} e^2_{s,\tau}$ denotes the mean-squared error of model s computed from the prediction errors $e_{s,\tau}$ for time horizon τ. A particular model, s, outperforms the benchmark model ($s = 0$), which is based on the average historical return if the OOS $R^2_{s,\tau}$ is significantly positive. OOS predictions are derived from a rolling window estimation of the predictive in-sample regression (6.1). This includes a time-specific intercept but excludes lagged values of IIL. The estimated coefficient $\beta_{ILL,t}$ and the current value of IIL_t are then used to generate the OOS return forecast $APR_{t\to t+\tau}$. Note that at date t, one uses only observations from the past to avoid any look-ahead bias.

Figure 6.11 presents the OOS R^2 for the predictions related to the APRs based on an initial 180-day estimation period where each day we predict the APR over a time horizon of 1 to 30 days. As visible, the OOS R^2s are positive but decreasing for longer prediction time horizons: For example, the OOS R^2s for the IIL peaks at 5% at the 1-day predictive time horizon. The IVs for BTC and ETH perform better than the IIL with OOS R^2s ranging from 18%-6%, while the IC performs similar to the IIL with a maximum R^2 of 4%. Similar to previous findings, realized quantities predict APRs less accurately OOS. In addition, we use the results of the linear regression model as a benchmark and implement various machine-learning models for comparison. The findings are discussed in the Robustness Section 7.

The OOS results using risk premia as a predictor are presented in Figure 6.12. Qual-

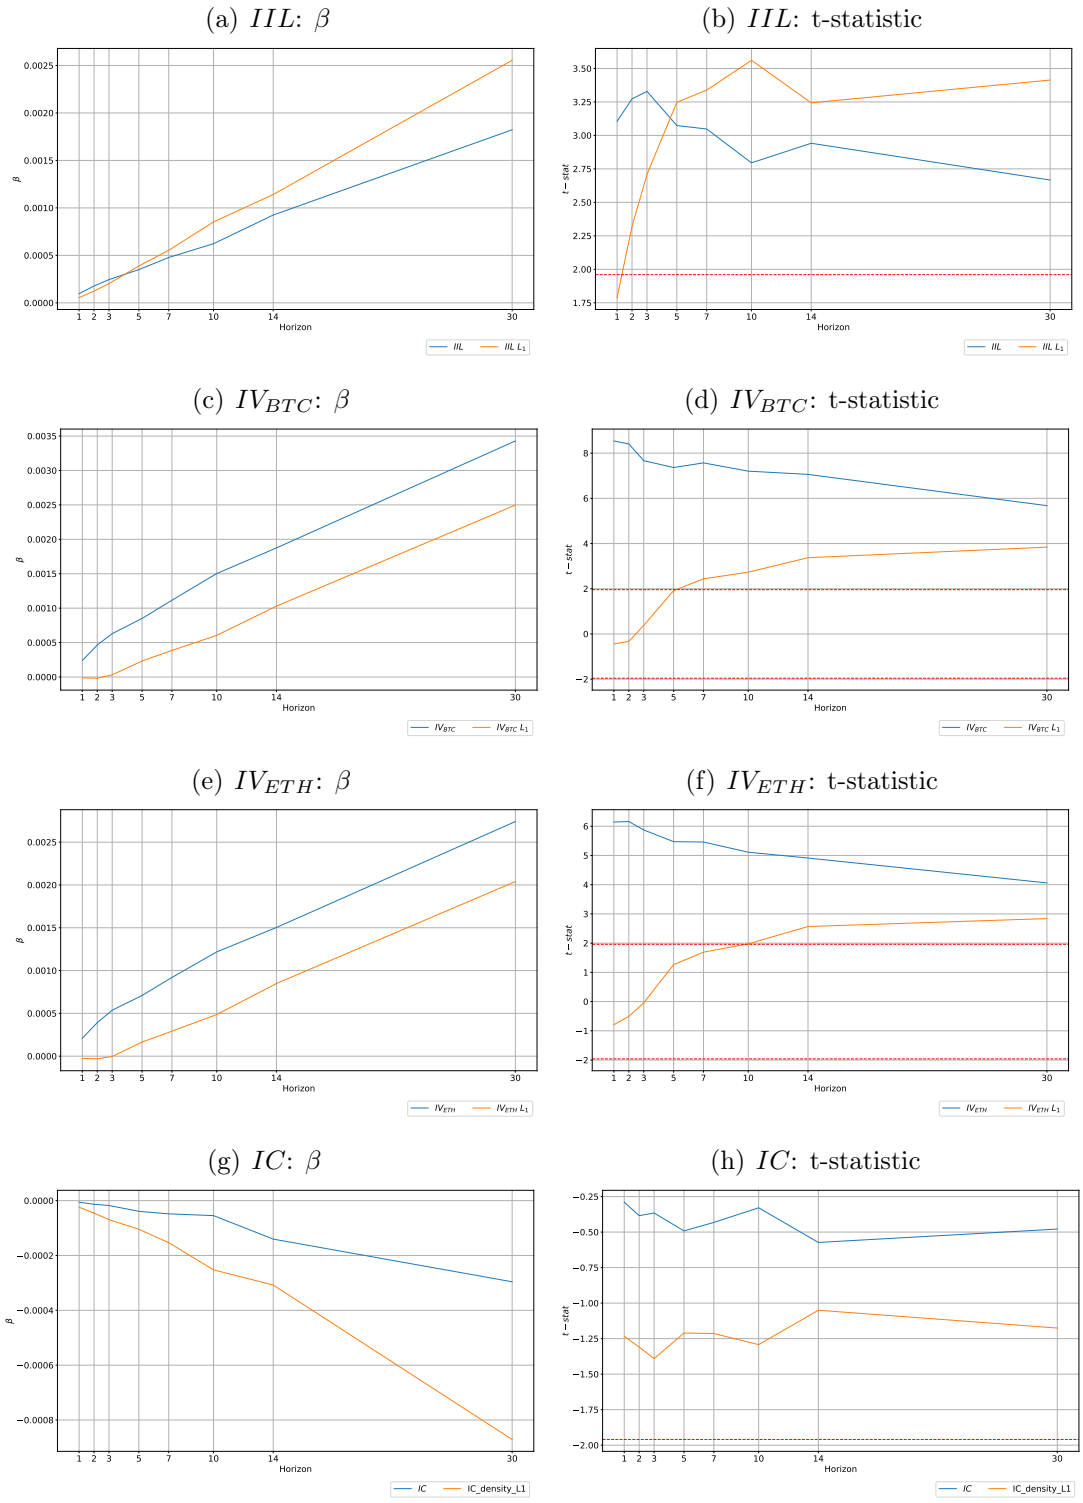

Figure 6.10: Predictive – Future APRs – Implied Measures. The plots report the β coefficients and their t-statistic for the time-series regression described by equation (6.1). Standard errors are computed with Newey and West (1987). As explanatory variables, IIL and its drivers IV_{BTC}, IV_{ETH}, IC, and their lagged values are considered. The implied quantities are calculated by the equations in Section 5.1. The sample period is from 2021-05 to 2025-02.

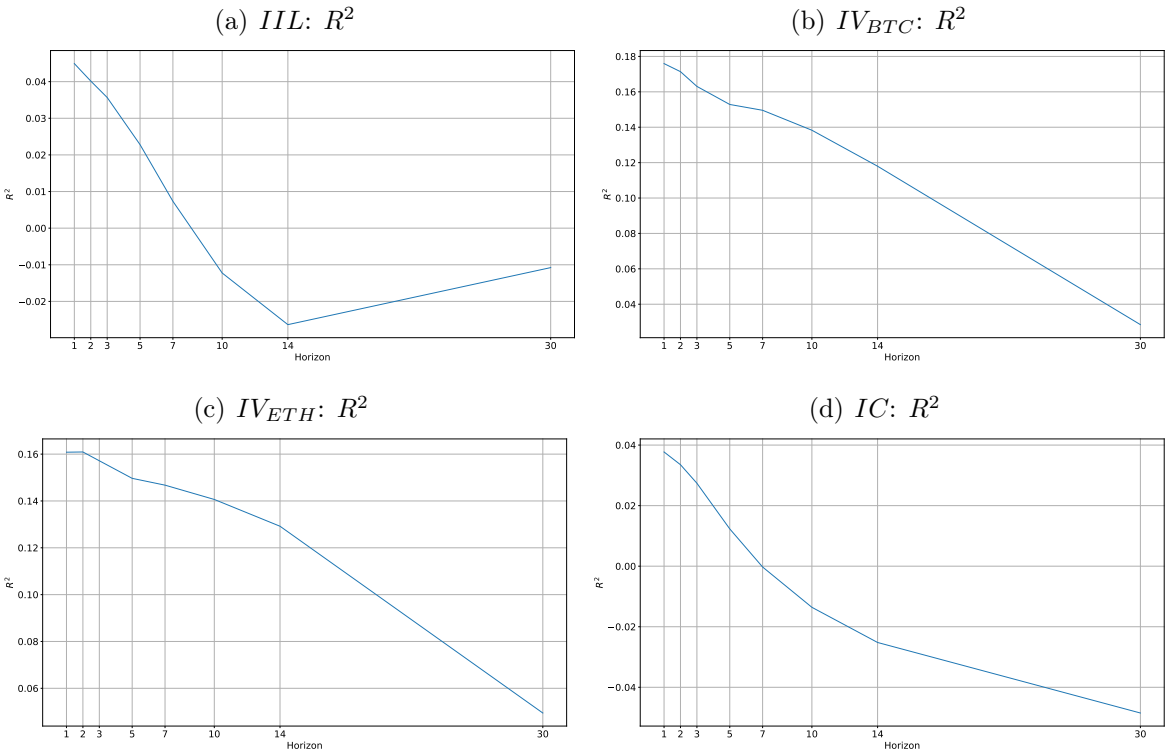

Figure 6.11: Out-of-Sample Predictions – Implied Measures. The plots report the OOS R^2 calculated by equation (6.2) for predictions based on a 180-day rolling estimation window. The implied quantities are calculated by the equations in Section 5.1. The sample period is from 2021-05 to 2025-02.

itatively, these results align with those obtained using only the implied variables, though the OOS R^2 has improved.

6.8 Cross-Sectional Implications for Liquidity Provision

To establish a risk-return relationship between the impermanent loss and the returns of the liquidity provision, we perform a two-stage Fama and MacBeth (1973) regression, where changes in the IIL and its drivers described in Section 5.1 are considered risk factors. As test assets, the pools' APRs are utilized, where their betas are estimated over a one-year rolling estimation window. Hence in the first stage, the APRs are regressed on a constant and the changes in the implied variables. In the second stage, the expected APRs are regressed on these betas from the first stage. Expected APRs are proxied as the realized APRs over the next $1, 2, 3, 5, 7, 10,$ and 14 days. The coefficients from the second stage regression (γ) represent the average risk premia of the risk factors.

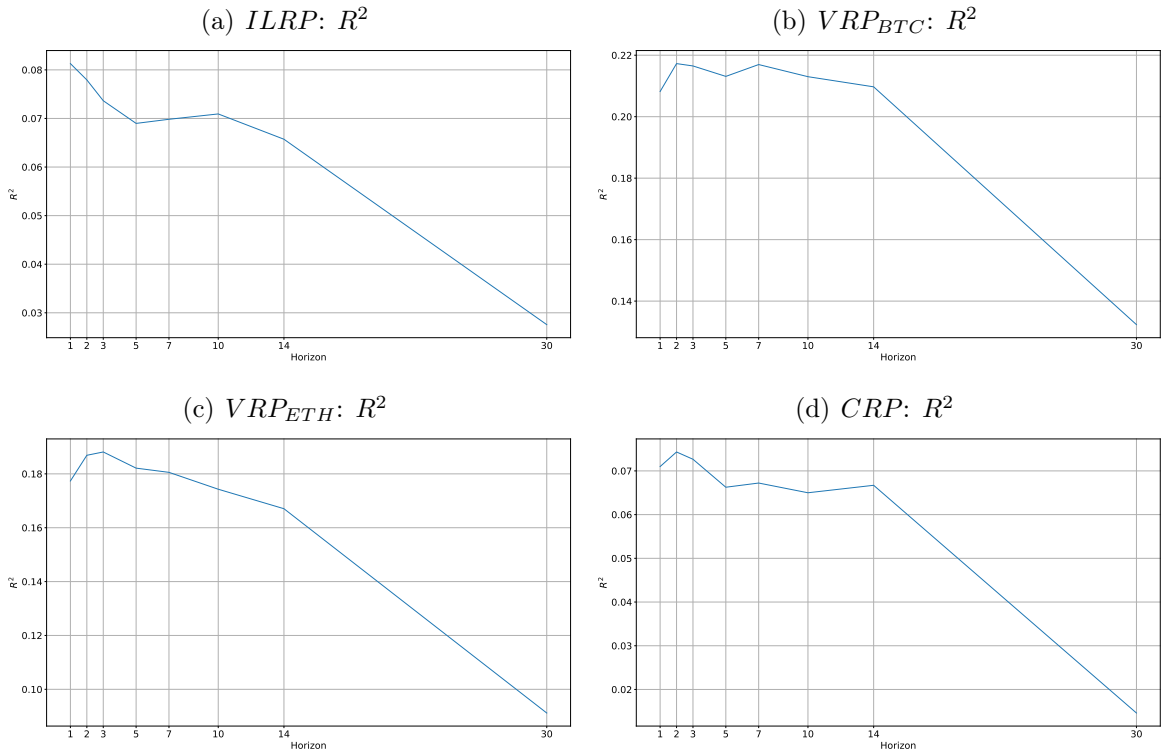

Figure 6.12: Out-of-Sample Predictions – Risk Premia. The plots report the OOS R^2 calculated by equation (6.2) for predictions based on a 180-day rolling estimation window. The risk premia are calculated following Section 5.3. The sample period is from 2021-05 to 2025-02.

