Journal of Financial Market Infrastructures **10**(2), 1–25
DOI: 10.21314/JFMI.2022.004

Research Paper

Do DEXs work? Using Uniswap V2 to explore the effectiveness of decentralized exchanges

Yuen C. Lo and Francesca Medda

Institute of Finance and Technology, Chadwick Building, University College London, Gower Street, London WC1 E6BT, UK; email: y.lo@ucl.ac.uk

(Received June 23, 2021; revised May 20, 2022; accepted May 23, 2022)

ABSTRACT

This paper investigates the function of a blockchain-based decentralized exchange, specifically the effectiveness of the Ether–Tether liquidity pool on the Uniswap V2. We note that cointegration between the price set by the liquidity pool and its price elsewhere is a necessary condition of effectiveness. The trading price offered by the liquidity pool is determined by the token reserves ratio. We apply autoregressive distributed lag, vector autoregressive and vector error correction model methodologies to 154 days of hourly data for the Ether–Tether trading pair. We find that liquidity providers and arbitrageurs ensure that the ratio of reserves is cointegrated with the trading pair price elsewhere, and therefore that Uniswap can be an effective financial market. This raises the possibility that such decentralized exchanges could be used to improve the completeness of financial markets.

Keywords: Uniswap; decentralized exchange (DEX); blockchain; Ethereum; tokenomics.

1 INTRODUCTION

This paper focuses on a growing application of blockchain: the decentralized exchange (DEX). On May 17, 2021, a record-breaking US$1.7 billion worth of

Corresponding author: Y. C. Lo

Print ISSN 2049-5404 | Online ISSN 2049-5412

digital tokens were traded on the Uniswap V2 DEX. These trades used almost US$9 billion of committed liquidity.[1] In the preceding year the platform's volumes at times exceeded those of the largest centralized cryptoasset exchange, Coinbase (Balakrishnan 2020).

Despite this progress, most cryptoasset trading takes place on centralized exchanges owned by a firm. DEXs have only recently gained a significant share of cryptoasset volumes relative to centralized exchanges. On the plus side, centralized exchanges offer consistent transaction costs, fast settlement and optimized user interfaces, while the downside of such exchanges is the regular hacks (and collapses) that jeopardize the assets of which they have custody. Gandal *et al* (2018) examine the fall of the Mt. Gox bitcoin exchange as well as the increasing price manipulation leading up to the actual event. The irony of this is that the record-keeping functionality of blockchains makes them natural payment- and token-transfer mechanisms. Blockchains such as Bitcoin are payment systems (Huberman *et al* 2019).

Lin (2019) identifies four dimensions across which exchanges can be decentralized:

(1) the blockchain platform,

(2) the mechanism for discovering a counterparty,

(3) the order matching algorithm, and

(4) transaction settlement.

Choices regarding these functions impact an exchange's trade-off between performance, privacy and capital intensity. The Uniswap V2 exchange is decentralized across all four dimensions. Lin (2019) enumerates the benefits of DEXs as lower counterparty risk, potentially lower fees and more trading pairs. Trends favoring a switch toward DEXs include

(1) the increasing quantity of distinct cryptoassets,

(2) the regulatory risk of listing a cryptoasset on a centralized exchange, and

(3) user preferences to avoid Know Your Customer and anti money laundering (AML) regulations required by a centralized exchange.

Relating to the last point, centralized exchanges are a focus of regulatory actions, with the US Commodity Futures Trading Commission and Securities and Exchange Commission charging the derivatives platform Bitmex with providing US-based customers with access to unregulated financial derivatives and not following AML

[1] URL: v2.info.uniswap.org/home.

requirements (Chicago Futures Trading Commission 2020). In the United Kingdom, FCA policy banned the sale of derivatives that reference cryptoassets to retail investors (Financial Conduct Authority 2020). Importantly, the FCA has not banned the trading of cryptoassets. Uniswap and other DEXs are not yet offering derivatives, but it is clear that both regulation and cryptoasset infrastructure continue to evolve at speed. Alexander and Heck (2020) detail the problems arising from inconsistent regulation of cryptoasset markets. The increasing significance of DEXs will make financial regulation more difficult.

Research into DEXs connects to the literature on financial market infrastructure and microstructure. Lees (2012) provides an overview of conventional capital markets. All financial markets seek to optimize the social welfare of transactions by bringing multiple parties to a single exchange. That electronic exchanges can be distributed geographically is not new. Biais *et al* (2005) review the microstructure literature including transaction costs and bid–ask spreads. Both centralized exchanges and early DEXs use order books of bids and asks. The bid consists of prices and volumes at which participants are openly willing to buy. The ask consists of prices at which participants are willing to sell. If the same party engages in the bid and the ask at the same time, they are a specialist or market maker, looking to profit on the spread. Comerton-Forde *et al* (2010) find that market-maker balance sheet and income statement variables impact time variation in liquidity (in other words, spreads widen when specialist participants have large positions or lose money) and the benefits of market makers vanish during times of stress. However, an alternative to a bid–ask-based financial market is a disintermediated reserve-based model that holds pools of assets traders can access. Uniswap V2 is such a model.

Liquidity providers (LPs) commit proportionate quantities of two cryptoassets to form the basis of a trading pair. Figure 1 shows the Ether and Tether reserves for the Ether–Tether (ETHUSDT) pair. In return they receive 0.3% of the value of trades. Angeris and Chitra (2020) note how Uniswap applies a constant product rule to these reserves to map them to a marginal price. Further detail on these mechanics is provided in Section 2.2. We use an hourly data set of 154 days of cryptoasset reserves for the ETHUSDT pair from Uniswap, and explore the following research question: are DEXs, in particular Uniswap, effective cryptoasset exchanges? If DEXs make markets more complete, they do so in two ways. First, they replace nonlinear-liquidity-providing agents with continuous pricing curves. Second, prices are less influenced by agent profit and loss. We examine this question with three testable hypotheses.

(H1) The price of the ETHUSDT Uniswap pair matches its exchange rate off Uniswap.

In a centralized exchange, market makers and participants ensure varying degrees of the efficient market hypothesis (Fama 1970). Uniswap uses passive liquidity pools

FIGURE 1 Ether and Tether reserves for the ETHUSDT pair on Uniswap.

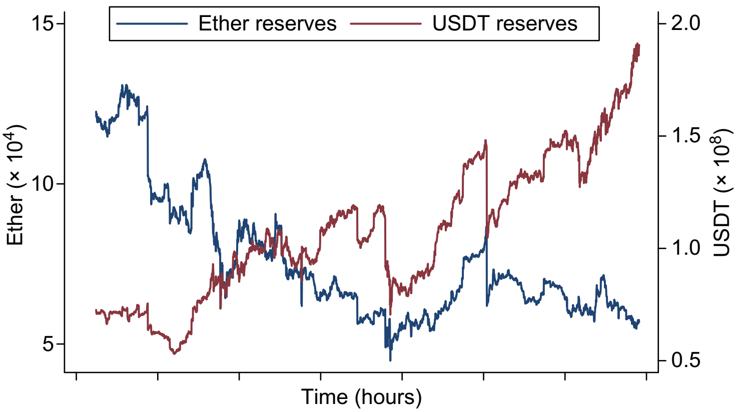

instead of active market makers, and therefore it is logical to test the connection between prices on and off Uniswap. Cointegration of the ratio of reserves to non-Uniswap pricing is a necessary, though not sufficient, condition of the effectiveness of Uniswap. It is where the pricing curve of Uniswap's constant product market maker equates to the price off-platform. A series of equilibrium correction autoregressive distributed lag (ARDL) models are formulated to test this hypothesis. We use a vector error correction model (VECM) as a robustness test.

