NBER WORKING PAPER SERIES

DECRYPTING NEW AGE INTERNATIONAL CAPITAL FLOWS

Clemens Graf von Luckner
Carmen M. Reinhart
Kenneth S. Rogoff

Working Paper 29337
http://www.nber.org/papers/w29337

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2021

We thank Neil Shephard, Kathryn Holston, Juan Farah Yacoub, Matthew Ferranti and Pierfrancesco Mei for very helpful comments and to the Molly and Dominic Ferrante Fund for research support. The findings, interpretations, and conclusions expressed in this paper are those of the authors. They do not necessarily represent the views of the institutions they are affiliated with. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Decrypting New Age International Capital Flows
Clemens Graf von Luckner, Carmen M. Reinhart, and Kenneth S. Rogoff
NBER Working Paper No. 29337
October 2021
JEL No. E42,E51,E58,F21,F24,F32,F38

ABSTRACT

This paper employs high frequency transactions data on the world's oldest and most extensive centralized peer-to-peer Bitcoin market, which enables trade in the currencies of more than 135 countries. We develop an algorithm that allows, with high probability, the detection of "crypto vehicle transactions" in which crypto currency is used to move capital across borders or facilitate domestic transactions. In contrast to previous work which has used "on-chain" data, our approach enables one to investigate parts of the vastly larger pool of "off-chain" transactions. We find that, as a conservative lower bound, over 7 percent of the 45 million trades on the exchange we explore represent crypto vehicle transactions in which Bitcoin is used to make payment in fiat currency. Roughly 20 percent of these represent international capital flight/flows/remittances. Although our work cannot be used to put a price on cryptocurrencies, it provides the first systematic quantitative evidence that the transactional use of cryptocurrencies constitutes a fundamental component of their value, at least under the current regulatory regime.

Clemens Graf von Luckner
Department of Economics
Sciences Po
28, Rue des Saints-Peres
75007 Paris France
and World Bank
clemens.grafvonluckner@sciencespo.fr

Carmen M. Reinhart
Kennedy School of Government
Harvard University
79 JFK Street
Cambridge, MA 02138
and World Bank
carmen_reinhart@harvard.edu

Kenneth S. Rogoff
Thomas D Cabot Professor of Public Policy
Economics Department
Harvard University
Littauer Center 216
Cambridge, MA 02138-3001
and NBER
krogoff@harvard.edu

Introduction

Are cryptocurrencies, particularly as epitomized by Bitcoin, genuinely a potential competitor for fiat money or gold, or are they mainly a speculative bubble akin to 17th century Dutch tulips? A critical question is the extent to which there is an underlying transactions demand, now or in the future, to justify their high if wildly fluctuating prices. The results of this paper provide perhaps the strongest quantitative evidence to date demonstrating that, the use of Bitcoin for both domestic and international payments has been significant worldwide.

The vast bulk of media focus has been on cryptocurrencies as an evolving asset class. And although there has been some previous quantitative research providing some insight into the use of Bitcoin in facilitating illegal activities, this work typically concentrates on the analysis of the users, not the type of use, classifying users into legal and illegal agents, rather than distinguishing between investment purchases and transactions use. Furthermore, these analyses have only been applied to on-chain transactions,[1] that is, those that are immutably registered on the blockchain. On-chain transactions, however, constitute only a small share of the universe of all Bitcoin trades, most of which are "off-chain" utilizing some form of exchange, some heavily regulated, some not so much.[2]

This paper thus develops a new methodology for examining off-chain transaction data that allows one to probabilistically tag individual trades as "crypto vehicle transactions," that are clearly distinguished from the use of Bitcoin as an investment vehicle. We do not know whether results such as ours can potentially explain prices for crypto currencies, but they certainly show that the popular media focus on Bitcoin as speculative asset ignores its very active use in payments and capital flows.

---

1 See Chung, 2019; Foley et al, 2019; Framewala et al, 2020; Ron and Shamir, 2013; Yang et al 2019; Zhao and Guan, 2015.

2 Though arriving at a precise estimate is difficult, as not all exchanges provide data on trade volume, while other exchanges might have incentives to provide inaccurate data, the off-chain transaction volume involving Bitcoin appears to have been at least 10 times the volume recorded on the Bitcoin blockchain (Sources: CryptoCompare.com API, Blockchain.com API)

Importantly, the data set we use allows one to track what fiat currency is used for payment (which on-chain transactions do not) and to show that a considerable fraction of crypto vehicle transactions involve cross border flows.

Our laboratory is the world's largest centralized peer-to-peer exchange market, that includes more than four years of extremely high frequency data for 135 currencies. The data cover all trades made over that window, including currency used, the size of the Bitcoins purchased (down to units of Satoshi – one hundred millionths), and a precise time stamp to the second. By matching identical-size (to eight digits) transactions that take place within a short window, we are able to identify with high probability trades out of one fiat currency into another, as well as transactions that very likely represent domestic payments.[3]

We estimate that at least 7.4% of all trades in our data are crypto vehicle transactions; this is only a lower bound, as we only include trades we can identify as likely transactions with very high probability. Of these, around 20% represent international currency movements. These are average numbers; for some currencies (notably the ruble and the renminbi, the share of purely domestic trades is much higher. For others, most notably countries with significant restrictions on international capital flows, the share of international trades is much higher. Our international payments and transfer estimate is especially likely to understate the true flows, as when trades involve countries with relatively inefficient financial systems, the long time lags required to clear fiat currency payments can elude our methodology, which focuses on short time windows. We show, for example, that a much larger share is suggested by the "natural experiment" embodied in the March 7, 2019 widespread power outage in the República Bolivariana de Venezuela that turned out to have massive spillovers on Bitcoin markets in several other Latin American currencies.

The first part of the paper describes our novel core data set, with details provided in the Data Appendix. The second section describes the probabilistic matching methodology we develop. In a significant number of cases, the individual trades we pair are like matching snowflakes that occur nowhere

[3] The 1 percent fee charged by the exchange we study makes the platform impractical for high-speed arbitrage use.

else (or almost nowhere else) in the data set. This is because Bitcoin prices are volatile and the unit of account for most crypto-currency trades is the underlying fiat currency (e.g., where someone sends dollars from the United States to relatives elsewhere). The third section presents the main results, which we interpret as compelling evidence that Bitcoin are being used both for domestic payments and for international transmission of funds (including the large global pool of international remittances). The next section discusses how our methodology can, in principle, be applied (or adapted) to other cryptocurrencies and exchanges, including many classes of stablecoins. This analysis can perhaps guide regulators in identifying the selection of off-chain transactions to focus on when requesting IP addresses for audit. It also highlights how individuals or institutions, who may not have access to any private information, might still be able to draw inferences on the transaction use of cryptocurrencies and cross-border transfers, including capital flight. The final section concludes.