6.8.1 Univariate Analysis

As visible from Figure 6.13 Panel (a) and Panel (b), the price of risk for the IIL is positive and displays significance for up to 14 days. The results for the drivers are displayed in Figure 6.13: The IV_{BTC} is positively priced and significant for up to 10 days (see Panel (c), and Panel (d)). The same holds for IV_{ETH}, where the significance is given for up to 14 days (see Panel (e), and Panel (f)). For IC, the regression coefficient is negative and statistically significant up to 7 days (Panel (g) and Panel (h)). In liquidity provision, LPs prefer tokens with high positive correlations to minimize impermanent loss. As a result, the price of risk, with IC as a factor, is negative. The cross-sectional findings for the RIL and RVs demonstrate qualitative similarity, whereas the RC lacks significance, as detailed in the Robustness Section 7.

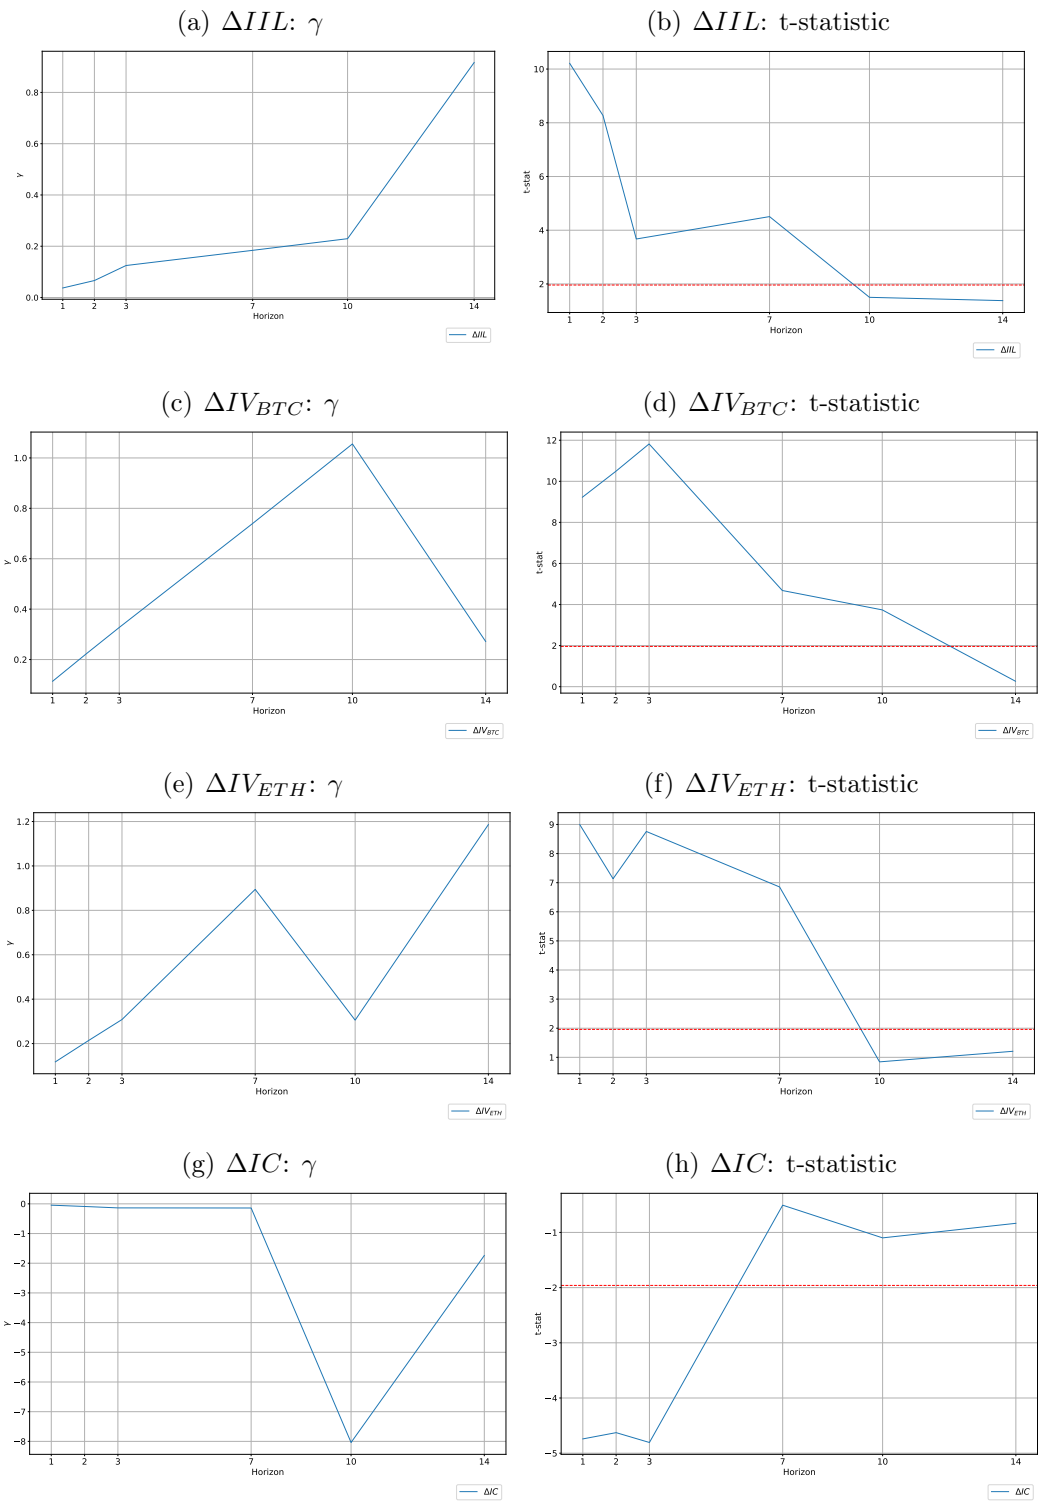

Figure 6.13: Cross-Section – Implied Measures. The table reports the γ coefficient and its t-statistic for the Fama and MacBeth (1973) two-stage cross-sectional regression. As risk factors, the drivers of IIL are considered, namely the changes in IV_{BTC}, IV_{ETH}, and IC. The implied quantities are calculated by the equations in Section 5.1. In each regression we control for the changes in the individual realized impermanent loss. The sample period is from 2021-05 to 2025-02.

6.8.2 Multivariate Analysis

We repeat the procedure for a multivariate setting where we simultaneously estimate the betas for more than one factor. Figure 6.14 displays the γ and its t-statistic over the predictive time horizons. In line with the univariate regressions, the volatilities (correlation) carry a positive (negative) price of risk.

(a) IIL Drivers: γ (b) IIL Drivers: t-statistic

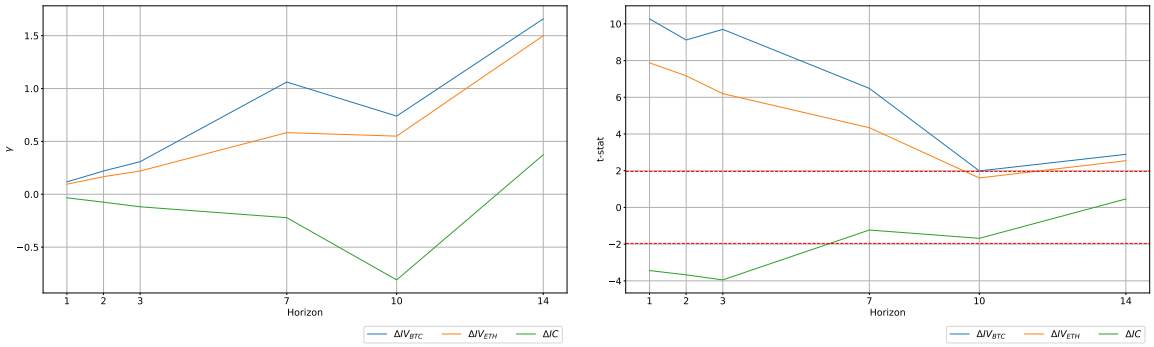

Figure 6.14: Cross-Section – Multivariate – Drivers of the IIL. This figure reports the γ coefficient and its t-statistic for the multivariate Fama and MacBeth (1973) two-stage cross-sectional regression. As risk factors, the drivers of IIL are considered, namely changes in IV_{BTC}, IV_{ETH}, and IC. The implied quantities are calculated by the equations in Section 5.1. In each regression we control for the changes in the individual realized impermanent loss. The sample period is from 2021-05 to 2025-02.

7 Robustness

To verify the robustness of the analysis to various specifications, a series of tests are carried out and reported in this section, Appendix D, and Internet Appendix II.

7.1 Non-Overlapping Predictions

To avoid autocorrelation caused by overlapping observations, the APRs are sampled in a non-overlapping fashion. In Figure D.3, the average slope and t-statistic for the APR predictability for each maturity are displayed. As visible, the signs of the coefficients align with the main results even though the coefficients are less significant: For example, the IIL is only significant up to 2 days. The results for IV_{BTC}, and IV_{ETH} are similar with significance given up to $5-7$ days.

7.2 OOS Predictions using Machine Learning Models

Taking the results of this linear regression (LR) model as a benchmark, we implement various machine learning models: linear regression with elastic net regularisation and Huber loss (LRENH), gradient boosted regression tree (GBRT), random forest regression (RFR), long short-term memory neural network (LSTM), and 1D convolutional neural network (CNN1D). The model settings are the same as described in Section 6.7.2, except for the two neural networks LSTM and CNN1D, we include the lagged values of the independent variables. Furthermore, besides testing the univariate models as the benchmark model, we also test multivariate models for all implied and realized variables. For training the models, we perform a grid search for the model hyperparameters, and we set ten percent of the data for model validations. The results are displayed in Figure D.4 and D.5 indicate that some machine learning models outperform linear regression in predictive accuracy measured by OOS R^2 and that implied variables exhibit greater predictive power compared to their realized counterparts.

7.3 Misspecification and Tradable Factor Risk Premia

The methodology proposed by Fama and MacBeth (1973) has been demonstrated to yield unreliable results, particularly when an asset pricing model may be misspecified or a candidate risk factor is only weakly correlated with returns, which are studied in Gospodinov et al. (2014), Kan et al. (2013), and Quaini et al. (2023).

We follow Quaini et al. (2023) and estimate the i) misspecification-robust stochastic discount factor (SDF) coefficients following Gospodinov et al. (2014) (GKR estimator), and ii) tradable factor risk premium following Kan et al. (2013) (KRS–FRP estimator) and Quaini et al. (2023) (TFRP and O–TFRP estimators), which is defined by the negative factor covariance with the projection of any SDF on the space of asset returns. The results are presented in Figure D.6. The left panels (a), (c), (e), and (g) display the misspecification-robust GKR estimator. The right panels (b), (d), (f), and (h) display the tradable factor risk premia calculated by the KRS–FRP, the TFRP, and the O–TFRP estimators. As visible, the coefficient and the risk premia are positive and significant for IIL, IV_{BTC}, and IV_{ETH}. For IC they are as expected negative but not significant.

7.4 Realized Quantities

Similar to the empirical analyses in Sections 6.5, 6.7, and 6.8 for the implied quantities, we repeat these analyses for the realized quantities from Section 5.2.

7.4.1 Realized Impermanent Loss and Liquditiy Pool Size

Figure II.1 displays the correlation coefficient of the realized measures and the aggregated TVL. The correlation behaves similarly as for implied variables (Figure 6.6), with the difference that the RC often displays a large negative value in the early sample.

7.4.2 In-Sample APR Predictability

For the APR in-sample prediction, we repeat the regression analysis described by equation (6.1) using the realized quantities RIL, RV_{BTC}, RV_{ETH}, and RC defined in Section 5.2. The results are displayed in Figure II.2. The RIL predicts the future APRs for a time horizon of three days. None of the other realized quantities display significance.

7.4.3 Out-of-Sample APR Predictability

For the APR OOS prediction, we repeat the procedures described in Section 6.7.2 using the realized quantities RIL, RV_{BTC}, RV_{ETH}, and RC defined in Section 5.2. The OOS R^2 is calculated by equation (6.2). The results are displayed in Figure II.3. The RVs predicts the future APRs OOS only for a time horizon of one day.

7.4.4 Cross-Sectional Prediction – Realized Variables

We repeat the cross-sectional tests described in Section 6.8 for realized variables. The results are displayed in Figure II.4 and are qualitatively similar to those for the implied variables, except for RC which does not show any significance. In addition, we repeat the ability to price the cross-section of pools using the changes in the individual daily realized impermanent loss (RIL_{ind}) of the respective pool (and not always the RIL for BTC-ETH). As displayed in Figure II.5 Panel (b) the associated t-statistic of the second-stage regression coefficient is only significant for the next two days.