(H2) The price of Ether, Bitcoin and the volume of transactions provide information that help predict changes in Uniswap reserves.

Here we examine which independent variables assist in predicting changes in reserve balances. Additionally, ARDL requires that there is at most one cointegrating relationship with the dependent variable, which testing this hypothesis can also check for.

(H3) Changes in one reserve balance in a pair cause changes in the other reserve balance.

ARDL does not prove causality. Therefore, we apply a vector autoregressive (VAR) model, and its test of Granger causality, to see if changes in one reserve balance of a pair influences the other reserve balance.

Our results contribute empirical evidence that liquidity pools on Uniswap V2 can be an effective cryptoasset exchange. They complement the work of Angeris *et al* (2020), who analyzes the mathematical implications of different constant function market-maker curves. Both our ARDL and VECM methodologies find in favor of the existence of a cointegrating vector between the derived ETHUSDT price on Uniswap and its price elsewhere. This cointegration is a necessary but not sufficient condition of effectiveness. We find a statistically significant relationship between the Ether and Tether reserves of the pool and the price of Bitcoin. This may indicate a connection between the liquidity pool and the wider cryptoasset space. Our VAR analysis suggests that, over the study period, changes in Tether reserves Granger cause changes in Ether reserves. This would be consistent with a specific type of arbitrage behavior that supports price cointegration.

The effectiveness of DEXs impacts both market completeness and cryptoasset regulation. Although blockchain promised the ability to digitally trade anything, in practice the liquidity may not have existed. Reserve-based markets imply that trades can now be carried out at any volume, making the financial markets more complete. Further, decentralized marketplaces will challenge the objectives and enforcement capabilities of regulators. In particular, as highlighted by Zetzsche *et al* (2020), decentralizing an institution eliminates the venture's need for a registered address and permanently located infrastructure, and therefore reduces the surface it exposes to the authorities. Section 2 provides a background to decentralized finance and Uniswap's pricing mechanism. Section 3 introduces the data, Section 4 details the methodology and Section 5 contains our results and discussion. Section 6 states our conclusions.

2 BACKGROUND

2.1 Blockchain, speculation and decentralized finance

Blockchain has become synonymous with digital tokens such as those traded on Uniswap. However, there is more to the technology than this. We highlight five functions.

The first function is as a mechanism to enable decentralized record-keeping, exemplified by Maersk and IBM's TradeLens project, which records the movement of 60% of the world's shipping containers (Jensen *et al* 2019). A record agreed by all is, by definition, accepted as "true". This reduces the need for trust, and at a minimum accelerates dispute resolution. In the future this may enable decentralized decision-making.

The second function is the smart contracts coded on the blockchain, which are commonly used to issue and manipulate third-party tokens. Shared code that all agree to be "true" can be thought of as shared rules. This may later open up new types of automation and agent relationships. Cong and He (2019) provide a formal proof of how a blockchain-based consensus, using smart contract-based prices contingent on delivery, can support new entrants. In their paper, new entrants signal quality by trustlessly guaranteeing buyers compensation if the product fails, making the contract space more complete. The shared computer code referred to as smart contracts does not come with guarantees. Rather, any consequences are public prior to interaction.

The third function is digital tokens. It is noted that both record-keeping and tokens can be used separately to enable payments and the transfer of value. However, it is with tokens that we enter the field of tokenomics and the ability of tokens to reduce project networking costs. Catalini and Gans (2016) implicitly divide these cost reductions into venture bootstrapping (where tokens are sold to investors or incentivize employees) and platform scaling (where tokens are offered to miners to process transactions or to evangelize users).

The fourth function is use of a distributed ledger as the payment infrastructure. There is limited need for a new electronic currency to substitute for bank deposits. However, there is demand for a novel payments infrastructure. The United States is a pivotal part of the international SWIFT payments system, which has been used to cut off Iran and sanction multinational companies (Majd 2018). Critically, a blockchain-based Chinese Central Bank Digital Currency (CBDC) would bootstrap a new payments system that can operate separately from existing infrastructures. Further, the Bank of England (2020) paper discusses the domestic resiliency benefit of a core payment network that sits outside the commercial banking system, but it only touches on why this facilitates features such as negative interest rates: a blockchain-based CBDC hands the payment system, user balances and its data to a single system owner. Kahn *et al* (2020) argue that distributed ledgers do not change the trade-offs of retail central bank accounts, but they do change the trade-offs of offering a token-based system.

The fifth function is, conversely, the ability to use decentralization to break rules and disrupt existing systems. The rise of blockchain tokens has facilitated online crime and money laundering. Foley *et al* (2019) use a variety of network analyses, such as transactions with known dark web wallets, to estimate that in 2019 one-quarter of Bitcoin users were involved with illegal activities (equating to US$76 billion in transactions). As Foley *et al* (2019, p. 1798) says, “cryptocurrencies are transforming … black markets by enabling black e-commerce”. However, the evolution and use of digital tokens suggest that illicit activities are not the primary use case of digital tokens. First, Brainard (2020) observes that the money-like use cases of means of exchange, store of value and unit of account have increasingly been taken over by stablecoins. Dwyer (2015) argues were never well addressed by Bitcoin. The US Financial Stability Board defines cryptoassets as “a type of private asset that depends primarily on cryptography and distributed ledger or similar technology as part of their perceived or inherent value” (Bank of England 2020, Footnote 6, p. 15), while Bank of England (2020, p. 15) defines stablecoins as a type of cryptoasset “whose value is linked to another asset”, ie, the US dollar. The most popular stablecoin is the Tether digital token (USDT). It makes up 5% of the value of all cryptoassets, compared with 60% for Bitcoin, but manages double the daily transaction value (Njuion 2020). Such stablecoins are unsuited to illicit activities, as they are typically centralized and easily frozen by their issuers.[2]

Despite the growth of cryptoassets for payments, arguably the leading use case for digital tokens is speculation. This is difficult to address empirically. Lo (2017) provides evidence that the price action of Bitcoin is consistent with it being traded as a proxy for the prototyping phase of a new technology. Ciaian *et al* (2017) use an ARDL methodology to find a variety of relationships between Bitcoin, altcoins and a set of macroeconomic variables. However, these intriguing papers reveal relatively little consistency or connection between any of these digital assets. Lo and Medda (2020) categorize and test a set of initial-coin-offering tokens, issued prior to 2017, by token function. They highlight the large quantity of funds directed to a set of ventures that consisted of little more than a white paper and a website. Although a number of these projects are still in operation, none has a noteworthy number of users. Other than Bitcoin, Ether and stablecoins, few cryptoassets have retained a significant share of the value of the space. Cumulatively, all this speaks to the speculative context of trading such vehicles. Arthur *et al* (2016) review the differences between gambling, speculation and investing. The key distinctions are the expected value (EV) and variability of returns. Speculation involves a higher EV than gambling (where a negative EV is the norm) and greater variability than investing. This is not to deride the importance of speculation. Both venture capital and oil drilling (especially prior to seismic surveys and shale drilling) observe a high number of project failures. Particularly in the crypto space, these flows of funds have been critical to the creation of decentralized building blocks, known as primitives.