I: Local Bitcoins, and off-chain exchange

The core data set, described in detail in the Data Appendix (Appendix A.7), makes use of data from LocalBitcoins.com, the world's largest peer-to-peer (P2P) Bitcoin exchange, which allows Bitcoin to be traded in 135 different fiat currencies. The data encompass more than 45 million trades over the period March 15, 2017 – July 23, 2021.[4] To contextualize, the universe of trades that occurred on-chain over the same period, encompasses 452 million transactions.[5] LocalBitcoins.com is able to operate in such a large number of countries because it only offers deposit services for members' Bitcoins, but does not involve national fiat currencies, thereby side-stepping national banking regulations. Although other P2P exchanges

[4] LocalBitcoins has existed since 2012. However, we limit our analysis to the period since March 2017, when the website revamped the exchange's back-end, guaranteeing consistency in the format of the data.

[5] One transaction recorded on the Blockchain can include the transfer from one sender to many recipients, thus, when instead counting these as individual transfers, the cumulative number of on-chain transactions from the inception of Bitcoin in 2009 through May 2021 amounts to 862 million (Source: Blockchain.com API, Authors' Calculations) As noted previously, the universe of off-chain transactions is easily an order of magnitude larger than on-chain transactions.

have since copied LocalBitcoins' model, none approach its scope, at least as of this writing. LocalBitcoins, registered in Finland, remains the P2P Bitcoin exchange with the broadest service worldwide.[6]

A peer-to-peer transfer is a trade of Bitcoin between two pseudonymous private wallets, typically in direct exchange for goods or fiat currency. Trades within an exchange such as LocalBitcoins are "off-chain" in that individuals buy and sell only their claims on Bitcoins that the intermediary houses. Despite its many account holders, LocalBitcoins represents only one node on the blockchain.[7] Nevertheless, with a LocalBitcoins account, it is possible to buy Bitcoin from an account holder in country A using country A's currency, and sell to a third account holder in country B in exchange for country B's currency, all within LocalBitcoins and without ever registering on the blockchain.[8] (Intra-country payments and transfers are similarly straightforward.)

In principle, since exchanges are required to collect information on account holders (albeit cross-country standards vary greatly), this means that a determined and well-resourced regulator can track individual activity related to Bitcoin much more easily on an exchange than activity related to a pure "peer-to-peer" Bitcoin transfer. However, particularly in the case of an exchange that allows citizens from many countries to trade, such as LocalBitcoins, the international dimension makes this much more difficult. Although access to private information, including IP addresses, would likely be granted by Finnish authorities in egregious cases of crime or terrorism, it is unlikely to be granted on a routine basis. This implies that information on most transactions is generally still beyond the reach of national authorities from other countries (for example to detect evasion of capital controls). Nevertheless, we highlight that the

[6] That is, LocalBitcoins.com intermediates trades only to extent of that it waits to clear the internal transfer of claims on the exchange to Bitcoin only after payment has been confirmed. The fact it is only matching parties and not intermediating the fiat money payments is what allows LocalBitcoins to operate outside the financial regulatory framework in most countries, answering only to its home base regulator in Finland.

[7] Nodes are also referred to as *wallets*. However, a *keychain* might be the better analogy, given that one such node can have an infinite number of addresses its users can use to send Bitcoin via the blockchain to their account with LocalBitcoins.

[8] Because LocalBitcoins, as a whole, represents only one node on the Bitcoin blockchain, the only transfers visible on the public blockchain are thus transfers of Bitcoin to, and withdrawal of Bitcoin from, LocalBitcoins.com.

publicly available data alone is sufficient to establish that the market is being used in a significant way as a "crypto vehicle currency" to make fiat currency payments (at home or abroad) or to send capital abroad.

Of course, there is substantial anecdotal evidence on the use of crypto for transactions purposes, and it is well known that Bitcoin is the medium of choice on the Dark Web, not to mention ransomware. But systematic analysis has been lacking to support that claim. Our methodology allows us to systematically establish a lower bound on the use of Bitcoin as a medium of exchange. It is quantitative, rather than anecdotal. In principle, the same algorithm can potentially be applied to data from any P2P exchange, should the researcher or regulator obtain (or demand) similar data to ours. This includes not only P2P exchanges that mainly match buyers and sellers (and hold Bitcoins in escrow until both sides confirm the fiat money payment has cleared), but centralized P2P exchanges as well.[9]

Table 1 provides some descriptive statistics of the core data set.

Table 1: Descriptive Statistics

Number of Trades	45 528 193
USD Trade Volume	USD 11.0 billion
Average Trade Size (USD)	242
Average Trade Size (BTC)	0.03734 Bitcoin
Largest Trade Size Recorded	USD 2.3 million
Number of Fiat Currencies	135

Source: LocalBitcoins.com API, authors' calculations

[9] In a centralized exchange, typically catering only to a single country, the exchange serves as an intermediary, buying and selling on its own account, with price moving to match demand and supply. Coinbase is the largest such exchange in the United States. Again, it is important to emphasize that although "on-chain" transactions are public and therefore directly observable, the vast majority of off-chain transactions are proprietary information, making the LocalBitcoins data set – which is publicly available –unique at present.

For each trade, the data include the time-stamp, trade-size, fiat currency used, and the price paid in fiat currency. (On-chain transactions do not record currency or price.) We note that the average trade sizes are relatively small, in part because agents – who must communicate anyway to make the P2P exchange – have an incentive to shift modalities for large trades after matching. This allows them to engage in a more efficient peer-to-peer exchange, which might involve trading paper currency for crypto in person, thereby avoiding the 1% fee charged by the platform.[10] However, even within the limited trade we can observe, a significant proportion turns out to be transactional, and often cross-border.