8 Conclusion

The IIL reflects the risk of a disparity in relative token prices as implied by option prices. This article quantifies this risk, investigates its fundamental economic drivers, and delves deeper into its predictive capabilities.

We assess the impermanent loss from a risk-neutral standpoint. Impermanent loss equates to one-eighth of the volatility of the relative price and can be computed by evaluating the log contract proposed by Carr and Madan (1998) using a portfolio of European options on the relative price. To tackle the primary challenges—namely, the lack of an options market for the relative price and the relative price not being a martingale—we take a two-step approach. First, we compute a joint density, minimizing the Hansen and Jagannathan (1991) bounds, from existing options on separate tokens. Second, we employ a change of numéraire. Our methodology integrates existing replication approaches that assume one of the tokens is a stablecoin, thereby offering methodological contributions as well.

We present empirical findings indicating that the IIL and its drivers are related to the pool size and subsequent pool returns. The predictability of IIL stems from its three primary components: implied individual token variances and their implied correlation. These components individually account for the variation in pool returns across different tokens.

References

Carol Alexander and Arben Imeraj. The bitcoin vix and its variance risk premium. *The Journal of Alternative Investments*, 23(4):84–109, 2021. ISSN 1520-3255.

Andrew Ang and Joseph Chen. Asymmetric correlations of equity portfolios. *Journal of Financial Economics*, 63(3):443–494, 2002. ISSN 0304-405X.

Andrew Ang, Robert J. Hodrick, Yuhang Xing, and Xiaoyan Zhang. The cross-section of volatility and expected returns. *The Journal of Finance*, 61(1):259–299, 2006. doi: https://doi.org/10.1111/j.1540-6261.2006.00836.x. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.2006.00836.x`.

Patrick Augustin, Roy Chen-Zhang, and Donghwa Shin. Reaching for yield in decentralized financial markets. In *Proceedings of the EUROFIDAI-ESSEC Paris December Finance Meeting 2023*, Paris, France, 2023. doi: https://dx.doi.org/10.2139/ssrn.4063228.

Gurdip Bakshi, Nikunj Kapadia, and Dilip Madan. Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, 16(1):101–143, 2015. ISSN 0893-9454.

Tim Bollerslev, George Tauchen, and Hao Zhou. Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11):4463–4492, 2009. ISSN 0893-9454.

Tim Bollerslev, Viktor Todorov, and Lai Xu. Tail risk premia and return predictability. *Journal of Financial Economics*, 118(1):113–134, 2015. URL `https://EconPapers.repec.org/RePEc:eee:jfinec:v:118:y:2015:i:1:p:113-134`.

Oleg Bondarenko and Carole Bernard. Option-implied dependence and correlation risk premium. *Journal of Financial and Quantitative Analysis*, 59(7):3139–3189, 2024. doi: 10.1017/S0022109023000960.

Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–651, 1978. ISSN 0021-9398.

Andrea Buraschi, Robert Kosowski, and Fabio Trojani. When there is no place to hide: Correlation risk and the cross-section of hedge fund returns. *The Review of Financial Studies*, 27(2):581–616, 2014. ISSN 0893-9454.

Adrian Buss, Lorenzo Schönleber, and Grigory Vilkov. Expected correlation and future market returns. Working Paper, FSFM, 2017.

Agostino Capponi and Ruizhe Jia. Liquidity provision on blockchain-based decentralized exchanges. *SSRN Manuscript*, 2024. doi: https://dx.doi.org/10.2139/ssrn.3805095.

Peter Carr and Roger Lee. Robust replication of volatility derivatives. In *MFA 2008 Annual Meeting*, 2008. URL `https://math.uchicago.edu/~rogerlee/rrvd.pdf`.

Peter Carr and Dilip Madan. *Volatility: New Estimation Techniques for Pricing Derivatives*, chapter 29 Towards a theory of volatility trading. Risk Books, 1998. ISBN 9781899332410.

Peter Carr and Liuren Wu. Variance risk premiums. *The Review of Financial Studies*, 22(3):1311–1341, 2009. ISSN 0893-9454.

Alvaro Cartea, Fayçal Drissi, and Marcello Monga. Decentralised finance and automated market making: Predictable loss and optimal liquidity provision. *arXiv Manuscript*, 2024. doi: https://doi.org/10.48550/arXiv.2309.08431.

Peter Christoffersen, Kris Jacobs, and Bo Young Chang. Forecasting with option implied information. *Handbook of Economic Forecasting*, 2(A):581–656, 2013. ISSN 1574-0706.

Joseph Clark. The replicating portfolio of a constant product market. *SSRN Manuscript*, 2020. doi: https://dx.doi.org/10.2139/ssrn.3550601.

Joseph Clark. The replicating portfolio of a constant product market with bounded liquidity. *SSRN Manuscript*, 2021. doi: https://dx.doi.org/10.2139/ssrn.3898384.

Jennifer S. Conrad, Robert F. Dittmar, and Eric Ghysels. Ex ante skewness and expected stock returns. *Journal of Finance*, 68(1):85–124, February 2013.

Simon Cousaert, Jiahua Xu, and Toshiko Matsui. Sok: Yield aggregators in defi. *arXiv Manuscript*, 2022. doi: https://doi.org/10.1109/ICBC54727.2022.9805523.

Kresimir Demeterfi, Emanuel Derman, Michael Kamal, and Joseph Zou. A guide to volatility and variance swaps. *The Journal of Derivatives*, 6(4):9–32, 1999. ISSN 1074-1240.

Joost Driessen, Pascal Maenhout, and Grigory Vilkov. Option-implied correlations and the price of correlation risk. Working paper, INSEAD, 2005.

Joost Driessen, Pascal J. Maenhout, and Grigory Vilkov. The price of correlation risk: Evidence from equity options. *Journal of Finance*, 64(3):1377–1406, 2009. ISSN 1540-6261.

Joost Driessen, Pascal Maenhout, and Grigory Vilkov. Option-implied correlations and the price of correlation risk. *SSRN Manuscript*, 2016. doi: https://dx.doi.org/10.2139/ssrn.2166829.

Eugene F. Fama and James D. MacBeth. Risk return and equilibrium: empirical tests. *Journal of Financial Political Economy*, 81(3):607–636, 1973. ISSN 0022-3808.

Stephen Figlewski. Risk-neutral densities: A review. *Annual Review of Financial Economics*, 10:329–359, 2018. ISSN 1941-1367.

Jean-Pierre Fouque, George Papanicolaou, Ronnie Sircar, and Knut Sølna. *Multiscale Stochastic Volatility for Equity, Interest Rate, and Credit Derivatives*. Cambridge University Press, Cambridge, U.K., 1 edition, 2011. ISBN 9780521843584.

Nikolay Gospodinov, Raymond Kan, and Cesare Robotti. Misspecification-robust inference in linear asset-pricing models with irrelevant risk factors. *The Review of Financial Studies*, 27(7):2139–2170, 2014. ISSN 0893-9454.

Paolo Guasoni and Eberhard Mayerhofer. Technical note—options portfolio selection. *Operations Research*, 68(3):733–740, 2020. ISSN 0030-364X.

Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of Political Economy*, 99(2):225–262, 1991. ISSN 0022-3808.

Campbell R. Harvey, Joel Hasbrouck, and Fahad Saleh. The evolution of decentralized exchange: Risks, benefits, and oversight, 2024.

Lioba Heimbach, Eric Schertenleib, and Roger Wattenhofer. Risks and returns of uniswap V3 liquidity providers. In *4th ACM Conference on Advances in Financial Technologies*, Cambridge, Massachusetts, U.S.A., 2022. ISBN 9781450398619.

Raymond Kan, Cesare Robotti, and Jay Shanken. Pricing model performance and the two-pass cross-sectional regression methodology. *The Journal of Finance*, 68(6):2617–2649, 2013. ISSN 1540-6261.

Roger Lee. *Encyclopedia of Quantitative Finance*, chapter Corridor Variance Swap. John Wiley & Sons, 2010a. ISBN 9780470057568.

Roger Lee. *Encyclopedia of Quantitative Finance*, chapter Weighted variance swap. John Wiley & Sons, 2010b. ISBN 9780470057568.

Alfred Lehar and Christine Parlour. Decentralized exchange: The uniswap automated market maker. *SSRN Manuscript*, 2023. doi: https://dx.doi.org/10.2139/ssrn.3905316.

Thomas Nanfeng Li, Siddharth Naik, Andrew Papanicolaou, and Lorenzo Schönleber. Yield farming for liquidity provision. *SSRN Manuscript*, 2024. doi: https://dx.doi.org/10.2139/ssrn.4422213.

François Longin and Bruno Solnik. Extreme correlation of international equity markets. *Journal of Finance*, 56(2):649–676, 2001. ISSN 1540-6261.

William Margrabe. The value of an option to exchange one asset for another. *The Journal of Finance*, 33(1):177–186, 1978. ISSN 1540-6261.

Ian Martin. What is the Expected Return on the Market? *The Quarterly Journal of Economics*, 132(1):367–433, 2016. ISSN 0033-5533.

Basile Maire Masaaki Fukasawa and Marcus Wunsch. Weighted variance swaps hedge against impermanent loss. *Quantitative Finance*, 23(6):901–911, 2023. ISSN 1469-7688.

Jason Milionis, Ciamac C. Moallemi, Tim Roughgarden, and Anthony Lee Zhang. Automated market making and loss-versus-rebalancing. *arXiv Manuscript*, 2024. doi: https://doi.org/10.48550/arXiv.2208.06046.

Whitney K. Newey and Kenneth D. West. A simple, positive-semidefinite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. ISSN 0012-9682.

Alberto Quaini, Fabio Trojani, and Ming Yuan. Tradable factor risk premia and oracle tests of asset pricing models. *Swiss Finance Institute Research Paper No. 23-81*, 2023. doi: https://dx.doi.org/10.2139/ssrn.4574683.

Lorenzo Schönleber. Correlations, value factor returns, and growth options. *SSRN Manuscript*, 2023. doi: https://dx.doi.org/10.2139/ssrn.3480606.

Steven E. Shreve. *Stochastic Calculus for Finance II Continuous-Time Models*. Springer Finance. Springer Science+Business Media, New York, U.S.A., 1 edition, 2004. ISBN 9780387401010.

Vasiliki D. Skinzi and Apostolos-Paul N. Refenes. Implied correlation index: A new measure of diversification. *Journal of Futures Markets*, 25(2):171–197, 2005. ISSN 0270-7314.

Appendix A Mathematics for Impermanent Loss

A.1 Proof of Proposition 3.1

By the relative price that is defined in equation (3.2) and Assumption 3.1, we have $\frac{N_2(t)}{N_1(t)} = R(t) = \frac{P_1(t)}{P_2(t)}$. Based on equations (3.3), (3.5), and (3.6), using Itô's lemma we can show that

$$d\left(N_1(t) P_1(t) + N_2(t) P_2(t)\right) - \left(N_1(t) dP_1(t) + N_2(t) dP_2(t)\right) \tag{A.1}$$

$$= -\frac{\sigma_R^2(t)}{4} N_1(t) P_1(t) dt.$$

Equation (A.1) shows that the numerator in equation (3.7) has only a dt term, and so we take the limit in the denominator of equation (3.7) as Δt goes to zero, giving us

$$dIL(t) = -\frac{\sigma_R^2(t) N_1(t) P_1(t) dt}{4(N_1(t) P_1(t) + N_2(t) P_2(t))} \tag{A.2}$$

Then again, from equation (3.2) and Assumption 3.1, we have $\frac{N_2(t)P_2(t)}{N_1(t)P_1(t)} = 1$, consequently equation (A.2) reduces to equation (3.8).

A.2 Change of Numéraire and Girsanov Theorem

Let us consider a risk-neutral probability measure $\mathbb{Q}$ that is equivalent to the physical probability measure $\mathbb{P}$ used in Section 3 over some finite time interval $t \in [0,\ T]$. Suppose r is the risk-free interest rate. The SDEs for the token price dynamics under $\mathbb{Q}$ are

$$dP_i(t) = rP_i(t) dt + \sigma_i(t) P_i(t) dB_i^Q(t) \quad i = 1,\ 2, \tag{A.3}$$

where $dB_i^Q(t) = dB_i(t) + \frac{\mu_i - r}{\sigma_i(t)} dt$ is a SBM under $\mathbb{Q}$, and $dB_1^Q(t) dB_2^Q(t) = \rho dt$. The change of measure also alters the probabilities for the volatilities $\sigma_i(t)$, but these processes remain $\mathcal{F}_t$-adapted. Details on measure changes for stochastic volatility models can be found in Chapter 2 of Fouque et al. (2011). It is straightforward to observe that the discounted token prices given by equation (3.5) are martingales under $\mathbb{Q}$. By utilizing Itô's lemma on $\ln P_i(t)$, their solutions are

$$P_i(t) = P_i(0) \exp\left(\int_0^t \left(r - \frac{1}{2}\sigma_i^2(s)\right) ds + \int_0^t \sigma_i(s) dB_i^Q(s)\right). \tag{A.4}$$

Then, by implementing Itô's lemma on equation (A.3) we can derive the risk-neutral SDE for the relative price under the same $\mathbb{Q}$,

$$dR(t) = d\left(\frac{P_1(t)}{P_2(t)}\right) = \mu_R^Q(t) R(t) dt + \sigma_R(t) R(t) R(t) B_R^Q(t),$$

where is $\mu_R^Q(t) = \sigma_2^2(t) - \rho\sigma_1(t)\sigma_2(t)$, $\sigma_R^2(t) = \sigma_1^2(t) + \sigma_2^2(t) - 2\rho\sigma_1(t)\sigma_2(t)$, and $B_R^Q(t) = \frac{\sigma_1(t)}{\sigma_R(t)} B_1^Q(t) - \frac{\sigma_2(t)}{\sigma_R(t)} B_2^Q(t)$ is a SBM under $\mathbb{Q}$. Please notice that $R(t)$ is not a martingale under $\mathbb{Q}$. The details of constructing a risk-neutral probability measure can

be found in Shreve (2004).