Uniswap is one of the primitives of the wider space known as decentralized finance (DeFi). The fund manager Kyle Samani defines DeFi as "Enforcing financial contracts through code running on censorship resistant and permissionless public blockchain" (Samani 2020). Other large players in DeFi include Compound in the lending and borrowing of cryptoassets, and Synthetix in cryptoasset derivatives. The DeFi space has become popular for liquidity mining or yield farming, where ether, stablecoins and other assets are committed and rewarded. Part of these rewards are payments, such as Uniswap's 0.3% fee for liquidity providers, but most are tokens handed out by the venture for platform scaling. Yearn.finance is an example of how primitives are used as building blocks. Smart contracts manage deposits on its platform, minting assets on Synthetix and trading on DEXs as required, to maximize potential rewards. The emergence of DeFi has exacerbated congestion and operation costs (ie, gas fees) on the Ethereum network, similar to the situation on the Bitcoin

[2] URL: https://trustnodes.com/2020/09/26/tether-freezes-30-million-usdt-after-kucoin-hack.

FIGURE 2 Uniswap constant product automated market maker.

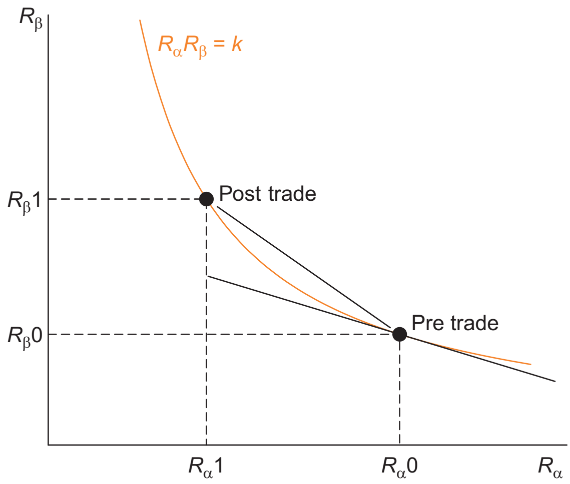

network in 2018. Early proof-of-work blockchain networks are capacity constrained by design (Lo and Medda 2018). Supporters might argue that this is how blockchains, such as Bitcoin, enable decentralization and censorship resistance. DeFi primitives are expanding the scope of both these functions.

2.2 Uniswap's constant product automated market maker

By construction, a constant product automated market maker (AMM) ensures that the reserves before and after the trade (assuming no fees) adhere to the function

$$k = R_\alpha R_\beta, \tag{2.1}$$

where k is a constant, R_α is the quantity of reserves of asset α, and R_β is the quantity of reserves of asset β. Equation (2.1) is plotted in Figure 2. If we differentiate both sides of $k = R_\alpha R_\beta = F(R_\alpha, R_\beta)$ to $0 = F_{R\alpha}\,\mathrm{d}R_\alpha + F_{R\beta}\,\mathrm{d}R_\beta$, we can rearrange this to show the price for any given ratio of reserves:

$$p_{\alpha\beta} = \frac{F_{R\alpha}}{F_{R\beta}} = -\frac{\mathrm{d}R_\beta}{\mathrm{d}R_\alpha}, \tag{2.2}$$

where $F_{R\alpha}$ is the partial derivative of the function F in terms of $R\alpha$. This price is only available where trades do not change the ratio of reserves (ie, he trade is small). Otherwise, the marginal price of a transaction is the relative change in quantity of the two reserves: $p'_{\alpha\beta} = -\Delta R_\beta / \Delta R_\alpha$. This is the slope of the line joining the before and after points on the curve. The slippage (difference between the realized price $-\Delta R_\beta/\Delta R_\alpha$ and the original price $-\mathrm{d}R_\beta/\mathrm{d}R_\alpha$) of a trade is positively correlated with trade size and inversely correlated to the size of reserves.

Angeris and Chitra (2020) generalize the mathematics of constant product market makers and argue that these market makers provide a tractable optimization problem for arbitrageurs to synchronize on- and off-chain prices. On a traditional exchange, the price of an asset lies between the bid and the ask, but this does not apply on DEXs. Market makers contribute to price discovery, but liquidity providers are price takers. LPs have no price protection other than the constant product function, which treats price as an output. Because arbitragers capture some of the value of price changes, the assets of an LP, excluding fees, will underperform a fixed portfolio of the original assets, unless the price reverts. This is deceptively referred to as “impermanent loss” yet, even if price reverts, LPs underperform a portfolio that actively rebalances. The CEO of Uniswap, Hayden Adams, has referred to LPs as “lo]ng fees/volatility and short volatility/fees” (Adams 2020). In other words, LPs benefit from fees that are a function of volatility but suffer from price change volatility. Separately, traders can specify a maximum deviation relative to an external price oracle, to protect themselves from short-term reserve fluctuations. Notably, large trades on Uniswap are vulnerable to front running, where bots watch Ethereum’s mempool of unprocessed trades, and buy and sell around market-moving transactions (Mierzwa 2020).

3 DATA

This study is based on closing hourly Uniswap data for the period from December 2, 2020 to May 5, 2021, via multiple queries of the Uniswap V2 subgraph.[3] Subgraphs are a way of storing public data and are accessible via Graph Query Language (GQL). The 3705 hours of data captured equate to 154 days. We note that, on May 5, Uniswap V3 (and its concentrated liquidity product) launched, so later data is not comparable. We acquire via an application programming interface the closing ETHUSDT and Bitcoin–Tether (BTCUSDT) prices from the Cryptocompare.com data aggregator, which is used by firms including Refinitiv and Quandl. The integration of the two data sets is based on the hourly Unix time stamps native to both. We do not know the exchange weights or methodology used by Cryptocompare’s benchmark exchange ETHUSDT rate. Descriptive statistics for a selection of data set variables are shown in Table 1. Total reserves for the pair in US dollars are plotted against trading volumes in Figure 3.

[3] URL: https://thegraph.com/explorer/subgraph/uniswap/uniswap-v2.

FIGURE 3 Total reserves and trading volumes for the ETHUSDT pair on Uniswap.

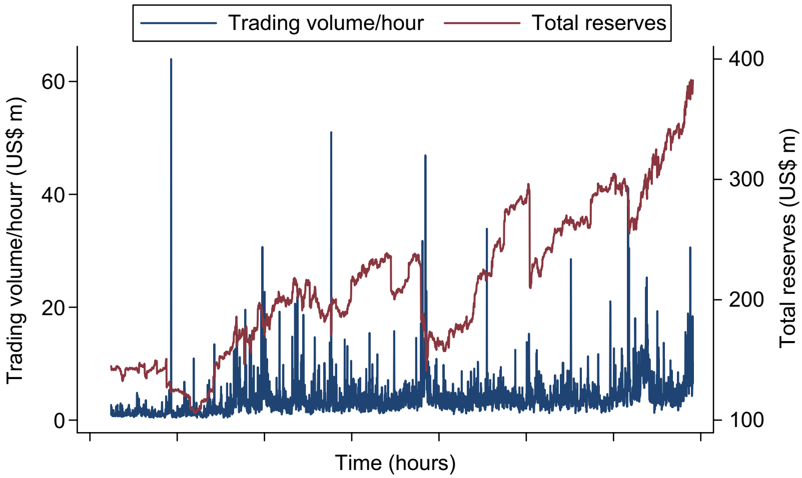

4 METHODOLOGY

Hypothesis (H1) requires us to test for cointegration between price and the ratio of reserves. This cointegration is central to the effective trading of cryptoassets on Uniswap, and can be thought of as a common stochastic trend. Within equilibrium correction ARDL, the test of cointegration is referred to as the bounds test. We proceed as follows:

(1) by categorizing the variables by their order of integration;

(2) by discussing the framework of the ARDL model; and

(3) by laying out the equilibrium correction ARDL to which the bounds test is applied.