II: Algorithm for detecting international crypto vehicle trades

In this section, we discuss our algorithm for (probabilistic) identification of cases where Bitcoin trades are likely being used for (cross-border) wealth transfers and payments. The mechanics are simple. Suppose an Argentine citizen wants to convert pesos in her Buenos Aires account to dollars in her Miami account, but evade Argentine capital controls. Or alternatively, she might want to buy a painting from a New York gallery to give to her sister who lives in New Jersey. Traditional interbank markets are often expensive and subject to capital controls, and whereas there are ways to avoid capital controls other than using crypto, many "traditional" ones have significant barriers to entry (e.g. trade mis-invoicing).[11] With the rise of Crypto markets all around the globe, there today exists a relatively easily and widely accessible alternative: The Argentine citizen can simply trade pesos for Bitcoin through an exchange or P2P platform (presumably from an Argentine resident), and then turn around and sell the same amount in exchange for dollars, presumably from an American resident. The dollars can then be used to make payments or be deposited in an American bank. The reverse process (from New York to Buenos Aires or elsewhere) encompasses the vast transactions associated with remittances. The process is illustrated in Figure 1. Crypto

[10] We have already distinguished P2P matching services such as LocalBitcoins from decentralized peer-to-peer trades which go through the blockchain; these involve paying miners a fee to verify the transaction, with the fees endogenously depending on congestion, but there is no centralized authority of any sort.
[11] Haibo 2008; Aizenman 2008; Coppola et al. 2020; Schneider 2003.

vehicle trades use the crypto asset only as a vehicle between fiat currencies, not as an investment asset or as legal tender itself.

Figure 1 – Crypto Vehicle Trade Flow Chart

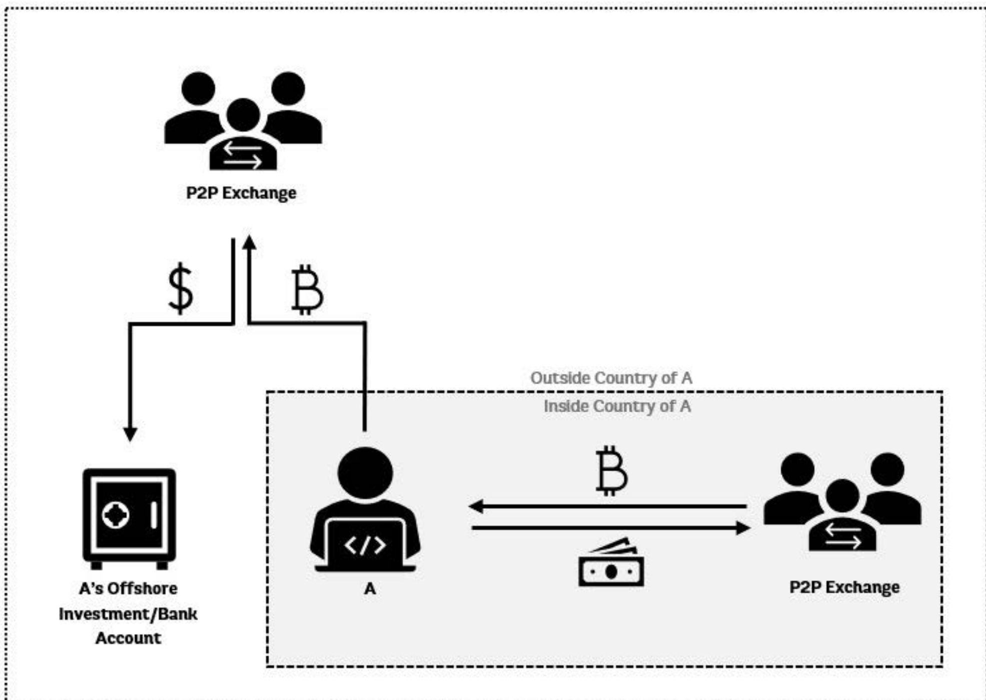

Given that Bitcoin prices are volatile, and that the fiat currency amounts being traded are highly varied, there is a very low probability of observing two identical-size matching trades (to eight digits), in and out of Bitcoin, within a relatively brief time window (say five hours) unless it is vehicle trade. Observing the same amount of bitcoin being traded twice within a short period of time, during which a countable number of trades take place, is thus akin to a probabilistic event, such as a die landing on the same side twice within a set number of throws – only that the Bitcoin *die* is not balanced, and has over 100 million sides.[12] We verify that a significant share of the 8-digit trades we document appear only twice in

[12] More precisely, when considering that a Bitcoin trade-size can indeed be greater than 1 Bitcoin, the "die" theoretically has 21 quadrillion sides - 21 million Bitcoins that can be mined, times 100 million Satoshi or decimal places.

the data set, and often within a very short time window. Because our data set contains quantity (in Satoshi), time stamp, price and fiat currency used, we can also see the flow of funds across countries and currencies, as well as use in domestic transactions.[13] Table 2 illustrates with an example from the data.

Table 2: Extract from the data

Timestamp	Trade Size	Price (Local Currency/Bitcoin)	Fiat Currency
2021-03-15 14:42:22	0.00098037	1.02E+11	Venezuelan Bolivar
2021-03-15 14:42:24	0.01157996	60449.26	US Dollar
2021-03-15 14:42:27	0.00022173	4509989.50	Indonesian Rial
2021-03-15 14:42:27	0.00047619	42000.04	British Pound
2021-03-15 14:42:28	0.00093023	6450017.50	Kenyan Shilling
2021-03-15 14:42:29	0.00063638	4321317.50	Russian Ruble
2021-03-15 14:42:33	0.00039107	1554708.87	Ukrainian Hryvnia
…	…	…	
2021-03-15 15:28:53	0.01157996	1.04E+11	Venezuelan Bolivar

Source: LocalBitcoins.com API

[13] As noted earlier, "on-chain" trades contain only the addresses involved, Bitcoin size and time, but not the fiat currency used or the price paid.

LocalBitcoins.com does charge a commission (averaging 1%) on both buy and sell trades, but this is generally paid directly by the market maker and does not affect the Bitcoin trade quantity or price reported, and therefore does not interfere with our matching algorithm. [14]

With this preamble, we now turn to the algorithm we use to identify crypto vehicle trades. Taking into account the large size of the data set, even with trade sizes documented out to eight digits, there is still a possibility of two identical-sized trades randomly appearing close together, especially as some trade sizes appear somewhat more frequently than others. The goal of our identification methodology is to arrive at an algorithm that identifies crypto vehicle transactions (both domestic and international) with a 95% confidence level. We aim to estimate the probability that two matching trades represent a crypto-currency vehicle trade, if they occur within, say, a five-hour window. The choice of window gives rise to the usual Type I and Type II errors trade-off. The shorter the window, the more currency vehicle trades we miss, the longer the window the more likely we are counting a random reoccurrence of an eight-digit match as and "in and out" vehicle trade. Our main results will turn out to be quite robust to the window choice; this is in part because the probabilistic approach directly controls for changes in the time window.