To convert the relative price $R(t)$ to a martingale, we can apply the technique of change of numéraire. We choose the token $P_2(t)$ as the numéraire, and define a new probability measure $\widetilde{\mathbb{Q}}$ by the following Radon-Nikodym density:

$$\left.\frac{d\widetilde{\mathbb{Q}}}{d\mathbb{Q}}\right|_T := \frac{P_2(T)}{\mathbb{E}^Q[P_2(T)]} = \frac{P_2(T)}{e^{rT}P_2(0)}. \tag{A.5}$$

We can show that the relative price $R(t)$ defined by equation (3.4) in Assumption 3.1 is a martingale under $\widetilde{\mathbb{Q}}$ defined by equation (A.5):

$$\widetilde{\mathbb{E}}^Q[R(T)] = \mathbb{E}^Q\left[\frac{P_2(T)}{e^{rT}P_2(0)}\frac{P_1(T)}{P_2(T)}\right] = \mathbb{E}^Q\left[\frac{P_1(T)}{e^{rT}P_2(0)}\right] = \frac{e^{rT}P_1(0)}{e^{rT}P_2(0)} = R(0).$$

Also, given the density of equation (A.5), we can derive the spread option formulas given by equation (**??**),

$$\begin{aligned}
Call_{sp}(T,\,K) &= e^{-rT}\mathbb{E}^Q\left[(P_1(T) - KP_2(T))^+\right], \\
&= P_2(0)\,\mathbb{E}^Q\left[\frac{P_2(T)}{e^{rT}P_2(0)}(R(T) - K)^+\right], \\
&= P_2(0)\,\widetilde{\mathbb{E}}^Q\left[(R(T) - K)^+\right], \\
Put_{sp}(T,\,K) &= P_2(0)\,\widetilde{\mathbb{E}}^Q\left[(K - R(T))^+\right].
\end{aligned}$$

Next, let us assume that $X(t)$ is a stochastic process defined as follows:

$$X(t) = \int_0^t \sigma_2(s)\,dB_2^Q(s).$$

Because $B_2^Q(t)$ is a martingale under $\mathbb{Q}$, therefore, it is obvious that $X(t)$ is a martingale under $\mathbb{Q}$ as well. The Doléans-Dade exponential of $X(t)$ at time $t = T$ is

$$\left.\frac{d\widetilde{\mathbb{Q}}}{d\mathbb{Q}}\right|_T = \mathcal{E}(X(T)) = \exp\left(X(T) - \frac{1}{2}\int_0^T \sigma_2^2(t)\,dt\right) = \frac{P_2(T)}{e^{rT}P_2(0)}.$$

Girsanov theorem states that given $W^Q(t)$ a SBM under $\mathbb{Q}$, we define the following stochastic process:

$$\widetilde{W}^Q(t) := W^Q(t) - \left\langle W^Q,\, X\right\rangle(t), \tag{A.6}$$

where $\left\langle W^Q,\, X\right\rangle$ denotes the quadratic variation of these two stochastic processes, and

then $\widetilde{W}^{Q}(t)$ is a SBM under $\widetilde{\mathbb{Q}}$. Consequently, $\widetilde{B}_2^{Q}(t) = B_2^{Q}(t) - \int_0^t \sigma_2(s)\,ds$ is a SBM under $\widetilde{\mathbb{Q}}$. Suppose that the two Brownian motions $W^{Q}(t)$ and $B_2^{Q}(t)$ are correlated $dB_2(t)\,dW^{Q}(t) = \rho dt$. Then by the Girsanov theorem, equation (A.6) becomes

$$\widetilde{W}^{Q}(t) = W^{Q}(t) - \left\langle W^{Q},\, X \right\rangle(t) = W^{Q}(t) - \rho \int_0^t \sigma_2(s)\,ds. \tag{A.7}$$

Therefore, $\widetilde{W}^{Q}(t)$ given by equation (A.7) is SBM under $\widetilde{\mathbb{Q}}$. Consequently, $\widetilde{B}_1^{Q}(t) = B_1^{Q}(t) - \rho \int_0^t \sigma_2(s)\,ds$ is a SBM under $\widetilde{\mathbb{Q}}$.

By applying the above results to the token prices given by equation (A.3), under $\widetilde{\mathbb{Q}}$ defined by equation (A.5) we have

$$\begin{aligned} dP_1(t) &= rP_1(t)\,dt + \sigma_1(t)\,P_1(t)\left(dB_1^{Q}(t) - \rho\sigma_2(t)\,dt + \rho\sigma_2(t)\,dt\right), \\ &= \left(r + \rho\sigma_1(t)\,\sigma_2(t)\right)P_1(t)\,dt + \sigma_1(t)\,P_1(t)\,d\widetilde{B}_1^{Q}(t), \\ dP_2(t) &= rP_2(t)\,dt + \sigma_2(t)\,P_2(t)\left(dB_2^{Q}(t) - \sigma_2(t)\,dt + \sigma_2(t)\,dt\right), \\ &= \left(r + \sigma_2^2(t)\right)P_2(t)\,dt + \sigma_2(t)\,P_2(t)\,d\widetilde{B}_2^{Q}(t). \end{aligned} \tag{A.8}$$

Please notice that $\frac{e^{rt}}{P_2(t)}$ is a martingale under $\widetilde{\mathbb{Q}}$, in other words, the discounted price of the base token is a martingale under this new measure. Finally, by utilizing Itô's lemma and equation (A.8) we have

$$dR(t) = d\left(\frac{P_1(t)}{P_2(t)}\right) = R(t)\left(\sigma_1(t)\,d\widetilde{B}_1^{Q}(t) - \sigma_2(t)\,d\widetilde{B}_2^{Q}(t)\right).$$

A.3 Variance Swap and Impermanent Loss

Assuming constant risk-free interest rate and zero dividend yield, as shown in Carr and Madan (1998), Demeterfi et al. (1999), Carr and Wu (2009), and Bakshi et al. (2015), the valuation of variance swap under the risk-neutral probability measure $\mathbb{Q}$ is

$$\begin{aligned} \mathbb{E}^{Q}\left[\int_0^T \sigma_i^2(t)\,dt\right] &= 2\mathbb{E}^{Q}\left[\int_0^T \frac{dP_i(t)}{P_i(t)}\right] - 2\mathbb{E}^{Q}\left[\ln\left(\frac{P_i(T)}{P_i(0)}\right)\right], \\ &= -2\mathbb{E}^{Q}\left[\ln\left(\frac{P_i(T)}{e^{rT}P_i(0)}\right)\right], \end{aligned} \tag{A.9}$$

because by equation (A.3) we have $\mathbb{E}^Q\left[\int_0^T \frac{dP_i(t)}{P_i(t)}\right] = rT$. Consequently, by using equation (A.4), the second line of equation (A.9) with respect to token price $P_1(T)$ under $\widetilde{\mathbb{Q}}$ is

$$\widetilde{\mathbb{E}}^Q\left[\ln\left(\frac{P_1(T)}{e^{rT}P_1(0)}\right)\right] = \widetilde{\mathbb{E}}^Q\left[-\int_0^T \frac{1}{2}\sigma_1^2(t)\,dt + \int_0^T \sigma_1(t)\,dB_1^Q(t)\right]. \tag{A.10}$$

As shown through the change of measure in Appendix A.2, we have $\widetilde{B}_1^Q(t) = B_1^Q(t) - \rho\int_0^t \sigma_2(s)\,ds$, therefore by using the property $\mathbb{E}^{\tilde{Q}}\left[\int_0^T \sigma_1(t)\,dB_1^{\tilde{Q}}(t)\right] = 0$, and substituting $B_1^Q(t)$ with $\widetilde{B}_1^Q(t)$, equation (A.10) becomes

$$\widetilde{\mathbb{E}}^Q\left[\ln\left(\frac{P_1(T)}{e^{rT}P_1(0)}\right)\right] = \widetilde{\mathbb{E}}^Q\left[\int_0^T \left(\rho\sigma_1(t)\,\sigma_2(t) - \frac{1}{2}\sigma_1^2(t)\right)dt\right]. \tag{A.11}$$

Similarly, the change of measure in Appendix A.2 also shows that $\widetilde{B}_2^Q(t) = B_2^Q(t) - \int_0^t \sigma_2(s)\,ds$, hence by using equation (A.4), the property $\mathbb{E}^{\tilde{Q}}\left[\int_0^T \sigma_2(t)\,dB_2^{\tilde{Q}}(t)\right] = 0$, and replacing $B_2^Q(t)$ with $\widetilde{B}_2^Q(t)$, the second line of equation (A.9) with respect to token price $P_2(T)$ under $\widetilde{\mathbb{Q}}$ is

$$\widetilde{\mathbb{E}}^Q\left[\ln\left(\frac{P_2(T)}{e^{rT}P_2(0)}\right)\right] = \widetilde{\mathbb{E}}^Q\left[\int_0^T \frac{1}{2}\sigma_2^2(t)\,dt\right]. \tag{A.12}$$

Furthermore, after subtracting equation (A.12) from equation (A.11) on the left-hand and right-hand sides respectively, we can get

$$\begin{aligned}\widetilde{\mathbb{E}}^Q\left[\ln\left(\frac{R(T)}{R(0)}\right)\right] &= -\frac{1}{2}\widetilde{\mathbb{E}}^Q\left[\int_0^T \left(\sigma_1^2(t) - 2\rho\sigma_1(t)\,\sigma_2(t) + \sigma_2^2(t)\right)dt\right], \\ &= -\frac{1}{2}\widetilde{\mathbb{E}}^Q\left[\int_0^T \sigma_R^2(t)\,dt\right].\end{aligned} \tag{A.13}$$

Eventually, by following equation (3.8) in Proposition 3.1, we can obtain the expectation of the impermanent loss under $\widetilde{\mathbb{Q}}$ using equation (A.13):

$$\widetilde{\mathbb{E}}^Q\left[IL(T)\right] = -\frac{1}{8}\widetilde{\mathbb{E}}^Q\left[\int_0^T \sigma_R^2(t)\,dt\right] = \frac{1}{4}\widetilde{\mathbb{E}}^Q\left[\ln\left(\frac{R(T)}{R(0)}\right)\right],$$

where $R(t)$ is the relative price defined by equation (3.6), $\sigma_R^2(t) = \sigma_1^2(t) - 2\rho\sigma_1(t)\,\sigma_2(t) + \sigma_2^2(t)$ is the variance of the relative price.

A.4 Margrabe Formula

As we can observe from equation (3.8), one of the key components is the correlation ρ between the token prices $P_1(t)$ and $P_2(t)$. To calculate ρ, we first determine the IV of the spread option using the formula provided by Margrabe (1978) for pricing European spread call and put options. The IV of a European spread option with maturity T at strike price K is $\widehat{v}(T,\, K)$ such that

$$\begin{aligned} Call_{sp}^{mk}(T,\, K) &= P_1(0)\,\Phi(d_1) - KP_2(0)\,\Phi(d_2), \\ Put_{sp}^{mk}(T,\, K) &= KP_2(0)\,\Phi(-d_2) - P_1(0)\,\Phi(-d_1), \end{aligned}$$

where $Call_{sp}^{mk}(T,\, K)$ and $Put_{sp}^{mk}(T,\, K)$ are the spread option prices given by the market data, $\Phi(\cdot)$ denotes the standard normal cumulative distribution function, d_1 and d_2 are given by

$$d_1 = \frac{\ln\frac{P_1(0)}{P_2(0)} + \frac{1}{2}\widehat{v}^2(T,\, K)\,T}{\widehat{v}(T,\, K)\sqrt{T}}, \qquad d_2 = d_1 - \widehat{v}(T,\, K)\sqrt{T}.$$

Appendix B Computation of the Joint Density

As in Section 4.2, let Π denote the excess return on a portfolio. Suppose that $(X,\, Y) \in \mathbb{R}^2$ are the two state variables, which are random variables $X \in \{x_1,\, x_2,\, \cdots,\, x_m\}$ and $Y \in \{y_1,\, y_2,\, \cdots,\, y_n\}$ on a probability space $(\Omega,\, \mathcal{F},\, \mathbb{P})$. Their joint probability distribution under this physical probability measure $\mathbb{P}$ is

$$p_{ij} = \mathbb{P}(X = x_i,\, Y = y_j). \tag{B.1}$$

Suppose that their joint probability distribution under an equivalent risk-neutral probability measure $\mathbb{Q}$ is

$$q_{ij} = \mathbb{Q}(X = x_i,\, Y = y_j). \tag{B.2}$$

By using equations (B.1) and (B.2), we then have a pricing kernel $M(x = x_i,\, y = y_j) \in \mathbb{R}^2$ such that

$$M_{ij} = \frac{q_{ij}}{p_{ij}},$$

We assume that the interest rate is zero such that $\mathbb{E}[M] = 1$. Because we know that $\operatorname{cov}(M,\, \Pi) = \mathbb{E}[M\Pi] - \mathbb{E}[M]\,\mathbb{E}[\Pi]$, then by utilizing the Cauchy–Schwarz inequality

we have $\left|\text{cov}\left(M,\ \Pi\right)\right|^2 \leqslant \text{std}^2\left(M\right)\text{std}^2\left(\Pi\right)$. Therefore, for any Π, p and M, we have

$$0 = \mathbb{E}\left[M\Pi\right] = \text{cov}\left(M,\ \Pi\right) + \mathbb{E}\left[\Pi\right] \geqslant -\text{std}\left(M\right)\text{std}\left(\Pi\right) + \mathbb{E}\left[\Pi\right],$$

where $\mathbb{E}$ is the expectation operator under $\mathbb{P}$, $\text{std}\left(\cdot\right)$ is the standard deviation, and $\text{cov}\left(\cdot\right)$ is the covariance. This yields the HJ bound in terms of Sharpe ratio $\frac{\mathbb{E}[\Pi]}{\text{std}(\Pi)}$,

$$\sup_{\Pi:\text{std}(\Pi)>0} \frac{\mathbb{E}\left[\Pi\right]}{\text{std}\left(\Pi\right)} \leqslant \inf_{M\in\mathcal{M}} \text{std}\left(M\right) \tag{B.3}$$

where $\mathcal{M}$ denotes the family of pricing kernels. The right-hand side of equation (B.3) motivates us to find a pricing kernel that minimizes the HJ bound.