Although Pesaran *et al* (2001) commented that ascertaining the order of integration was unnecessary prior to testing for cointegration under ARDL, this was asserted in a bounded fashion: the framework does not extend directly to variables that are integrated of order 2, ie, $I(2)$. Therefore, we test for unit roots using augmented Dickey–Fuller (ADF), Phillips–Perron (PP) and Dickey–Fuller generalized least squares (DFGLS) tests. We use the Akaike information criterion (AIC) to determine the appropriate number of lags.

The results shown in Table 2 indicate that our sample contains a mix of integration orders. The reserves, ratio of reserves and prices are stationary in the first differences

TABLE 1 Descriptive statistics: 154 day snapshot of Uniswap ETHUSDT pair.

	N	Mean	SD	Min	Median	Max
Ether reserves (tokens)	3705	75 754	19 562	44 800	69 271	130 929
USDT reserves (tokens)	3705	107 042 548	30 741 537	52 994 920	103 405 152	191 274 496
Total reserves (US$m)	3705	214	61.5	106	207	383
Ether hourly transaction volume	3705	2 630	3 824	573	2 010	194 929
USDT hourly transaction volume	3705	3 840 317	3 285 136	405 076	3 116 107	50 567 096
ETH reserves $\times$ USDT reserves	3705	7.76e+12	1.74e+12	3.16e+12	7.89e+12	1.26e+13
Ratio of reserves USDT to ETH	3705	1 542	631	538	1 629	3 473
ETHUSDT close price (US$)	3705	1 542	632	539	1 627	3 484
BTCUSDT close price (US$)	3705	43 210	13 800	17 649	47 455	64 568
Difference in log Ether reserves	3704	−0.000207	0.0131	−0.319	−0.000137	0.14
Difference in log USDT reserves	3704	0.00026	0.0127	−0.311	0.000495	0.142

SD, standard deviation.

TABLE 2 Stationarity test results.

	ADF		PP		DFGLS	
	Level	**1st diff.**	**Level**	**1st diff.**	**Level**	**1st diff.**
Ether reserves	NS	S	NS	S	NS	S
USDT reserves	NS	S	NS	S	NS	S
Ether volumes	S	S	S	S	S	S @ < 18 lags
USDT volumes	S	S	S	S	S	S @ < 21 lags
ETHUSDT price	NS	S	NS	S	NS	S
BTCUSDT price	NS	S	NS	S	NS	S
Ratio of reserves	NS	S	NS	S	NS	S

Three tests of stationarity applied to seven time series, on levels and first differences. NS, nonstationary at the 5% statistical significance level. S, stationary at the 5% statistical significance level.

($I(1)$), while volumes will likely be stationary in levels $I(0)$. The DFGLS test applies a GLS detrending on the series prior to running an ADF test, which can improve the power of the test (Elliott *et al* 1996). Although both OLS- and GLS-based tests see a declining power in the presence of level or trend breaks, the risk is in misidentifying a stationary time series with a structural break as nonstationary ie, that the order of integration is overestimated (Cook and Manning 2004). Therefore, ARDL is appropriate and can be represented by

$$y_t = c_0 + c_1 t + \sum_{i=1}^{p} \phi_i y_{t-i} + \sum_{j=0}^{q} \beta_j x_{t-j} + u_t, \tag{4.1}$$

where y_t is the dependent variable at time t, with up to p lags included in the model and x_t is the $k \times 1$ vector of independent variables, u_t is a random error term and c_0 and c_1 are the deterministic intercept and time trend coefficients. For simplicity we show here the lag order q as being the same for all the independent variables (this does not have to be the case).

An extension of the model in (4.1) estimates the long-run relationships as an equilibrium correction process (Pesaran *et al* 2001). It frames the independent variables as long-run forcing of the dependent variable (Kripfganz and Schneider 2020). This assumes the independent variables are weakly exogenous, and models should consider the directionality of effects during formulation (eg, it may be plausible for transactions to drive changes in reserves, but it is less likely that reserves force transactions). With respect to hypothesis (H1), y_t becomes the ratio of reserves R_t, while

x_t are the exchange rates of ETH and BTC with Tether, ie,

$$\begin{aligned}\Delta R_t &= c_0 + c_1 t + \alpha(R_{t-1} - \theta_1 \mathrm{ETHUSDT}_{t-1} - \theta_2 \mathrm{BTCUSDT}_{t-1}) \\ &\quad + \sum_{i=1}^{p-1} \varphi_{Ri} \Delta R_{t-i} + \omega_1 \Delta\mathrm{ETHUSDT}_t + \omega_2 \Delta\mathrm{BTCUSDT}_t \\ &\quad + \sum_{j=1}^{q-1} \varphi_{\mathrm{ETH}j} \Delta\mathrm{ETHUSDT}_{t-j} + \sum_{k=1}^{r-1} \varphi_{\mathrm{BTC}k} \Delta\mathrm{BTCUSDT}_{t-k} + u_t,\end{aligned} \tag{4.2}$$

where α is the adjustment coefficient, θ are the long-run coefficients on the first lags of $\mathrm{ETHUSDT}_t$ and $\mathrm{BTCUSDT}_t$, ω are the short-run coefficients on the first differences of $\mathrm{ETHUSDT}_t$ and $\mathrm{BTCUSDT}_t$, and φ are the short-run coefficients on the lagged differences of R_t, $\mathrm{ETHUSDT}_t$ and $\mathrm{BTCUSDT}_t$.

This choice of methodology benefits from its ability to estimate both short-run and long-run parameters at the same time. Further, Pesaran and Shin (1999) observe that an appropriate estimation of the orders of the extended $\mathrm{ARDL}(p, m)$ model is sufficient to correct for both the residual serial correlation and the problem of endogenous regressors. The ARDL models and coefficients are estimated in STATA using the `ARDL` package, which is based on Kripfganz and Schneider (2020). These models are subjected to two parts of the ARDL bounds test. Note that if there is no cointegration, then the ARDL model in (4.1) is used to estimate relationships between variables and their lags. Hypothesis (H1) is investigated via a variety of specifications that look for cointegration between the ratio of Ether to USDT reserves and the exchange rate of ETHUSDT. Hypothesis (H2) uses the same methodology and searches for the presence of cointegrating and auto regressive relationships between reserves, transactions and price.

Cointegration implies that there are stationary equilibrium relationships between separate nonstationary variables. A corollary of this is that, when these variables diverge, at least one of the cointegrated variables reconverges to return the system to a long-run equilibrium. In (4.2) the rate of this is estimated by the coefficient α. The bounds test begins with a Wald test (F-statistic) of the joint hypothesis

(H_0^F) $\alpha = 0$ and $\sum_{i=0}^{q} \varphi_{xi} = 0$,

versus the alternative hypothesis

(H_1^F) $\alpha \neq 0$ and $\sum_{i=0}^{q} \varphi_{xi} \neq 0$.

If the null hypothesis is rejected, then the t-statistic is used to test

(H_0^t) $\alpha = 0$

versus

$(\mathrm{H}_1^t)\ \ \alpha \neq 0.$

The distributions of these test statistics are nonstandard and depend on the integration order of the independent variables. Kripfganz and Schneider (2020) extend the set of available critical values for the bounds test by estimating response surface models, with each significance level showing four critical values based on $I(0)$ and $I(1)$ for the F- and t-tests. There can be at most one cointegrating relationship between the independent variables and the dependent variable (although there may be additional cointegrating relationships between the independent variables). The validity of the bounds test depends on normally distributed error terms that are homoscedastic and serially uncorrelated. For the equilibrium correction ARDL model for the ratio of ETH/USDT reserves to ETHUSDT price, we carry out the Breusch–Godfrey Lagrange multiplier (LM) test for autocorrelation and the Breusch–Pagan test for heteroscedasticity. Kripfganz and Schneider (2020) note that bounds testing with a higher lag order can be useful for addressing the remaining serial error correlation, with a more parsimonious model applied after testing for forecasting purposes. Across our analysis, AIC, which indicates the optimality of a model, is used to select the set of variables and the number of lags. AIC is less parsimonious than Schwarz's Bayesian information criterion (BIC), but in ARDL lowers the risk of serial correlation.