Our algorithm is constructed to generate an unbiased estimate of the share of trades on LocalBitcoins that are Crypto Vehicle Trades, while controlling for potential false discoveries

Let S be the set of all I individual trades in our dataset, i, each of which has a trade-size, x_i. Distinct trade sizes, denoted by x_k, are assumed be an element of

(1) $$X = \{x_1, \dots, x_K\},$$

[14] The 1% transaction fee (2% to both buy and sell) makes LocalBitcoins.com a relatively unattractive vehicle for arbitrageurs, and in any event they would most likely be larger market makers whose trades our algorithm misses anyway because sizes would no longer match because of fees.

which is fixed and known, $K \leq I$.[15] The number of distinct trade sizes K will be massively large, in our data greater than 7 million. Our null hypothesis (H_0) corresponds to a model of what one would expect if there were no vehicle trades, and any exact matching transactions were solely random.

<u>Assumption 1 (Null Model)</u> *Assume that under the model implied by the null hypothesis, trades of any given size appear as an independent Poisson process. The number of times any unique trade size, x_k occurs from time 0 to time t is thus defined as $PP(\vartheta p_k)$, where $\vartheta > 0$ and $p_k \geq 0$.[16] The Poisson process' intensity, ϑp_k, is the product of p_k, the probability of any new trade having the size x_k, such that $\sum_{k=1}^{K} p_k = 1$, and ϑ, the number of arrivals of trades over the time period of interest.*

Note that we will estimate p_k on the basis of our data set, making use of the frequency of each individual trade. The probability of a trade size that never occurs in the data set, thus has a probability, p_k, equal to zero.[17]

Consider the benchmark case of a 5-hour period following any given trade i, where N_i total trades happen and let n_i denote the number of times trade size x_i occurs. Under the null hypothesis, n_i follows a single multinomial draw,

$$(2) \qquad (n_1, \dots, n_I) \mid N_i \sim MultiN(N_i; p_1, \dots, p_K)$$

This is because conditioning on $N_{i,}$ removes ϑ from the conditional probability distribution.[18]

---

[15] We denote trade sizes by subscript k, as certain trade sizes occur more than once in the data so that $I \geq K$.
[16] We thus arrive at a K-dimensional vector of independent Poisson Processes for each trade size x_k.
[17] In contrast, when applying an alternative Bayesian approach, one might instead attach positive probabilities even to trades that did not occur.
[18] For every trade i, we are interested in the number of times the corresponding trade size x_i occurs in the subsequent five-hour windows. We denote this n_i. Since n_i is a function of the trade-size, and all trade-sizes that occur within the five hours after a trade, n_i is distinct to each trade i, and not solely a function of the trades-size x_i.

Rejection of the null thus implies the presence of vehicle trades in the data. To assess whether the null hypothesis holds or not, define

$$\text{(3)} \qquad \theta_i \; = \; P\,((\,n_i > 1)\,|\,N_i\,) \in [0,1]\,.$$

Note that without imposing any underlying structure to the data, θ_i can only be observed as a categorical variable, taking on the values $\{0,1\}$. Meanwhile, under the null model, the multinomial structure implies that $P\,((\,n_i > 1)\,|\,N_i\,)$ would equal

$$\text{(4)} \qquad \theta_i^* = 1 - \{(1 - p_i)^{N_i} + N_i p_i\}\,(1 - p_i)^{N_i - 1}\;, \qquad i = 1, \dots, I.$$[19]

When K is very large, N_i is relatively modest (typically we looking at window of 2, 5 or 10 hours out of several years of data), and p_i is very small (since prices are in Satoshi), then $\theta_i^* \cong 1 - (1 - p_i)^{N_i}$ as the chance $n_i > 2$ is close to zero.

To detect vehicle trades, we test for departures from the model under the null hypothesis of purely random pairings.[20]

Note that if p_i is not very small (as we shall see later, in Figure 2 below, certain size trades are common), then $n_i > 1$ becomes much more likely under the null. We formalize this in the following way.

[19] Note that $p_i = p_j$, as x_i is drawn from X with probability p_j

[20] As a robustness check, the Appendix introduces a similar approach using independent, distinct time windows, which are thus not prone to potential serial correlation.

Definition 2 *Let* $\Theta_\theta \in$ *[0; 1] be some preset number. The trade i is not a candidate for a statistical vehicle trade of size* x_i*, if*

$$(5) \qquad H_{0,i}\colon\ \theta_i^* \geq \Theta_\theta\,, \qquad i = 1, \dots, I\,,$$

Otherwise i is potentially a statistical vehicle trade of size x_i*.*

Recalling that I is the total number of trades in the entire data set, the "vehicle trade share estimand" is thus

$$(6) \qquad \varphi = \frac{2\sum_{i=1}^{I} \boldsymbol{\alpha}_i\left(\theta_i - \theta_i^*\right)}{I},\ with\ \boldsymbol{\alpha}_i = \begin{cases} 1\ if\ \theta_i^* < \Theta_\theta \\ 0\ if\ \theta_i^* \geq \Theta_\theta \end{cases}$$

Note that because every vehicle trade consists of two trade legs and since our algorithm identifies the first leg of any vehicle trade, the share of trades that we identify is multiplied by two.

Under the null hypothesis $\varphi = 0$, since φ captures the excess clustering of trades at particular trade sizes, that is, in excess of what one would expect under H_0.

We set Θ_θ to 0.05 to remove trades which are relatively common across the entire data set, so $n_i = 2$ provides little evidence of a vehicle trade.

Given that I is large, it is important to recognize there are I hypotheses, relating our approach to the multiple hypotheses testing literature (Efron (2007), Wakefield, J. (2007)), which has been extensively applied to genomic sequencing and chromosome segmentation. However, in this literature, because of technological/economic constraints, one is unlikely to possess the population of the underlying data. By contrast, in this exercise we have access of the full distribution of trades (the full LocalBitcoins trade data set). This allows us to control for "false positives" by estimating the expected share misidentified as vehicle trades by our algorithm through randomized sampling from our overall data set, in a way that by construction, does not contain any vehicle trades.

Definition 3 *We declare a "discovery" of a vehicle trade of size* x_i*, if*

$$n_i > 1 \quad and \quad \theta_i^* \leq \Theta_\theta \ . \tag{7}$$

This is recorded by $d_i = \boldsymbol{\alpha}_i \, \boldsymbol{\phi}_i \qquad with \quad \boldsymbol{\phi}_i = \begin{cases} 1 \; if \;\; n_i > 1 \\ 0 \; otherwise \end{cases}$.