B.1 The Unconstrained Case

We start from an unconstrained case of quadratic programming (QP) problem. For the joint probability distributions given by equations (B.1) and (B.2), define $\phi_{ij} := \frac{q_{ij}}{\sqrt{p_{ij}}}$, therefore we have $p_{ij} = \sum_{i,j}\phi_{ij}^2 - 1$ and the variance, denote by $\text{var}\left(\cdot\right)$, of M is $\text{var}\left(M\right) = \sum_{i,j}\left(\frac{q_{ij}}{p_{ij}} - 1\right)^2$. We then construct an unconstrained QP problem for the HJ bound defined by equation (B.3):

$$\begin{aligned} \underset{\phi:\phi_{ij}>0}{\text{minimize}}\ & \sum_{i,j}\phi_{ij}^2, \\ \text{subject to}\ & \sum_{i,j}\phi_{ij}\sqrt{p_{ij}} = 1 \end{aligned} \tag{B.4}$$

There exists a solution to the QP problem (B.4). We can rewrite the problem in a matrix form as

$$\begin{aligned} \underset{\boldsymbol{\phi}}{\text{minimize}}\ & \boldsymbol{\phi}^\top\boldsymbol{\phi}, \\ \text{subject to}\ & \boldsymbol{\phi}^\top\sqrt{\boldsymbol{p}} = 1, \\ & \boldsymbol{\phi}_i > 0 \quad \forall i \leqslant mn, \end{aligned} \tag{B.5}$$

where $\boldsymbol{\phi}$ stacks all the columns of matrix $\left[\boldsymbol{\Phi}\right]_{ij} = \phi_{ij}$ into an $mn \times 1$ vector as follows:

$$\boldsymbol{\phi} = \begin{bmatrix} \phi_{11} & \phi_{21} & \cdots & \phi_{m1} & \phi_{12} & \cdots & \phi_{m2} & \cdots & \phi_{1n} & \phi_{2n} & \cdots & \phi_{mn} \end{bmatrix}^\top,$$

and $\boldsymbol{p}$ stacks all the columns of matrix $[\boldsymbol{P}]_{ij} = p_{ij}$ into an $m \times n$ vector as follows:

$$\boldsymbol{p} = \begin{bmatrix} p_{11} & p_{21} & \cdots & p_{m1} & p_{12} & \cdots & p_{m2} & \cdots & p_{1n} & p_{2n} & \cdots & p_{mn} \end{bmatrix}^\top,$$

and $\sqrt{\boldsymbol{p}}$ denotes that the square root function is applied on every entry of vector $\boldsymbol{p}$, in other words, $\left[\sqrt{\boldsymbol{p}}\right]_{ij} = \sqrt{p_{ij}}$. This matrix form given by equation (B.5) of the QP (B.4) can be solved with standard software packages such as SciPy, CVXPY, and MATLAB.[13]

B.2 Constrained Case for Marginals

Suppose that marginal probability distributions under $\mathbb{Q}$ are given as follows $\sum_j q_{ij} = \mu_i$ and $\sum_i q_{ij} = \nu_j$. We then construct a constrained QP problem for the HJ bound given by equation (B.3):

$$\begin{aligned} \underset{\phi:\phi_{ij}>0}{\text{minimize}} \quad & \sum_{i,j} \phi_{ij}^2, \\ \text{subject to} \quad & \sum_j \phi_{ij}\sqrt{p_{ij}} = \mu_i, \\ & \sum_i \phi_{ij}\sqrt{p_{ij}} = \nu_j. \end{aligned} \tag{B.6}$$

For each $i \leqslant m$ and each $j \leqslant n$, we define the following two matrices:

$$\boldsymbol{A}_i = \boldsymbol{I}_{n\times n} \otimes \left(\boldsymbol{e}_i \boldsymbol{e}_i^\top\right), \qquad \boldsymbol{B}_j = \operatorname{diag}\left(\boldsymbol{e}_j\right) \otimes \boldsymbol{I}_{m\times m}, \tag{B.7}$$

where $\boldsymbol{I}_{a\times b}$ is the identity matrix with dimension $a \times b$, $\boldsymbol{e}_i \in \mathbb{R}^m$ and $\boldsymbol{e}_j \in \mathbb{R}^n$ are the i^{th} and j^{th} canonical (or standard) basis vectors respectively, and $\otimes$ denotes the Kronecker product. To gain a sense for the operations performed by matrices $\boldsymbol{A}_i$ and $\boldsymbol{B}_j$ defined by equation (B.7), please take a look at the following matrix multiplications as examples:

$$\begin{aligned} \boldsymbol{A}_i\boldsymbol{\phi} &= \operatorname{diag}\left(\begin{bmatrix} \boldsymbol{e}_i\boldsymbol{e}_i^\top & \cdots & \boldsymbol{e}_i\boldsymbol{e}_i^\top & \cdots & \boldsymbol{e}_i\boldsymbol{e}_i^\top \end{bmatrix}\right)\boldsymbol{\phi}, \\ &= \begin{bmatrix} 0 & \cdots & 0 & \phi_{i1} & 0 & \cdots & 0 & \phi_{i2} & 0 & \cdots & 0 & \phi_{in} & 0 & \cdots & 0 \end{bmatrix}^\top, \quad \forall i \leqslant m, \\ \boldsymbol{B}_i\boldsymbol{\phi} &= \operatorname{diag}\left(\begin{bmatrix} \boldsymbol{0}_{m\times m} & \cdots & \boldsymbol{I}_{m\times m} & \cdots & \boldsymbol{0}_{m\times m} \end{bmatrix}\right)\boldsymbol{\phi}, \\ &= \begin{bmatrix} 0 & \cdots & 0 & \phi_{1j} & \phi_{2j} & \cdots & \phi_{mj} & 0 & \cdots & 0 \end{bmatrix}^\top, \quad \forall j \leqslant n, \end{aligned} \tag{B.8}$$

[13] See https://docs.scipy.org/doc/scipy/tutorial/optimize.html, https://www.cvxpy.org/, and https://www.mathworks.com/.

where diag $(\cdot)$ denotes a diagonal matrix, $\mathbf{0}_{m\times m}$ is a zero matrix with dimension $m \times m$, the dimension of matrix $\boldsymbol{A}_i$ and $\boldsymbol{B}_i$ is $(mn) \times (mn)$.

By observing equation (B.8), we can easily see that matrices $\boldsymbol{A}_i$ and $\boldsymbol{B}_j$ are symmetric $\boldsymbol{A}_i = \boldsymbol{A}_i^\top$ and $\boldsymbol{B}_j = \boldsymbol{B}_j^\top$, and are also idempotent $\boldsymbol{A}_i^2 = \boldsymbol{A}_i$ and $\boldsymbol{B}_j^2 = \boldsymbol{B}_j$. Thus, we can use equation (B.7) to construct the constraints of the QP problem (B.6) as follows:

$$\boldsymbol{\phi}^\top \boldsymbol{A}_i \sqrt{\boldsymbol{p}} = \sum_j \phi_{ij} \sqrt{p_{ij}} = \mu_i, \qquad \boldsymbol{\phi}^\top \boldsymbol{B}_j \sqrt{\boldsymbol{p}} = \sum_i \phi_{ij} \sqrt{p_{ij}} = \nu_j.$$

Furthermore, let us define two sparse vectors $\boldsymbol{u}_i = \boldsymbol{A}_i \sqrt{\boldsymbol{p}} \in \mathbb{R}^{mn}$ and $\boldsymbol{v}_j = \boldsymbol{B}_j \sqrt{\boldsymbol{p}} \in \mathbb{R}^{mn}$, then similar to equation (B.5), we can then rewrite the QP problem (B.6) in a matrix form as follows:

$$\begin{aligned} \underset{\boldsymbol{\phi}}{\text{minimize}} \;\; & \boldsymbol{\phi}^\top \boldsymbol{\phi}, \\ \text{subject to} \;\; & \boldsymbol{u}_i^\top \boldsymbol{\phi} = \mu_i \quad \forall i \leqslant m, \\ & \boldsymbol{v}_j^\top \boldsymbol{\phi} = \nu_j \quad \forall j \leqslant n, \\ & \boldsymbol{\phi}_k > 0 \qquad \forall k \leqslant mn. \end{aligned} \tag{B.9}$$

These sparse vectors $\boldsymbol{u}_i$ and $\boldsymbol{v}_j$ can be further integrated into a single constraint matrix $\boldsymbol{C} \in \mathbb{R}^{(m+n)\times(mn)}$ as follows:

$$\boldsymbol{C} = \begin{bmatrix} \boldsymbol{u}_1^\top & \boldsymbol{u}_2^\top & \cdots & \boldsymbol{u}_m^\top & \boldsymbol{v}_1^\top & \boldsymbol{v}_2^\top & \cdots & \boldsymbol{v}_n^\top \end{bmatrix}^\top. \tag{B.10}$$

It is easy to check that matrix $\boldsymbol{C}$ is $2mn$-sparse. By using equation (B.10), equation (B.9) can be rewritten succinctly as

$$\begin{aligned} \underset{\boldsymbol{\phi}}{\text{minimize}} \;\; & \boldsymbol{\phi}^\top \boldsymbol{\phi}, \\ \text{subject to} \;\; & \boldsymbol{C}\boldsymbol{\phi} = \begin{bmatrix} \boldsymbol{\mu} & \boldsymbol{\nu} \end{bmatrix}^\top, \\ & \boldsymbol{\phi}_k > 0 \quad \forall k \leqslant mn, \end{aligned} \tag{B.11}$$

where $[\boldsymbol{\mu}]_i = \mu_i$ and $[\boldsymbol{\nu}]_j = \nu_j$ are vectors of marginal probability distributions under $\mathbb{Q}$.

For the implementation in solving the QP problem (B.11), the entries of matrix $\boldsymbol{C}$ can be loaded before running the optimization. A software package is required to solve QP problems with $(m + n)$ sparse constraints. For instance, with $m = n = 100$, the problem involves 200 constraints and 10000×10000 sparse matrices. Suitable options

include SciPy, CVXPY, and MATLAB.

The Lagrange multipliers of the optimization problem are discussed in Internet Appendix I.

Appendix C Implied Impermanent Loss for Uniswap V3 and Implied Loss-Versus-Rebalancing

In the following, we briefly outline the application of our methodology, which is broadly applicable, for the valuation of impermanent loss in Uniswap V3—where liquidity providers can allocate tokens within specified price ranges—and to other adverse selection metrics, such as loss-versus-rebalancing (LVR), as discussed by Milionis et al. (2024).