Our study uses a VECM as a robustness check of our hypothesis (H1). VECMs are an extension of the VAR model we use to test for Granger causality as part of hypothesis (H3). We explain how VAR models address directional changes in cryptoasset reserves before moving on to discussing VECM. VAR modeling specifies as many models as there are dependent variables (Enders 1995). We use the first difference of logs to ensure the linearity of changes in the two rapidly increasing reserve balances. In its basic form of two variables with a single lag, the VAR model defines the following:

$$\begin{aligned}\Delta(\ln \mathrm{ETH}_t) &= \alpha_u + \beta_{u1}\Delta(\ln \mathrm{USDT}_{t-1}) + \varepsilon_u,\\ \Delta(\ln \mathrm{USDT}_t) &= \alpha_e + \beta_{e1}\Delta(\ln \mathrm{ETH}_{t-1}) + \varepsilon_e.\end{aligned}$$

The variables are considered to be endogenous. Although it is possible to use lags selectively, typically each model repeats the same lagged explanatory variables symmetrically. The Granger causality tests within the VAR model examine whether the prior-period first difference of the log of one cryptoasset reserve provides information about the value of current-period first difference of the log of the other cryptoasset reserve. Tests of Granger causality exploit the directionality of time to imply the directionality of the relationship. Changes in reserve balances are a corollary of trades on the Uniswap platform and, following such trades, the mechanism by which arbitrageurs cointegrate the reserve ratio and price.

VAR models require stationary time series. Earlier, we used the first difference of the logs of the original $I(1)$ time series to ensure this. VECMs add back in some of the information of the undifferenced time series. First, they estimate the long-run equilibrium using ordinary least squares. Note that the VAR model is applied to changes in reserves, but hypothesis (H1) and this VECM relate to the ratio of reserves and the ETHUSDT price. If they are cointegrated, the residuals are stationary and the estimators are super consistent (Enders 1995):

$$R_t = \alpha + \beta' \,\mathrm{ETHUSDT}_t + \varepsilon. \tag{4.3}$$

The differences between actual observations and modeled observations are then included in the VECM. These residuals are the deviation from the long-run equilibrium. One form of this, with one lag and no deterministic trend, is the following:

$$\Delta R_t = \alpha + \lambda(R_{t-1} - \beta' \,\mathrm{ETHUSDT}_{t-1}) + \beta_1 \Delta\mathrm{ETHUSDT}_{t-1} + \varepsilon_1. \tag{4.4}$$

Note that, in multivariate notation, a cointegration matrix Π is typically used to represent the potentially complex nature of the cointegrating relationship, whereas here it is written out explicitly. The error correction term, λ, estimates how changes in R_t vary when one of the variables deviates from the common stochastic trend. As with VAR modeling, VECM is symmetric, and $\Delta\mathrm{ETHUSDT}_t$ is also estimated as a function of R_t. In the next section we examine the results.

5 RESULTS AND DISCUSSION

The results of applying ARDL to our dependent variable, the ratio of Ether to USDT reserves, with the exchange rate of Ether and the exchange rate of Bitcoin (both priced in USDT) are shown in Table 3 ([A] and [B] are two alternative specifications of the model). As all three variables in this model are $I(1)$, the bounds test statistics are compared with the $I(1)$ critical values. The F-statistic and the t-statistic are more extreme than the related critical values (p-value $=$ 0.000), which reject the null hypothesis of no level relationship. This provides evidence in favor of (H1), ie, the price of the ETHUSDT Uniswap pair matches its exchange rate off Uniswap.

This result confirms empirically the effectiveness of Uniswap's reserve-balance-based Ether and USDT exchange pair on an hourly time frame. These results are supported graphically by Figure 4. Part (b) of this figure indicates that some of the arbitrage opportunity is visible in the data but exceeds 1% only five times over the sample period. We note that, because of fees, arbitrage is unlikely to take place when the difference between on- and off-Uniswap prices is less than 0.3%.

TABLE 3 ARDL ratio of reserves to ETHUSDT price.

	[A]	**[B]**
Adjustment factor		
L (difference from equilibrium)	−0.900***	-0.904***
(a) Long-run effects		
L (ETHUSDT price)	1.000***	1.000***
L (BTCUSDT price)		0.000
(b) Short-run effects		
LD (ratio of reserves)	−0.063***	−0.063***
D (ETHUSDT price)	0.951***	0.937***
LD (ETHUSDT price)	0.069***	0.063***
D (BTCUSDT price)		0.001***
LD (BTCUSDT price)		0.000
AIC	20 387.975	20 376.652
BIC	20 425.273	20 432.599
N	3701	3701
(c) Bounds test results		
F-statistic	798.271	536.459
t-statistic	−39.956	−40.116
F-test p-value $I(1)$	0.000	0.000
t-test p-value $I(1)$	0.000	0.000

L, lagged variable. LD, lagged first difference of the variable. D, first difference of the variable. The bounds test rejects (H_0), ie, that there is no level relationship at the 5% significance level. *, ** and *** denote $p < 0.05$, $p < 0.01$ and $p < 0.001$, respectively.

Returning to Table 3, during the study time period the adjustment factor α is 0.9. This suggests that 90% of the difference between the ratio of reserves and the ETHUSDT price is adjusted back to long-run equilibrium over the course of the subsequent hour. The long-run effects are the coefficients θ of the lagged exchange rates of ETHUSDT and BTCUSDT. In both specifications, the coefficient of the lagged ETHUSDT price is 1. Of the long-run coefficients, only ETHUSDT is statistically significant. The short-run effects are φ and ω from (4.2), which are the coefficients on the first and lagged differences of our variables. All of the short-run effects are statistically significant except for the lagged difference of BTCUSDT. The lower AIC value and the statistical significance of the first difference of BTCUSDT suggest the

FIGURE 4 The ratio of Ether and Tether reserves (on the ETHUSDT pair on Uniswap) versus the ETHUSDT price.

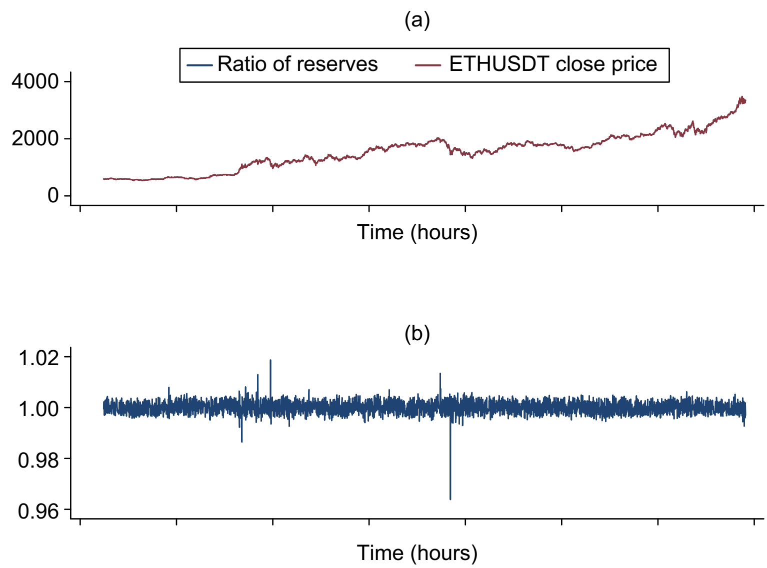

(a) Ratio of USDT/ETH reserves and ETHUSDT price. (b) Reserve ratio deviation (ETHUSDT price/ratio of USDT/ETH reserves).