For each trade *i*, this discovery algorithm is a single hypothesis test with a size (under the null hypothesis) of θ_i^*.

To control for false "discoveries", and to establish our estimate of the share of crypto vehicle trades, φ, we introduce a measure of the number of trades one would expect to falsely discover as vehicle trades, c_i.

$$\hat{\varphi} = \frac{2\sum_{i=1}^{I} (d_i - c_i)}{I}, \tag{8}$$

where $c_i = \boldsymbol{\alpha}_i \, \theta_i^*$. Over *i*, c_i describes the probabilities of a trade seeing a random match within five hours under the null model. Summing c_i over *i* thus represents a measure of the expected number of matches, assuming the null model holds.

Theorem 4 *Under an arbitrary data generating process for* $(n_1, \ldots, n_I)$,

$$\text{(9)} \qquad \mathrm{E}[\,\hat{\varphi} \mid N_{i,\ldots,I}\,] = \varphi$$

Proof of Theorem 4:

$$\text{(10)} \qquad \mathrm{E}[\,\hat{\varphi} \mid N_{i,\ldots,I}\,] = \frac{2\sum_{i=1}^{I}(E\,[\,d_i \mid N_i] - c_i)}{I}$$

$$\text{(11)} \qquad = \frac{2}{I}\sum_{i=1}^{I} \boldsymbol{\alpha}_i \;(E\;[\,\boldsymbol{\phi}_i \mid N_i] - \theta_i^*)$$

$$\text{(12)} \qquad = \varphi$$

Where we make use of the fact that for any single i, $E\;[\,\boldsymbol{\phi}_i \mid N_i\,] = 1 * P\,((\,n_i > 1) \mid N_i\,)$.

The algorithm's capacity to identify capital flows (cross-border transactions) thus relies on two assumptions: First, for trades to be matched with a degree of confidence, individual crypto trade sizes must be sufficiently unique. If the majority of the trades had the exact same nominal size, the matching algorithm would be of limited use. This feature does not, however, characterize our data, where there exist 7.4 million different trade sizes (more precisely 7,405,680).[21] More than two-thirds of these occur twice or less. The distribution of trade sizes for the data set is illustrated in Figure 2.

[21] The value which occurs most often in the dataset is 0.00010000 BTC, with 44,544 trades having that nominal value.

Figure 2 - The historical distribution of trade-sizes

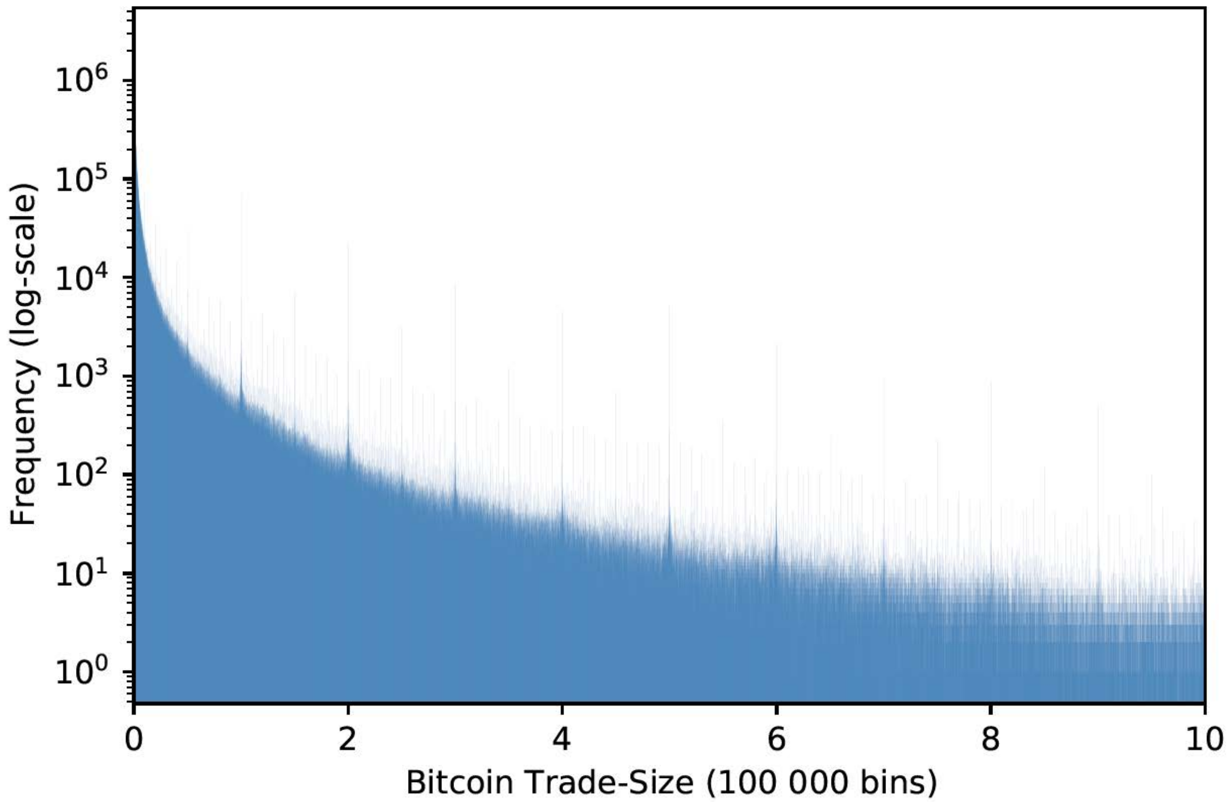

The figure includes trades with sizes between 0.00000001 Bitcoin and 10 Bitcoin. We apply this distribution to derive the applied probability density function that is used to estimate the probability of a trade size occurring under the null model, and thus to control for the probability of false discovery. Source: LocalBitcoins.com API, Authors' Calculations

The second assumption the algorithm relies on is that market participants who are aiming to use Bitcoin as a crypto vehicle to make transactions in fiat currency, will have strong incentives to minimize their holding times. The main constraint being that the requisite fiat money transfers on both ends, buying and selling, can take time, which is especially an issue for some less liquid developing economy currencies. Indeed, for many trades – though depending on the fiat payment channel used – the time between the trade legs of a vehicle trade is typically only a few hours or less (see Appendix A.2). This is likely driven by the high volatility of Bitcoin prices. Since significant delays between the purchase and the execution of trades would risk leading to losses on the buyer or seller side, depending on whether Bitcoin prices rise or fall.