C.1 Implied Impermanent Loss for Uniswap V3

The valuation formula in Proposition 4.1 applies to the V2 protocol, but many pools now operate under the V3 protocol, where LPs allocate tokens within a specific price range. Consider the ticks $(r_i)_{i=1,\,2,\,\ldots,\,m}$ where $r_{i+1} = r_i \times 1.0001$, and let $C_i = [r_i,\, r_{i+1})$ denote the band. If $R(t) \in C_i$ then the LP earns a reward but is also exposed to impermanent loss, in particular, it can be shown that the V3 analog to equation (3.8) is

$$dIL_i(t) = -\frac{\sigma_R^2(t)\sqrt{R(t)}}{4\left(2\sqrt{R(t)} - \sqrt{r_i} - \frac{R(t)}{\sqrt{r_{i+1}}}\right)}\mathbf{1}_{\{R(t)\in C_i\}}dt.$$

This increment of impermanent loss can be derived using Itô's lemma as done in deriving equation (3.8). Integrating over time yields the total impermanent loss, which is equivalent to a corridor variance swap (Lee (2010a,b)),

$$\begin{aligned}\widetilde{\mathbb{E}}^Q\left[IL_i(T)\right] &= -\frac{1}{4}\widetilde{\mathbb{E}}^Q\left[\int_0^T \frac{\sigma_R^2(t)\sqrt{R(t)}}{2\sqrt{R(t)} - \sqrt{r_i} - \frac{R(t)}{\sqrt{r_{i+1}}}}\mathbf{1}_{\{R(t)\in C_i\}}(t)\,dt\right], \qquad \text{(C.1)}\\ &= -\frac{1}{2}\int_{K\in C_i}\frac{\widetilde{\mathbb{E}}^Q\left[(K - R(T))^+\right]\mathbf{1}_{\{K<R(0)\}} + \widetilde{\mathbb{E}}^Q\left[(R(T) - K)^+\right]\mathbf{1}_{\{K\geqslant R(0)\}}}{\sqrt{K^3}\left(2\sqrt{K} - \sqrt{r_i} - \frac{K}{\sqrt{r_{i+1}}}\right)}dK.\end{aligned}$$

The pool-wide impermanent loss is a weighted sum of individual corridor variance swaps, with weights proportional to the TVL in each band, $\sum_i \mathrm{TVL}_i\widetilde{\mathbb{E}}^Q\left[IL_i(T)\right] / \sum_i \mathrm{TVL}_i$, where TVL_i is the TVL in band i.

C.2 Implied Loss-Versus-Rebalancing

Our framework can be used to assess loss-versus-rebalancing (LVR) Milionis et al. (2024), with the advantage of valuing LVR under $\mathbb{Q}$ introduced in Section 4.

Denote the TVL of a liquidity pool as $V(t) := N_1(t) P_1(t) + N_2(t) P_2(t)$. In our continuous-time framework, an increment of LVR in the V2 protocol is

$$dLVR(t) = \frac{1}{8} V(t) \sigma_R^2(t) dt,$$

where $N_i(t)$ and $P_i(t)$ for $i = 1, 2$ are the notation from Section 3.1 for number of tokens in the pool and the market prices. Consider the constant product rule as given in equation (3.1), where the parameter L remains constant. We can express the N_i in terms of L and $R(t)$, namely, $N_1(t) = \frac{L}{\sqrt{R(t)}}$ and $N_2(t) = L\sqrt{R(t)}$, and from Assumption 3.1 it follows that

$$V(t) = 2L\sqrt{P_1(t) P_2(t)}.$$

Then the valuation of LVR under $\mathbb{Q}$ is computed using the change of numéraire introduced in Section 4.3,

$$\mathbb{E}^Q[LVR(T)] = \frac{e^{rT} L P_2(0)}{4} \int_0^T \widetilde{\mathbb{E}}^Q \left[\sqrt{R(t)} \sigma_R^2(t) \right] dt.$$

Using the weighted variance swap framework of Lee (2010b), it follows that

$$\mathbb{E}^Q[LVR(T)] = -2L \left(\mathbb{E}^Q \sqrt{P_1(T) P_2(T)} - e^{rT} \sqrt{P_1(0) P_2(0)} \right) \tag{C.2}$$

Equation (C.2) can be computed using the estimated RND obtained by minimizing the HJ bound, as proposed in Section 4.2.

Implied LVR for non-constant L can be managed under assumptions such as L being uncorrelated with $R(t)$ and behaving as a martingale. In V3, implied LVR can be computed as a portfolio of options, similar to equation (C.1) for IL. Like V2 LVR, the implied V3 LVR is expressed in units of the original numéraire.

Appendix D Additional Figures

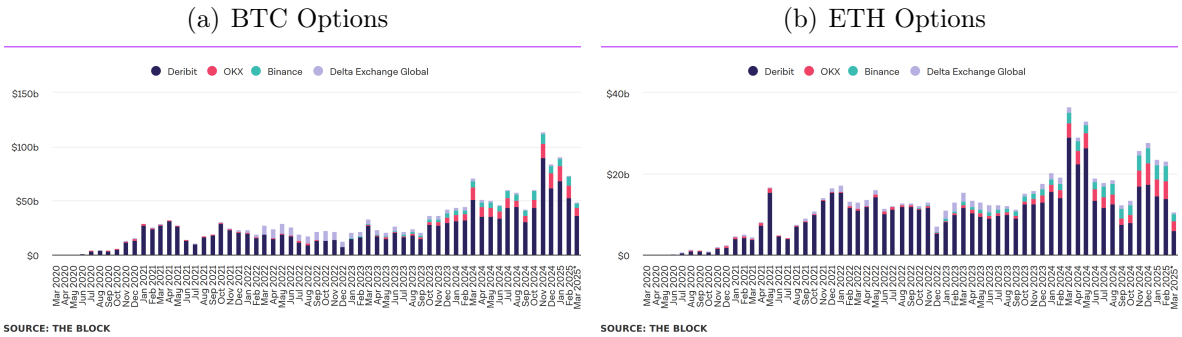

Figure D.1: Option Volumes across Exchanges. The figure reports the BTC and ETH options trading volume, in dollar terms, across cryptocurrency exchanges. The data is obtained from The Block.

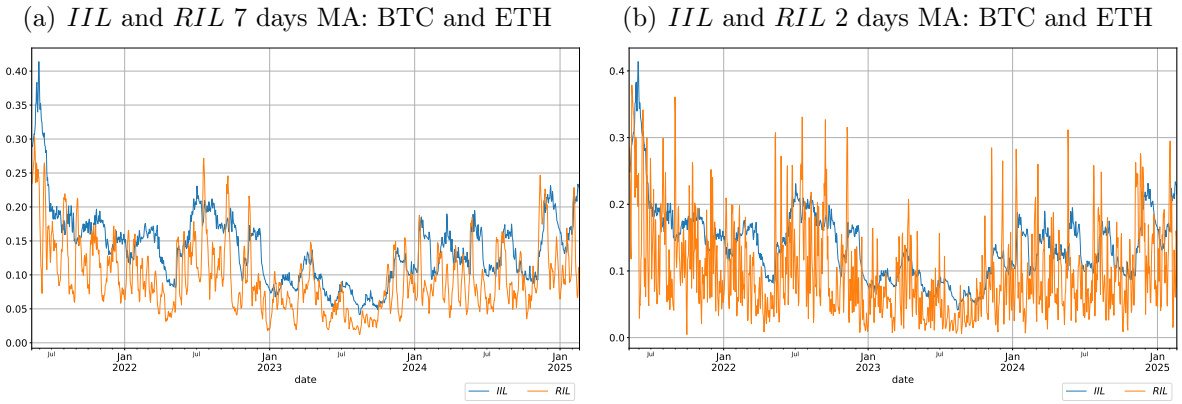

Figure D.2: IIL and RIL – BTC and ETH. The figure shows different moving averages for IIL and RIL of BTC and ETH. The implied quantity is calculated by the equation in Section 5.1. The realized quantity is calculated by the approach in Section 5.2. The data is sampled daily, and the sample period is from 2021-05 to 2023-11.

(a) IIL: β

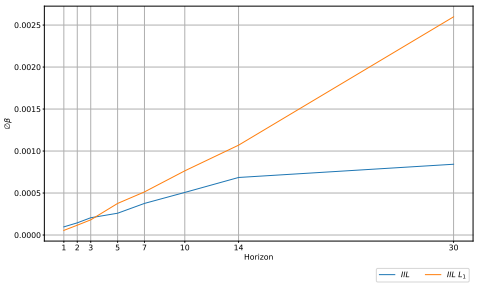

(b) IIL: t-statistic

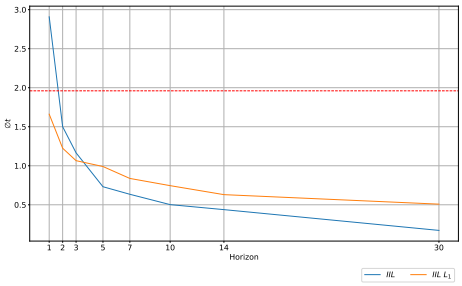

(c) IV_{BTC}: β

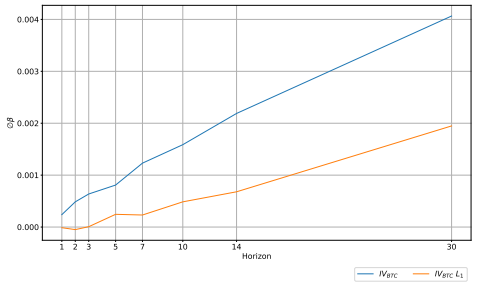

(d) IV_{BTC}: t-statistic

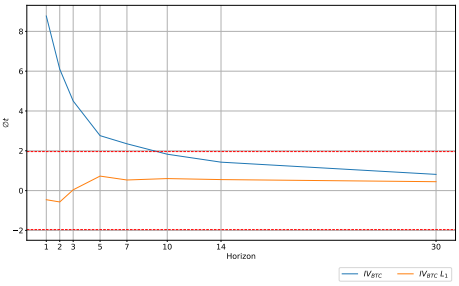

(e) IV_{ETH}: β

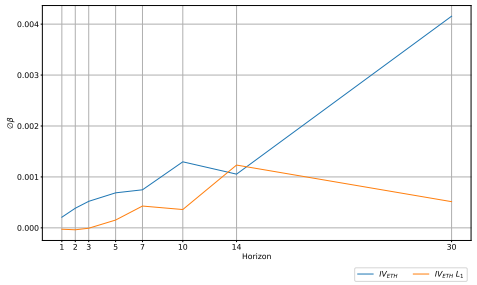

(f) IV_{ETH}: t-statistic

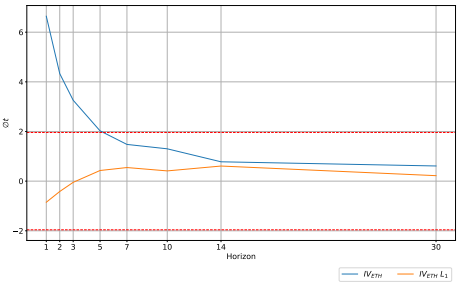

(g) IC: β

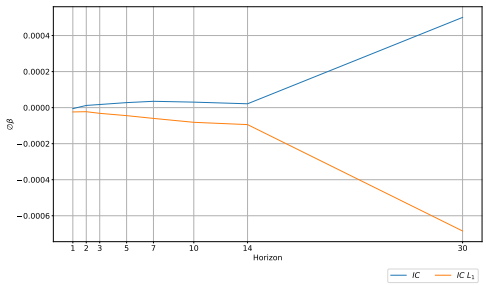

(h) IC: t-statistic

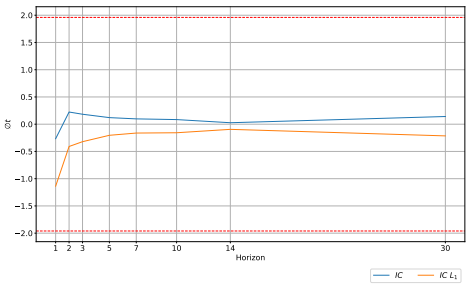

Figure D.3: Predictive – Non-Overlapping – Future APRs – Implied Measures and Lagged Values. The plots report the average β coefficients and their average t-statistic for the time-series regression as outlined in equation (6.1). The data is sampled at a frequency equal to the predictive time horizon (i.e., non-overlapping). As explanatory variables, IIL and its drivers IV_{BTC}, IV_{ETH}, IC, and their lagged values are considered. The implied quantities are calculated by the equations in Section 5.1. The sample period is from 2021-05 to 2025-02.

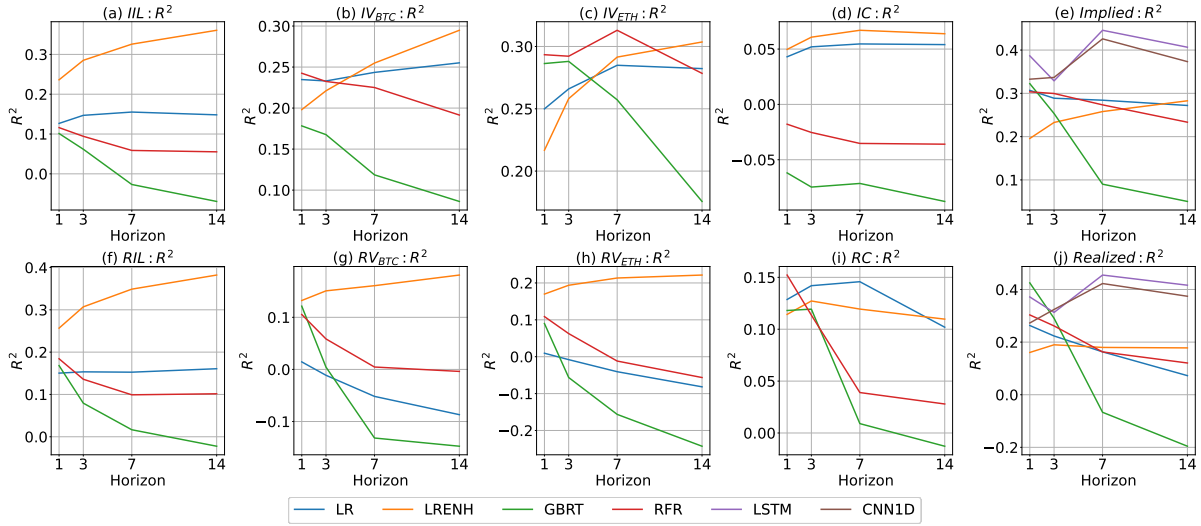

Figure D.4: Out-of-Sample Predictions – Machine Learning Models. The plots report the OOS R^2 obtained from predictions using the machine learning models for each predictor, calculated by equation (6.2) for predictions based on a 180-day rolling estimation window. We deploy the linear regression (LR) model, linear regression with elastic net regularisation and Huber loss (LRENH), gradient-boosted regression tree (GBRT), random forest regression (RFR), long short-term memory neural network (LSTM), and 1D convolutional neural network (CNN1D). The implied and realized quantities are calculated by the equations in Section 5. The sample period is from 2021-05 to 2023-11.