Bitcoin price does contain information on predicting changes in the ratio of reserves. This may be because of Bitcoin's importance in the cryptoasset space, its impact on trader wealth or some residual use as a unit of account. We run a Breusch–Godfrey LM test for autocorrelation, which does not reject the null hypothesis of no serial correlation for 1–10 lags at the 5% significance level. The Breusch–Pagan test for heteroscedasticity has a χ^2 test statistic of 0.24 and a p-value of 0.6269. Therefore, we do not reject the null hypothesis of constant variance at the 5% significance level.

As a robustness check, we execute a VECM to complement our ARDL model. It is an alternative way to examine our two time series, the ratio of reserves between Ether and USDT, and the ETHUSDT price. As required, both are integrated of order 1. The first differences are taken and regressed on zero or one lagged difference, as suggested by selection order information criterions. The error correction coefficient is the critical output and indicates whether and how the two time series converge. The results in Table 4 ([C] and [D] are two alternative specifications of the model) indicate that the reserve ratio moves toward the model equilibrium, in both specifications, at the 99.9% statistical significance level. We do not find evidence that the ETHUSDT price moves toward the ratio of reserves. This supports the case that the

TABLE 4 Robustness check: vector error correction model.

	[C]	[D]
D (Ratio of reserves)		
L (Error correction coefficient)	−0.893***	−0.853***
LD (Ratio of reserves)		−0.023
LD (ETHUSDT price)		0.038
D (ETHUSDT price)		
L (Error correction coefficient)	0.079	0.049
LD (Ratio of reserves)		0.043
LD (ETHUSDT price)		−0.033
AIC	52 822.888	52 798.676
BIC	52 847.757	52 848.411
N	3704	3703

L, lagged variable. LD, lagged first difference of the variable. D, first difference of the variable. Models are ordered by descending AIC. *, ** and *** denote $p < 0.05$, $p < 0.01$ and $p < 0.001$, respectively.

two time series are cointegrated using a second methodology and offers evidence that Uniswap pricing moves to match the price elsewhere.

We note that a finding of cointegration is a necessary but not sufficient condition for the effectiveness of Uniswap and its automated market maker. If they are not cointegrated, then one of these prices is wrong for a prolonged period, and an arbitrage opportunity for risk-free profits would be sustained. Drilling further down into the efficiency of the ETHUSDT pair is a direction for future research as more data becomes available. Additionally, analyzing the effectiveness and efficiency of other markets on Uniswap is a open problem. The issue of the 0.3% trading fee is universal. But the plethora of rarely traded token pairs on Uniswap results in variations in the available data. This paper focuses on a token pair, where off-DEX pricing is liquid and high frequency. Yet, this is not the case for many token pairs, and we highlight the difficulty in the empirical analysis of illiquid markets that may exist solely because of an LP-based platform such as Uniswap (eg, where there is no off-Uniswap benchmark price). However, we observe that this is an opportunity as well as a constraint. Anecdotally, it is now possible to observe changes in liquidity as prices change, which opens up a largely unexplored space for empirical researchers.

In order to explore hypothesis (H2), we put the ratio of reserves to one side and run ARDL models with Ether reserves and USDT reserves as our dependent variables. The bounds tests on these equilibrium correction models (not shown) do not reject the null hypothesis of no level relationship; we find no evidence of cointegration. Because of this, the equilibrium correction models are not appropriate, and the results of the standard ARDL model are presented in Tables 5 and 6. For both dependent

TABLE 5 Short-run ARDL model of Ether reserves within the ETHUSDT Uniswap pair.

	[E]	[F]	[G]
L (ETH reserves)	0.903***	0.903***	0.899***
L2 (ETH reserves)	0.095***	0.095***	0.099***
(USDT reserves)	0.001***	0.001***	0.001***
L (USDT reserves)	−0.001***	−0.001***	−0.001***
L2 (USDT reserves)	−0.000***	−0.000***	−0.000***
L3 (USDT reserves)	0.000*	0.000*	0.000*
(ETH price)	−38.534***	−38.519***	−35.147***
L (ETH price)	32.288***	32.243***	28.683***
L2 (ETH price)	6.746***	6.663***	6.420***
L3 (ETH price)	−1.266*	−1.208*	
L4 (ETH price)	0.731	0.790	
(ETH volume)		−0.001	−0.000
(USDT volume)		0.000	0.000
L (USDT volume)		−0.000*	−0.000*
(BTCUSDT price)			−0.203***
L (BTCUSDT price)			0.201***
AIC	56 105.509	56 106.529	56 058.911
BIC	56 180.105	56 199.775	56 152.157
N	3701	3701	3701

L, lagged variable. LD, lagged first difference of the variable. D, first difference of the variable. Models are ordered by descending AIC. *, ** and *** denote $p < 0.05$, $p < 0.01$ and $p < 0.001$, respectively.

variables, we execute three models ([E], [F], [G], and [H], [I], [J], respectively) with specifications that go from specific to general. The lower the AIC, the more appropriate the specification of the model. For both Ether reserves and USDT reserves the most general models with the most variables appear to be preferred in predicting changes in the dependent variables. The suggestion that the price of Ether impacts reserves makes sense, as reserves are a function of both liquidity provision in a ratio set by price and trades that exchange one reserve for another at a price dependent on impact. The statistical significance on volumes is somewhat weaker. Notably, the statistical significance of Bitcoin is unexpected. Together, these results find in favor of hypothesis (H2). We test the other variables to ensure there are no additional cointegrating relationships that may impact our earlier analysis. Mostly there is no logic for such directionality, and we do not find such evidence. Over the study time period we also do not find cointegration between the price of Ether and the price of Bitcoin (not shown). The result of this may be different over longer time periods.

TABLE 6 Short-run ARDL model of USDT reserves within the ETHUSDT Uniswap pair.

	[H]	[I]	[J]
L (USDT reserves)	0.862***	0.865***	0.862***
L2 (USDT reserves)	0.160***	0.161***	0.163***
L3 (USDT reserves)	−0.025**	−0.029***	−0.029***
(ETH reserves)	1 138.738***	1 135.284***	1 140.416***
L (ETH reserves)	−983.720***	−982.353***	−984.978***
L2 (ETH reserves)	−153.029***	−150.898***	−152.422***
(ETH price)	52 473.761***	52 261.263***	48 944.982***
L (ETH price)	−42 700.000***	−42 600.000***	−39 100.000***
L2 (ETH price)	−9 620.627***	−9 504.396***	−9 637.803***
(ETH volume)		0.947	0.899
(USDT volume)		−0.011*	−0.010*
L (USDT volume)		0.016***	0.016***
L2 (USDT volume)		−0.004	−0.003
L3 (USDT volume)		0.008	0.008
L4 (USDT volume)		−0.009*	−0.009*
(BTCUSDT price)			208.153***
L(BTCUSDT price)			−206.205***
AIC	1.09e+05	1.09e+05	1.09e+05
BIC	1.10e+05	1.10e+05	1.10e+05
N	3701	3701	3701

L, lagged variable. LD, lagged first difference of the variable. D, first difference of the variable. Models are ordered by descending AIC. *, ** and *** denote $p < 0.05$, $p < 0.01$ and $p < 0.001$, respectively.