To highlight how volatile Bitcoin prices are, Figure 3 shows the intraday volatility of Bitcoin markets since 2016, measured by the daily maximum drawdown – that is, what share of one's investment one could have lost if one happened to buy at the high and sell at the low on any given day – compared to the same measure for two highly liquid FX-pairs. The point made by Figure 3 is in line with recent studies finding Bitcoin-fiat volatility to exceed the volatility of major currency pairs by a factor of ten (Baur and Dimpfl, 2021).[22] Even when assessing longer time windows, the volatility of Bitcoin markets, and thus the risk of holding Bitcoin for extended time periods, remains high in comparison: The annualized standard deviation of the USD/BTC since 2014 has been 93%, compared to 8% and 12% for the exemplary USD/EUR and USD/MXN exchange rate pairs. Our assumption that market participants not interested in exposure to Bitcoin per se, would try to trade in and out of the digital currency as quickly as possible, is thus likely to hold in most cases.[23]

[22] Of course, especially in the post-pandemic context, there exist some developing economy currencies which are equally spectacularly volatile.

[23] It is certainly possible that some percentage of agents using Bitcoin mainly as transfer vehicle do not mind – or possibly even prefer -- some exposure to price volatility. To the extent we are too conservative in picking a relatively short time window, this constitutes another reason why our estimates are a lower bound on crypto vehicle trades.

Figure 3 - Bitcoin price volatility compared to volatility of two forex pairs

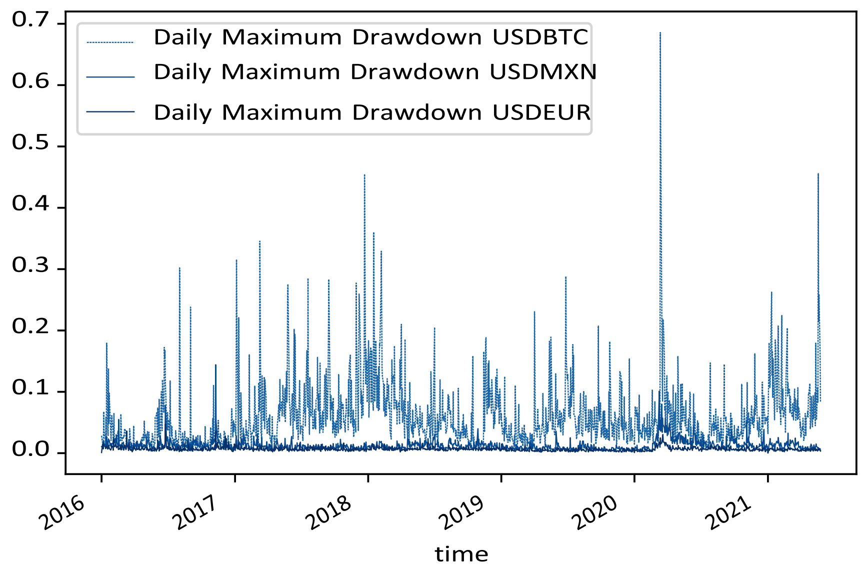

Source: CryptoCompare.com API, Bloomberg Terminal, Authors' Calculations

However, even though participants engaged in crypto vehicle trades have a strong incentive to get in and out quickly, in practice there can be speed limits. Within the LocalBitcoins platform, trades can clear very quickly; the time between an order being made and the escrow being released is typically very short. However when crypto vehicle trades involve two different fiat currency transfers, these can cause a slightly longer delay, or in the case of an international transfer, potentially a much longer delay, especially if the market (or the banking system) for at least one of the currencies has inefficient clearing facilities.[24] Indeed, as we shall discuss, our findings suggest that the time elapsed between two trades making up cross-border vehicle trade tends to be longer than for domestic crypto vehicle trades. The five-hour time window chosen (we will also report results for shorter and longer windows) reflects the findings from the analysis of trade delays, as well as the trade-off described above. For this reason, we also consider other time windows as

[24] For example, when in a given market there exist no fintech alternatives to the interbank market to make domestic transactions, transactions would usually take at least one working day to clear, meaning the Bitcoin would remain in the escrow for that long, meaning in turn that the time between two trades must exceed the five-hour window we consider.

robustness checks. Further, by means of a Monte Carlo simulation, applying this methodology to a randomly shuffled data set, we will present evidence on how applying different time windows affects the likelihood of false positives produced by the algorithm.

III: Results

Table 3 gives results for our matching algorithm using a five-hour window to identify crypto vehicle trades. Of the 45 million trades, 10.5% or 4.78 million trades have an exact match (or are matched) in terms of Satoshi size. However, running our algorithm very conservatively excludes 0.89 million of these trades because the trade size is sufficiently common (for example 0.1 Bitcoin) that the algorithm cannot attach a 95% confidence interval to a matched pair of being a crypto vehicle trade. Further, we deduct the number of trades one would expect to match with a 95% confidence interval even in a data set without any real vehicle trades (i.e. the False Discovery Rate). This brings the percentage of crypto vehicle trades in the data set down to a conservative lower bound of 7.4%.

Table 3: Crypto Vehicle Trades

Total Number of trades	45 528 193
Number of trades with one match	4 782 374
Number of trades identified as vehicle trades ((P(Match is Random) < 0.05), d_i	3 895 362
Expected "False Discoveries" in a data set without vehicle trades, c_i	522 251
Share of total trades identified as crypto vehicle trades, φ	**7.4%**

Share of trade volume identified as vehicle trades, γ (net of False Discoveries) [25]	**3.8%**
Cross-border trades[26] / identified crypto vehicle-trades	17%
Cross-border trades / identified crypto vehicle-trades (in USD) Volume)	20%
Cross-border trades / identified crypto vehicle-trades (in USD) Volume) when excluding RMB and RUB	27%

Source: LocalBitcoins.com API, Authors' Calculations

Table 3 also shows that of the 7.4% of all trade we identify as crypto vehicle trades, just under one-fifth are cross-currency/cross border; that is, where two different fiat currencies are used in the two transactions of exactly matching Satoshi size. In Appendix A.2, we perform a Monte Carlo simulation where we draw random samples of trades from the data set.[27] Applying the algorithm on these randomly constructed data sets, only about 1% are individually identified as vehicle trades. And exactly 0% remain identified as vehicle trades after deducting the false discovery control discussed in section II.