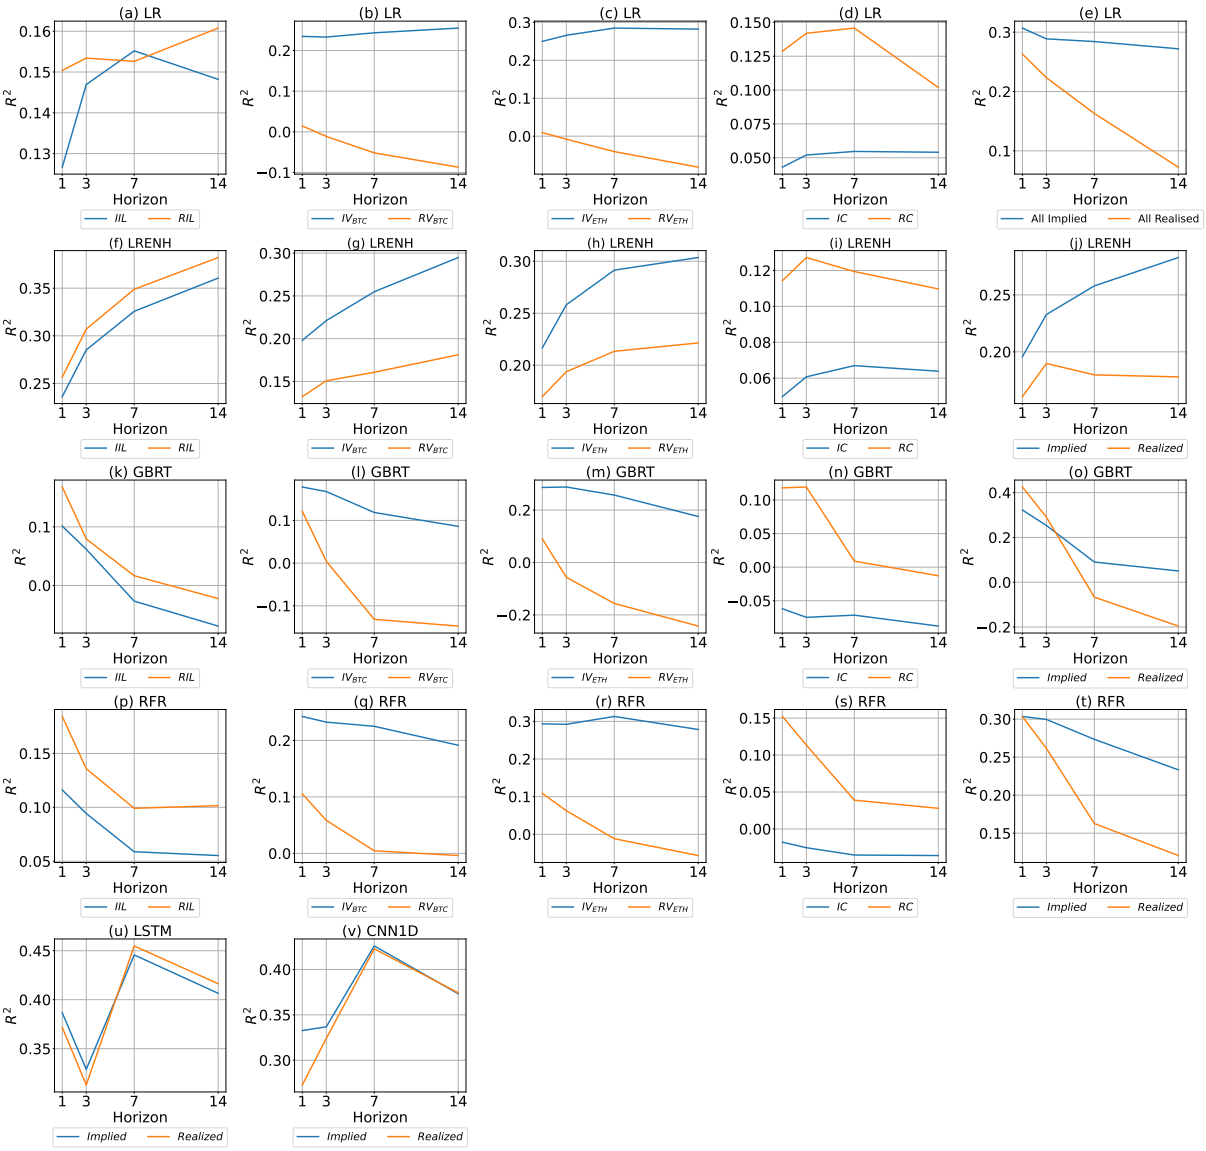

Figure D.5: Out-of-Sample Predictions – Machine Learning Models – Implied vs. Realized. The plots report the OOS R^2 obtained from predictions using the machine learning models, calculated by equation (6.2) for predictions based on a 180-day rolling estimation window. The OOS R^2 of the implied and realized counterparts are displayed in each panel. We deploy the linear regression (LR) model, linear regression with elastic net regularisation and Huber loss (LRENH), gradient-boosted regression tree (GBRT), random forest regression (RFR), long short-term memory neural network (LSTM), and 1D convolutional neural network (CNN1D). The implied and realized quantities are calculated by the equations in Section 5. The sample period is from 2021-05 to 2023-11.

(a) IIL: SDF Coefficient

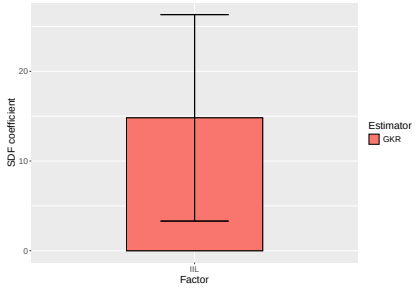

(b) IIL: Tradable Risk Premia

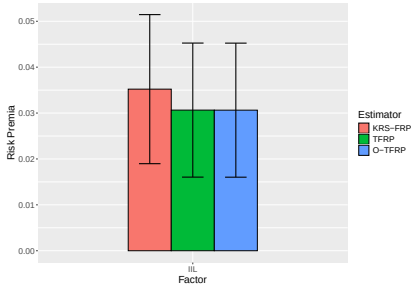

(c) IV_{BTC}: SDF Coefficient

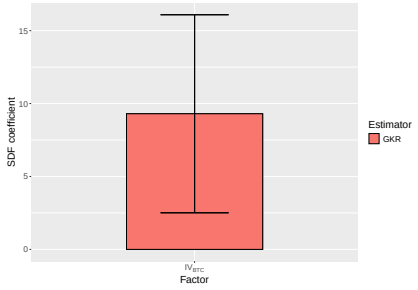

(d) IV_{BTC}: Tradable Risk Premia

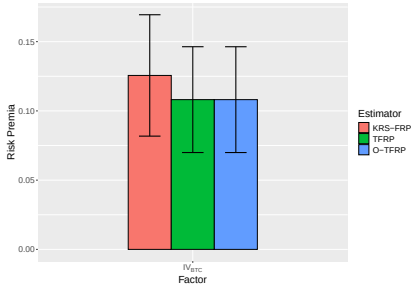

(e) IV_{ETH}: SDF Coefficient

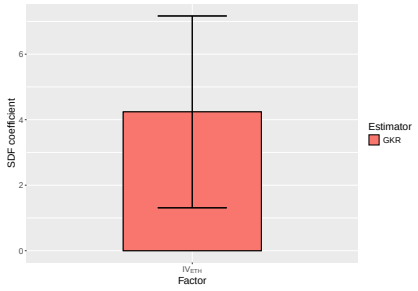

(f) IV_{ETH}: Tradable Risk Premia

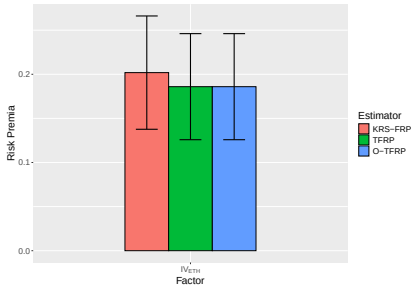

(g) IC: SDF Coefficient

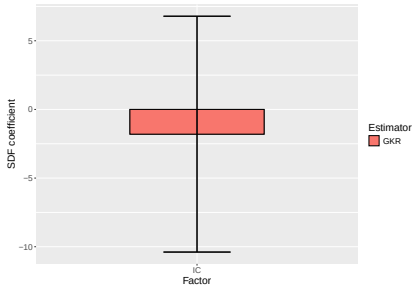

(h) IC: Tradable Risk Premia

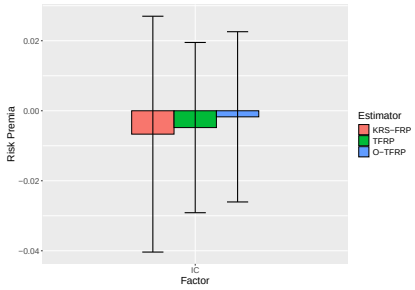

Figure D.6: Cross-Section – Univariate – Tradable Risk Premia. The figure presents the results of the SDF procedure described in Section 7.3. The left figures report the SDF coefficient and its standard error bars calculated using the GKR estimator. The right figures display the tradable risk premia calculated using the KRS–FRP, the TFRP, and the O-TFRP estimators. As risk factors, the levels of IIL and its drivers are considered, namely IV_{BTC}, IV_{ETH}, and IC. The implied quantities are calculated by the equations in Section 5.1. The sample period is from 2021-05 to 2023-11.

Internet Appendix

I Lagrange Multipliers

Proposition I.1. *Assume under the physical probability measure $\mathbb{P}$, the joint density is $p(x,\, y) > 0$ for all $(x,\, y)$. Then, the optimal pricing kernel given by equation (4.8) from the HJ upper bound in equation (4.6) is arbitrage-free if and only if $\lambda_1(x) + \lambda_2(y) > 0$ for all $(x,\, y)$.*

It is not necessary to demonstrate a proof for Proposition I.1 because it is a direct invocation of the fundamental theorem of asset pricing.

Proposition I.2. *Given optimal Lagrange multipliers $\lambda_1(x)$ and $\lambda_2(y)$ in equation (4.7) from the minimization of the HJ upper bound in equation (4.6), $\lambda_1(x)$ can be replicated by a portfolio of European call and put options on $P_1(T)$, forward contracts on $P_1(T)$, and cash. Similarly, $\lambda_2(y)$ can be replicated by a portfolio of European call and put options on $P_2(T)$, forward contracts on $P_2(T)$, and cash.*

Proof. Regularity of the Lagrange multiplier functions as proved in Guasoni and Mayerhofer (2020), where if μ and ν are sufficiently differentiable then so is the solution to the dual problem. Given this regularity, the formula of Carr and Madan (1998) is used to replicate the Lagrange multipliers,

$$\lambda_1(x) = \lambda_1(x_0) + \lambda_1'(x_0)(x - x_0) + \int_0^{x_0} \lambda_1''(K)(K - x)^+ dK + \int_{x_0}^{\infty} \lambda_1''(K)(K - x)^+ dK,$$

$$\lambda_2(y) = \lambda_2(y_0) + \lambda_2'(y_0)(y - y_0) + \int_0^{y_0} \lambda_2''(K)(K - y)^+ dK + \int_{y_0}^{\infty} \lambda_2''(K)(K - y)^+ dK,$$

where x_0 and y_0 are the chosen reference points, usually the forward prices of $P_1(T)$ and $P_2(T)$, respectively. For x, $\lambda_1(x_0)$ is a position in cash, $\lambda_1'(x_0)(x - x_0)$ is the net position in $\lambda_1'(x_0)$ many forward contracts, and the integrals are portfolios of out-of-the-money European call and put options; the same breakdown applies for y and $\lambda_2(y)$. □

II Realized Measures

(a) Realized Measures and Aggregated TVL

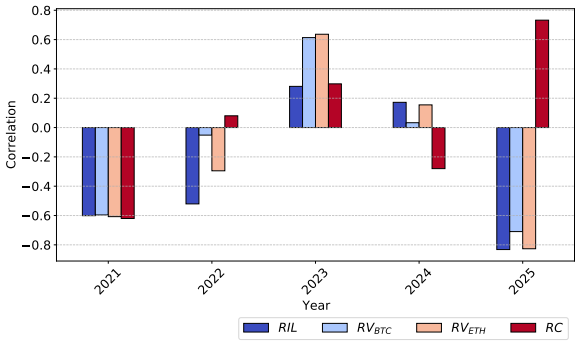

Figure II.1: Contemporaneous Correlations – Realized Measures and Aggregated Pool Size (TVL). The plots report the correlation coefficients calculated for each year separately between the realized quantities, and the TVL which is calculated as the aggregate TVL of all pools. The realized quantities are calculated as discussed in Section 5.2. The sample period is from 2021-05 to 2025-02.