Hypothesis (H3) examines how the Uniswap ETHUSDT reserves return to equilibrium. We investigate this with a VAR model. We begin by reviewing the order selection statistics for our two variables. The lag order selection information criterions suggest one and four lags. We run two models, the first with one lag [L] and the second with four lags [K]. The results of this modeling are shown in Table 7. Tests of model stability suggest that the eigenvalues are appropriately within the unit circle.

When the dependent variable is the first difference in the log of Ether reserves, the lagged first difference in the log of USDT reserves is statistically significant under both specifications. Although the four-lag model identifies a number of other statistically significant autoregressive relationships, the AIC and BIC are very slightly higher, so do not appear to boost predictiveness.

At the 5% statistical significance level, we reject the null hypothesis that the first differences of the log of the USDT reserves do not Granger-cause changes in the first differences in the log of the Ether reserves. For one lag the χ^2 test statistic is 5.14

TABLE 7 VAR model of Ether and USDT reserves.

	[K]	[L]
(a) First difference of log Ether reserves		
LD (log ETH reserves)	−0.008	−0.013
L2D (log ETH reserves)	−0.018	
L3D (log ETH reserves)	0.033	
L4D (log ETH reserves)	−0.079***	
LD (log USDT reserves)	−0.045*	−0.047*
L2D (log USDT reserves)	0.066**	
L3D (log USDT reserves)	−0.077***	
L4D (log USDT reserves)	0.019	
(b) First difference of log USDT reserves		
LD (log ETH reserves)	−0.026	−0.030
L2D (log ETH reserves)	−0.004	
L3D (log ETH reserves)	0.043*	
L4D (log ETH reserves)	−0.002	
LD (log USDT reserves)	−0.023	−0.022
L2D (log USDT reserves)	0.044*	
L3D (log USDT reserves)	−0.090***	
L4D (log USDT reserves)	−0.045*	
AIC	−4.50e+04	−4.50e+04
BIC	−4.49e+04	−4.49e+04
N	3700	3703

LD, lagged difference of the variable. L2D, lagged second difference of the variable. L3D, lagged third difference of the variable. L4D, lagged fourth difference of the variable. Models are ordered by descending AIC. *, ** and *** denote $p < 0.05$, $p < 0.01$ and $p < 0.001$, respectively.

with a p-value of 0.023. For four lags the χ^2 test statistic is 29.24 with a p-value of 0.00. However, we do not reject the null hypothesis that the first differences of the log of the Ether reserves do not Granger-cause changes in the first differences in the log of the USDT reserves ($p = 0.125$ and $p = 0.154$ for one and four lags, respectively). Overall, we find evidence in favor of (H3): changes in one reserve balance (USDT) of a pair Granger causes changes in the other reserve balance (Ether). It is hard to explain definitively why this would be the case. However, we can make inferences, because on Uniswap every trade has a price impact. Ceteris paribus, arbitrage trades following off-Uniswap price changes should not impact the next period. Only arbitrage trades following trading-induced reserve changes should link two time periods.

Arguably, arbitrage should lead to bidirectional Granger causality. As this is not the case, it may simply be that nonarbitrage trades tend to be purchases of Ether. Because of the nature of the automated market maker, in this case, there are larger changes in the USDT balance. In other words, our Granger causality results are consistent with a reserve ratio at equilibrium impacted by a first trade buying Ether, which pushes USDT reserves out of balance. Afterward, an arbitrage trade sells Ether (buys USDT) to bring the reserve ratio back into equilibrium with benchmark pricing. This sequence sees a change in USDT reserves leading a change in Ether reserves.

Bringing together the various findings, the error correction ARDL and VECM results support our case that ETHUSDT prices on and off Uniswap V2 are cointegrated. The VECM results suggest that the on-Uniswap reserve ratio and price move toward the off-Uniswap price, hinting that price discovery for ETHUSDT occurs on centralized exchanges. Hasbrouck's information share measure (Hasbrouck 1995) would be a suitable method for analyzing this further. The VAR results delve further into the equilibrium process, showing that changes in the USDT reserves Granger-cause changes in the Ether reserve balances.

6 CONCLUSION

This research provides empirical evidence regarding the effectiveness of reserve-based asset exchanges. We find that, for the sample period, the ratio of Ether and USDT reserves to the ETHUSDT pair is cointegrated with a third-party ETHUSDT exchange rate benchmark. For a constant product automated market maker, this cointegration is a necessary condition of the exchange rate on-platform approximating the exchange rate off-platform. The success of Uniswap is a rare example of a financial market operating without the classical features of bids and asks, market makers or auctioneers. It is a clarion call to regulators, governments and financial market participants that the innovation and decentralization promised by blockchain-based systems are starting to gain traction. It is easy to discount the long-term impact of new highly speculative trading instruments but less easy to deride new financial infrastructure that improves the completeness of markets. DEX structures may be able to complement traditional bid–ask based capital markets. An argument made by Lo and Medda (2020) is that blockchain does not build strictly superior systems, but rather builds alternative systems that are attractive along uncommon dimensions (eg, no single point of control (political decentralization) and censorship resistance). Yet, more complete markets would constitute a quantitative benefit of blockchain. Further, DEXs have important implications for regulation, as decentralized exchanges do not require a legal form or fixed geographical infrastructure. This raises the question of how regulators and governments should respond to a marketplace that does not need a registered address and geographically fixed physical infrastructure. To date, rule makers have focused on regulating the institutions of the emerging cryptoasset space (Blandin *et al* 2019). This may no longer be possible.

Directions for future research include the potential to add an uncorrelated LP asset to investor portfolios, to test whether decentralized exchanges are more or less risky than centralized exchanges, and to examine whether decentralized exchanges can exist without centralized exchanges providing price discovery.

DECLARATION OF INTEREST

The authors report no conflicts of interest. The authors alone are responsible for the content and writing of the paper.

REFERENCES

Adams, H. (2020). None of the above! You're long fees/volatility and short volatility/fees. Twitter (@haydenzadams replying to @zhusuApp), September 24, 18:04. URL: https://twitter.com/haydenzadams/status/1309176877869826048?s=20.

Alexander, C., and Heck, D. F. (2020). Price discovery in Bitcoin: the impact of unregulated markets. *Journal of Financial Stability* **50**, 1–18 (https://doi.org/10.1016/j.jfs.2020.100776).

Angeris, G., and Chitra, T. (2020). Improved price oracles: constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pp. 80–91. ACM Press, New York (https://doi.org/10.1145/3419614.3423251).

Angeris, G., Evans, A., and Chitra, T. (2020). When does the tail wag the dog? Curvature and market making. Preprint (arXiv:2012.08040 [q-fin.TR]) (https://doi.org/10.48550/arXiv.2012.08040).

Arthur, J. N., Williams, R. J., and Delfabbro, P. H. (2016). The conceptual and empirical relationship between gambling, investing, and speculation. *Journal of Behavioral Addictions* **5**(4), 580–591 (https://doi.org/10.1556/2006.5.2016.084).

Balakrishnan, A. (2020). Uniswap's daily volume just overtook Coinbase by more than $80 million. Blog Post, August 31, CryptoBriefing.com. URL: https://cryptobriefing.com/uniswaps-daily-volume-overtook-coinbase-more-80-million/.

Bank of England (2020). Central bank digital currency: opportunities, challenges and design. Discussion Paper, Bank of England. URL: https://bit.ly/3K2HYbX.

Biais, B., Glosten, L., and Spatt, C. (2005). Market microstructure: a survey of microfoundations, empirical results, and policy implications. *Journal of Financial Markets* **8**(2), 217–264 (https://doi.org/10.1016/j.finmar.2004.11.001).