Table A.3 in the Appendix looks at the robustness of our results to the time window, using two-hour, five-hour and ten-hour windows. As can be seen in the table, as the time-window increases, the number of candidate vehicle trades increases. However, because the number of trades encompassed

[25] See the Appendix for an explanation of how the methodology is adjusted to calculate the share of the trade volume.
[26] Note that here we consider only trades with two different currencies as cross-border. However, during limited field experiments one of the authors found the US dollar to be used for cross-border crypto vehicle trades in countries whose primary legal tender is not the US dollar.
[27] For details see Appendix A2.

increases as N_i increases, ceteris paribus, P increases, so that the number of trades that we identify as matched vehicle trades with at least a 95% confidence level decreases (Appendix A.1, Figure A.2)

In our baseline methodology, to assess the unconditional probability of any trade size occurring, we apply a uniform historical distribution of trade sizes over 2017-2021. One might conjecture that this distribution could be affected by changing dynamics over time (see Appendix A.3) and that the methodology should potentially control for this by embedding a time-series dimension in the assessment of the values of $\{p_j, \ldots, p_Q\}$. To assess whether such a dynamic is likely to significantly affect our point estimates, we compare the share of trades that are identified by our core methodology (uniform distribution of trade-sizes) versus vehicle trades for each calendar year. We derive the p-values from a year-specific distribution function. Table 4 presents evidence that the impact of such dynamic is indeed limited, and that the computationally significantly more efficient approach applying a uniform trade-size distribution appears to provide a close enough approximation.

Table 4: Comparing results using uniform and year-specific trade-size distributions

	(1) Uniform trade-size distribution	(2) Year-specific trade size distribution
2017	6.3 %	5.8 %
2018	8.1 %	7.6 %
2019	7.8 %	7.5 %
2020	6.6 %	6.3 %
2021	7.4 %	6.9 %

Source: LocalBitcoins.com API, Authors' Calculations

Importantly, at longer time windows, as the number of trades identified as vehicle trades decreases, the share of trades that are cross-border trades increases. This observation is driven by two separate dynamics:

1. As the difference between Figure 4a and Figure 4b below, the average (median) time that passes between two trades that are identified as two legs of one vehicle trade differs significantly between domestic and international transfers: 44 (13) vs 83 (45) minutes.
2. With an average (median) trade-size of 139 (45) USD of cross-border vehicle trades, compared to 115 (28) USD for domestic trades, cross-border trade sizes are significantly larger.[28]

Figure 4 - Time between two trade legs of vehicle trades.

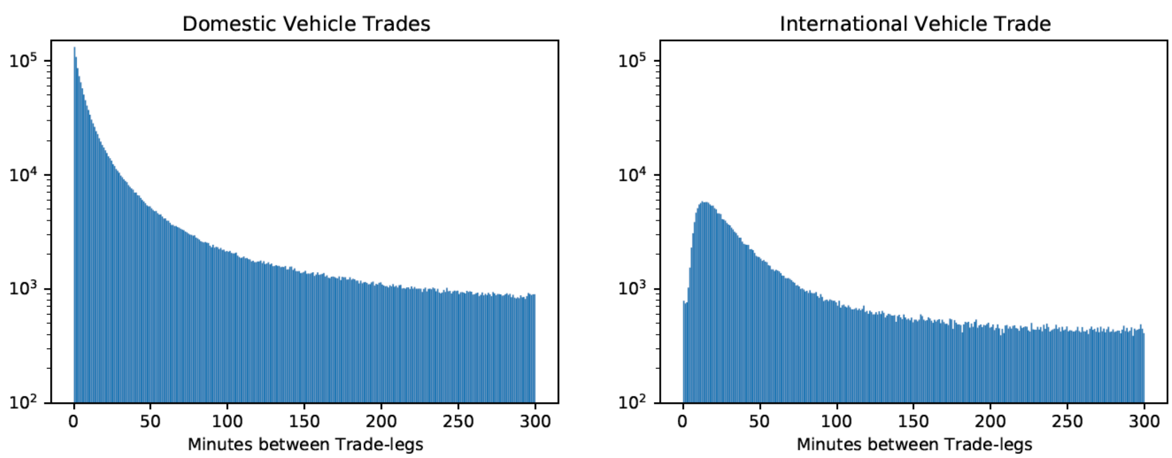

(a) Domestic Vehicle Trades (b) Cross-Border Vehicle Trades

Source: LocalBitcoins.com API, Authors' Calculations

[28] These statistics are based on the identified trades using a five-hour window. Whenever we present ratios or shares of vehicle trades, we moreover consider any trade whose individual hypothesis test leads to a rejection of the match not being a vehicle trade with a 95% confidence level. Under the assumption that false-discoveries are homogenously distributed across our sample, the false-discoveries and numerators and denominators cancel out, so that the false-discovery-control is not required.

Table 5 breaks down the 3.9 million crypto vehicle trades we have identified into the most common origin-destination pairs. Notably, the Russian Federation and China stand out as having an exceptionally high share of same currency trades, that is where the ruble or renminbi, respectively, is both the origin and destination currency, accounting for 89.8% and 84% of all trades involving these currencies. In the case of Russia, US financial sanctions may be responsible for the popularity of Bitcoin as a vehicle for domestic transactions, and it may also offer some limited capacity to evade state surveillance. In the case of the US dollar, euro and pound, internal transactions are still quite significant, accounting for around 70% of all transactions.[29] Because these are international currencies, it is a possible that a percentage of these represent transactions involving residents of other countries.

[29] Note that the instances where the same currency occurs twice are also unlikely to be simply due to tumbling – the practice of mixing bitcoins in a pool so to obscure a transactional trail. Because LocalBitcoins represents a *Webwallet*, transactions by multiple customers can be arriving or departing from the same node on the blockchain. Hence, simply sending Bitcoins to one's LocalBitcoins account, and later withdrawing the Bitcoins from the account already suffices to tumble the coins. In other words, the tumbling does not require a transaction involving any fiat currencies.

Table 5: Largest Currency Pairs in Crypto Vehicle Trades

Origin Currency	Destination Currency	Trade Volume as share of total identified vehicle trade volume in Origin Currency (in %)
Russian Ruble	Russian Ruble	89.8
Venezuelan Bolivar	Venezuelan Bolivar	51.7
Yuan Renminbi	Yuan Renminbi	84.0
US Dollar	US Dollar	60.1
Euro	Euro	76.5
Colombian Peso	Colombian Peso	62.1
Pound Sterling	Pound Sterling	76.9
US Dollar	Nigerian Naira	24.3
Nigerian Naira	Nigerian Naira	26.3
Colombian Peso	Venezuelan Bolivar	24.8

Source: LocalBitcoins.com API, Authors' Calculations

Out of the 10 top crypto vehicle trades listed in Table 5, two are cross border: US dollars to Nigerian naira, and Colombian pesos to Venezuelan bolivars, both likely capturing remittance flows around fairly strict Venezuelan controls (although these can also represent official Venezuelan entities avoiding US sanctions.) Figure 5 below and Table A1 in the Appendix expands the list to the top 10 cross currency crypto vehicle trades, with Peru, Chile and Argentina alongside Colombia, Venezuela and Nigeria.