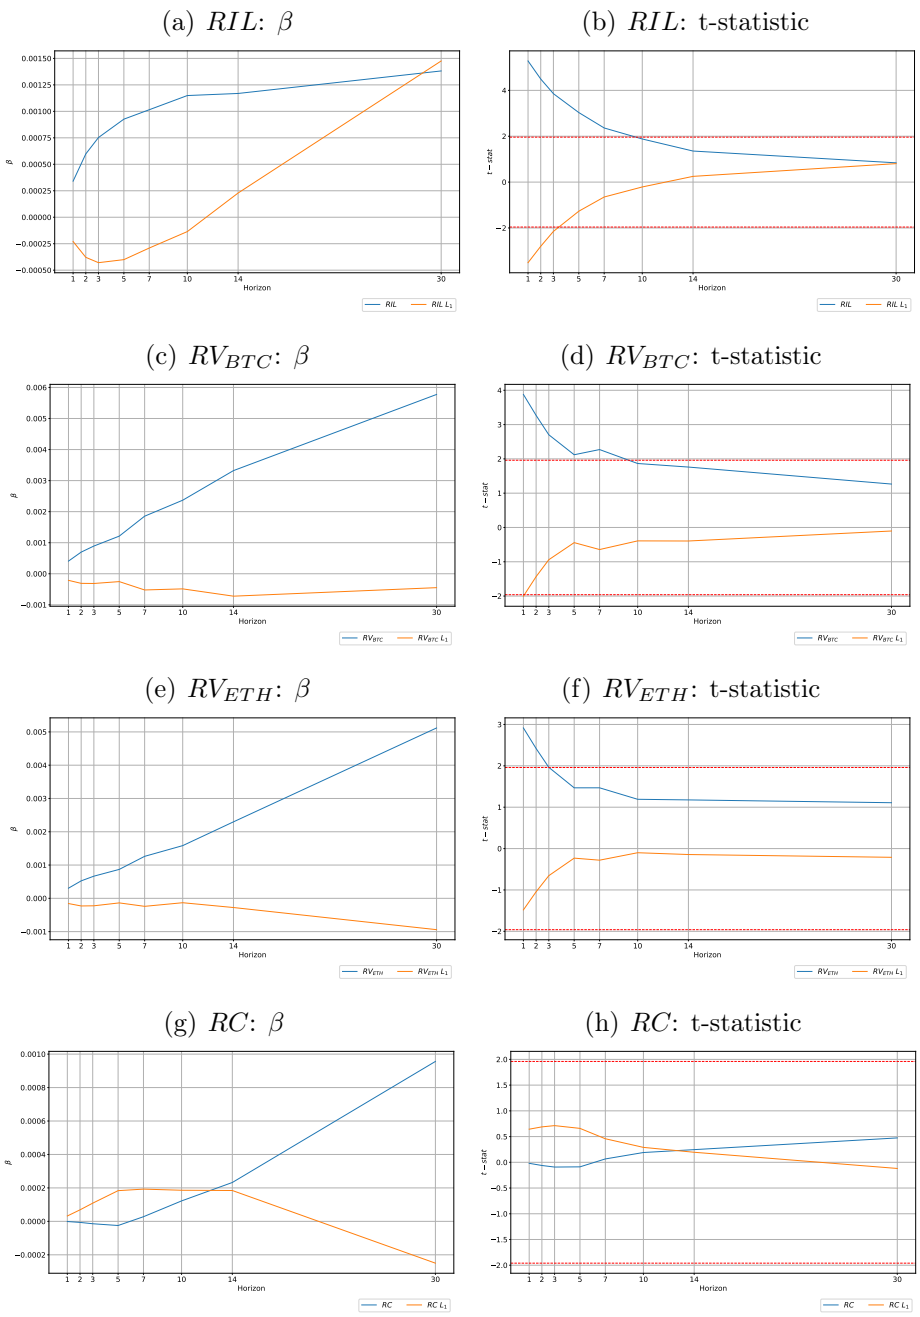

Figure II.2: Predictive – Future APRs – Realized Measures. The table reports the β coefficient and its t-statistic for the time-series regression described by equation (6.1). Standard errors are computed with Newey and West (1987). As explanatory variables, RIL, RV_{BTC}, RV_{ETH}, RC and their lagged values are considered. The realized quantities are calculated by the approaches in Section 5.2. The sample period is from 2021-05 to 2023-11.

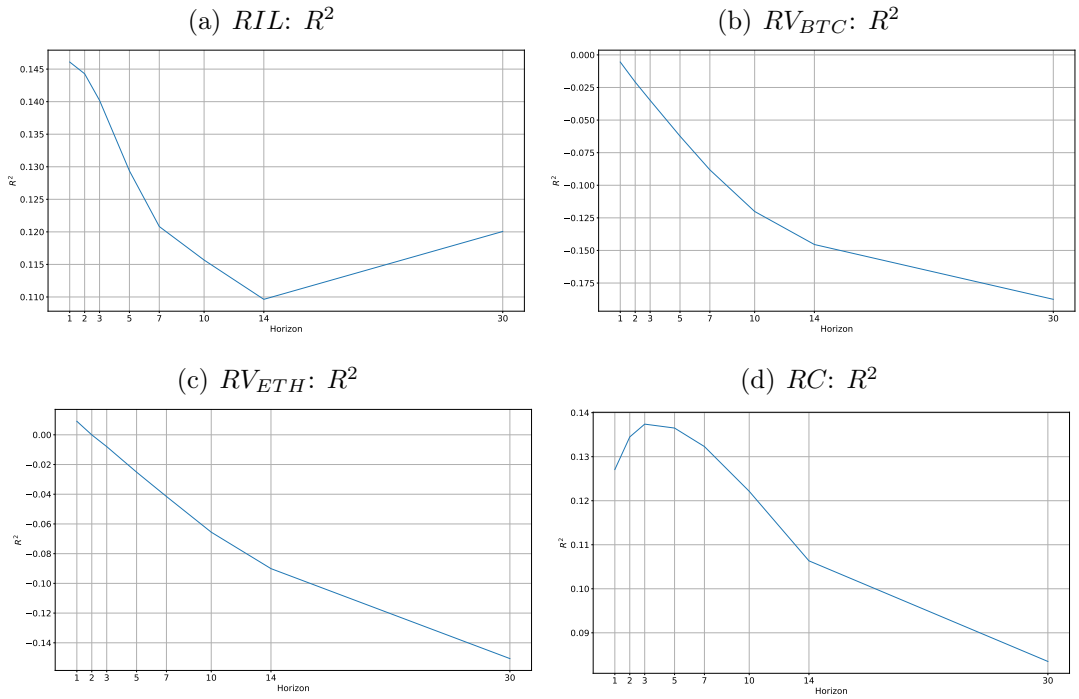

Figure II.3: Out-of-Sample Predictions – Realized Measures. The plots report the OOS R^2 calculated by equation (6.2) for predictions based on a 180-day rolling estimation window. The realized quantities are calculated by the approaches in Section 5.2. The sample period is from 2021-05 to 2023-11.

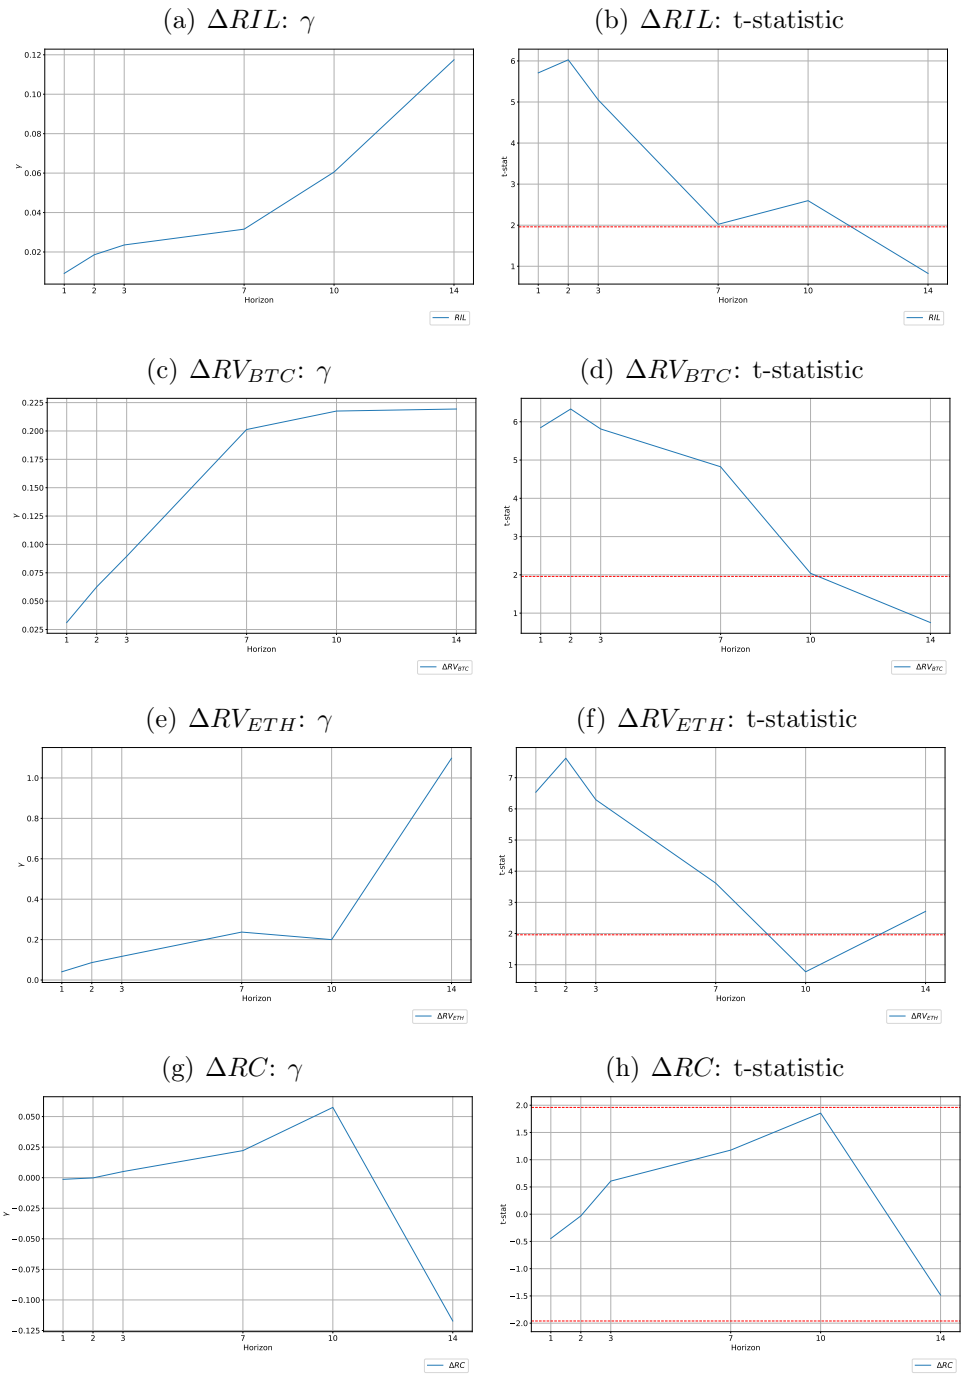

Figure II.4: Cross-Section – Univariate – Realized Measures. The table reports the γ coefficient and its t-statistic for the univariate Fama and MacBeth (1973) two-stage cross-sectional regression. As risk factors, the drivers of RIL are considered, namely the changes in RV_{BTC}, RV_{ETH}, and RC. In each regression we control for the changes in the individual realized impermanent loss. The realized quantities are calculated by the approaches in Section 5.2. The sample period is from 2021-05 to 2025-02.

(a) ΔRIL_{ind}: γ (b) RIL_{ind}: t-statistic

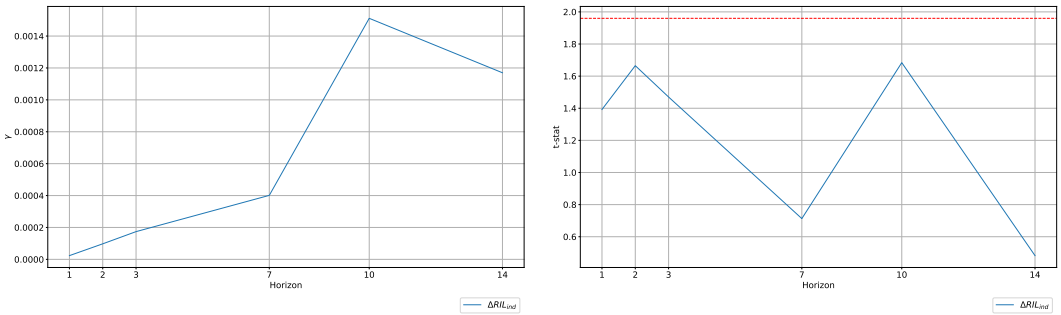

Figure II.5: Cross-Section – Univariate – Realized Individual Impermanent Loss. The table reports the γ coefficient and its t-statistic for the univariate Fama and MacBeth (1973) two-stage cross-sectional regression. As risk factors, changes in the individual daily RIL_{ind} given by equation (5.3) are considered. The sample period is from 2021-05 to 2025-02.