Blandin, A., Cloots, A. S., Hussain, H., Rauchs, M., Saleuddin, R., Allen, J. G., Cloud, K., and Zhang, B. Z. (2019). Global cryptoasset regulatory landscape study. Research Paper 23/2019, University of Cambridge Faculty of Law (https://doi.org/10.2139/ssrn.3379219).

Brainard, L. (2020). The digitization of payments and currency: some issues for consideration. Speech, Symposium on the Future of Payments, Stanford, CA. URL: www.bis.org/review/r200205j.pdf.

Catalini, C., and Gans, J. S. (2016). Some simple economics of the blockchain. Working Paper 22952, NBER (https://doi.org/10.3386/w22952).

Chicago Futures Trading Commission (2020). Complaint for injunctive and other equitable relief and civil monetary penalties under the commodity exchange act and commission regulations. Civil Action 20-cv-8132, Chicago Futures Trading Commission.

Ciaian, P., Rajcaniova, M., and Kancs, d'A. (2017). Virtual relationships: short- and long-run evidence from Bitcoin and altcoin markets. *Journal of International Financial Markets, Institutions and Money* **52**, 173–195 (https://doi.org/10.1016/j.intfin.2017.11.001).

Comerton-Forde, C., Hendershott, T., Jones, C. M., Moulton, P. C., and Seasholes, M. S. (2010). Time variation in liquidity: the role of market-maker inventories and revenues. *Journal of Finance* **65**(1), 295–331 (https://doi.org/10.1111/j.1540-6261.2009.01530.x).

Cong, L. W., and He, Z. (2019). Blockchain disruption and smart contracts. *Review of Financial Studies* **32**(5), 1754–1797 (https://doi.org/10.3386/w24399).

Cook, S., and Manning, N. (2004). The disappointing properties of GLS-based unit root tests in the presence of structural breaks. *Communications in Statistics: Simulation and Computation* **33**(3), 585–596 (https://doi.org/10.1081/SAC-200033390).

Dwyer, G. P. (2015). The economics of Bitcoin and similar private digital currencies. *Journal of Financial Stability* **17**, 81–91 (https://doi.org/10.1016/j.jfs.2014.11.006).

Elliott, G., Stock, J., and Rothenberg, T. (1996). Efficient tests for an autoregressive unit root. *Econometrica* **64**, 813–36 (https://doi.org/10.2307/2171846).

Enders, W. (1995). *Applied Econometric Time Series*. Wiley (https://doi.org/10.2307/2291367).

Fama, E. F. (1970). Efficient capital markets: a review of theory and empirical work. *Journal of Finance* **25**(2), 383–417 (https://doi.org/10.2307/2325486).

Financial Conduct Authority (2020). Prohibiting the sale to retail clients of investment products that reference cryptoassets. Policy Statement PS20/10, FCA, London.

Foley, S., Karlsen, J. R., and Putnins, T. J. (2019). Sex, drugs, and Bitcoin: how much illegal activity is financed through cryptocurrencies. *Review of Financial Studies* **32**(5), 1798–1853 (https://doi.org/10.1093/rfs/hhz015).

Gandal, N., Hamrick, J., Moore, T., and Oberman, T. (2018). Price manipulation in the Bitcoin ecosystem. *Journal of Monetary Economics* **95**, 86–96 (https://doi.org/10.1016/j.jmoneco.2017.12.004).

Hasbrouck, J. (1995). One security, many markets: determining the contributions to price discovery. *Journal of Finance* **50**(4), 1175–1199 (https://doi.org/10.1111/j.1540-6261.1995.tb04054.x).

Huberman, G., Leshno, J. D., and Moallemi, C. (2019). An economist's perspective on the Bitcoin payment system. In *AEA Papers and Proceedings*, Volume 109, pp. 93–96. American Economic Association, Nashville, TN (https://doi.org/10.1257/pandp.20191019).

Jensen, T., Hedman, J., and Henningsson, S. (2019). How TradeLens delivers business value with blockchain technology. *Management Information Systems Technology Quarterly* **18**(4), 221–243 (https://doi.org/10.17705/2msqe.00018).

Kahn, C. M., Rivadeneyra, F., and Wong, T.-N. (2020). Should the central bank issue e-money? *The Journal of Financial Market Infrastructures* **8**(4), 1–22 (https://doi.org/10.21314/jfmi.2019.117).

Kripfganz, S., and Schneider, D. C. (2020). Response surface regressions for critical value bounds and approximate p-values in equilibrium correction models. *Oxford Bulletin of Economics and Statistics* **82**, 1456–1481 (https://doi.org/10.1111/obes.12377).

Lees, F. A. (2012). *Financial Exchanges: A Comparative Approach*. Routledge, London (https://doi.org/10.4324/9780203805893).

Lin, L. X. (2019). Deconstructing decentralized exchanges. *Stanford Journal of Blockchain Law and Policy* **2**(1), 58–77. URL: https://stanford-jblp.pubpub.org/pub/deconstructing-dex/release/1.

Lo, Y. C. (2017). Blockchain and Bitcoin: technological breakthrough or the latest tulip price bubble? Working Paper, Social Science Research Network (https://doi.org/10.2139/ssrn.3198530).

Lo, Y. C., and Medda, F. (2018). Bitcoin mining: converting computing power into cash flow. *Applied Economics Letters* **26**(14), 1171–1176 (https://doi.org/10.1080/13504851.2018.1540841).

Lo, Y. C., and Medda, F. (2020). Assets on the blockchain: an empirical study of tokenomics. *Information Economics and Policy* **53**, Paper 100881 (https://doi.org/10.1016/j.infoecopol.2020.100881).

Majd, M. (2018). The cost of a SWIFT kick: estimating the cost of financial sanctions on Iran. In *The Political Economy of International Finance in an Age of Inequality: Soft Currencies, Hard Landings*, Epstein, G. A. (ed), Chapter 9, pp. 175–193. Edward Elgar (https://doi.org/10.4337/9781788972635.00017).

Mierzwa, T. (2020). How to munch on pickles from a whale dinner. Article, September 25, Token Flow Insights, Medium.com. URL: https://medium.com/token-flow-insights/how-to-munch-on-pickles-from-a-whale-dinner-edb5628cc769.

Njuion, J. P. (2020). Tether's (USDT) daily trade vol. eclipses BTC's, MarketCap hits $13b. Blog Post, August 27, Ethereum World News. URL: https://en.ethereumworldnews.com/tethers-usdt-daily-trade-vol-eclipses-btcs-marketcap-hits-13b/.

Pesaran, M. H., and Shin, Y. (1999). An autoregressive distributed lag modelling approach to cointegration analysis. In *Econometrics and Economic Theory in the 20th Century: The Ragnar Frisch Centennial Symposium*, Strøm, S. (ed), pp. 371–413. Econometric Society Monographs. Cambridge University Press (https://doi.org/10.1017/CCOL521633230.011).

Pesaran, M. H., Shin, Y., and Smith, R. J. (2001). Bounds testing approaches to the analysis of level relationships. *Journal of Applied Econometrics* **16**(3), 289–326 (https://doi.org/10.1002/jae.616).

Samani, K. (2020). Enforcing financial contracts through code running on censorship resistant and permissionless public blockchains. Twitter (@KyleSamani replying to @patrick_oshag), September 22, 06:40. URL: https://twitter.com/KyleSamani/status/1308280047984242688.

Zetzsche, D. A., Arner, D. W., and Buckley, R. P. (2020). Decentralized finance. *Journal of Financial Regulation* **6**(2), 172–203 (https://doi.org/10.1093/jfr/fjaa010).