Figure 5 - The World's 25 biggest Crypto Vehicle Channels.

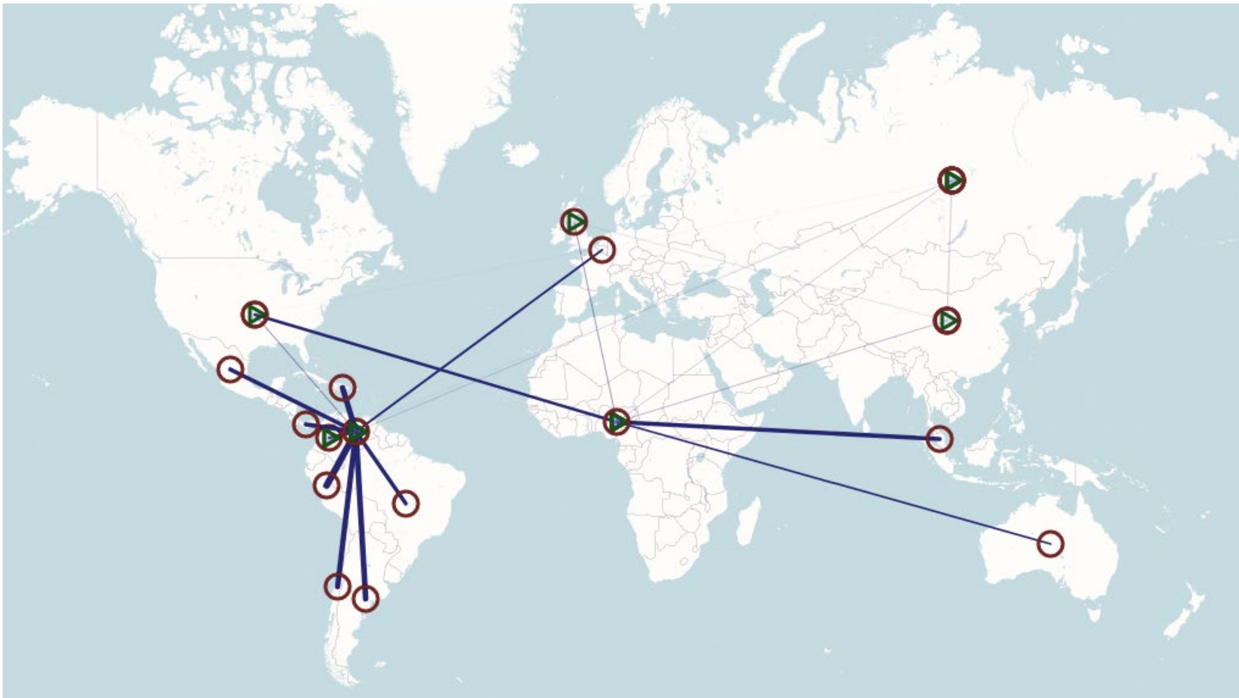

Circles: Origin, Triangles: Destination. Line-width: Channel volume as share identified trade volume in Origin Currency. See Table A.1 in the Appendix for further statistics. Sources: LocalBitcoins.com API, Authors' Calculations

It is important to emphasize that whereas LocalBitcoins is a truly worldwide market, the relative trade volume across different currencies has naturally shifted over time as the crypto market has evolved, depending on a host of factors. These include the rise and fall of competing local exchanges, fluctuating regulatory constraints on crypto, the changes in capital controls (sometimes sudden during financial or debt crises), market size, etc. Table 6 gives the 10 most active markets, measured by annual LocalBitcoins transactions per capita, for the years 2018-2020. The results overlap with Table 5, which does not control for population size.

Table 6: Ten most active markets, scaled by population, 2018-2020.

2018	2019	2020
Russian Ruble	Venezuelan Bolivar	Venezuelan Bolivar
Venezuelan Bolivar	Russian Ruble	Russian Ruble
Belarusian Ruble	Belarusian Ruble	Belarusian Ruble
Swedish Krona	Panamanian Balboa	Colombian Peso
Panamanian Balboa	Peruvian Sol	Chilean Peso
British Pound	Swedish Krona	Peruvian Sol
Peruvian Sol	Colombian Peso	Swedish Krona
New Zealand Dollar	Chilean Peso	Panamanian Balboa
Colombian Peso	British Pound	British Pound
Nigerian Naira	Kazakhstani Tenge	Argentine Peso

Source: LocalBitcoins.com API, Authors' Calculations

The question of the determinants of crypto use in different markets merits further investigation, which will be facilitated by expanding the number of exchanges incorporated, as comparable data becomes available. Yet, the notion that the use of crypto currencies is changing the economics of capital control evasion, and that, in turn, capital control evasion is an important driver for the expansion of crypto markets is broadly in line by the list of countries in Table 6. Though many factors render a more systematic analysis impossible,[30] Figure 6 presents the results of an exemplary case study, further illustrating the link between capital controls and crypto market

[30] The most important factor being the small sample size of countries introducing capital controls over the period of our sample, Additionally, when demand for crypto currencies rises, as the need to evade capital controls becomes imminent, often, alternative and more cost-efficient exchanges open, which then reduces trade volumes on LocalBitcoins.com. Finally, one of the authors' field research suggests that the recent rise of stable coins has at least in some cases, such as Lebanon, led to the former taking over the role of Bitcoin as the major crypto vehicle to evade capital..

expansion. Namely, the figure presents the number of trades in our data set where the fiat currency used is the Argentine peso around the 2019 instance, when the Argentine government imposed capital controls. As is discussed by Ilzetzki, Reinhart and Rogoff (2019), inter alia, the rise of a parallel market for hard currency is the best market-based indicator for the existence of capital controls. The rapid expansion of the Bitcoin market in Argentina, occurring in lockstep with the rise of the parallel market premium, thus is consistent with crypto vehicle trades having become the 21st century's novel channel for capital control evasion.

Figure 6 – Example of Bitcoin market expansion coinciding with capital controls

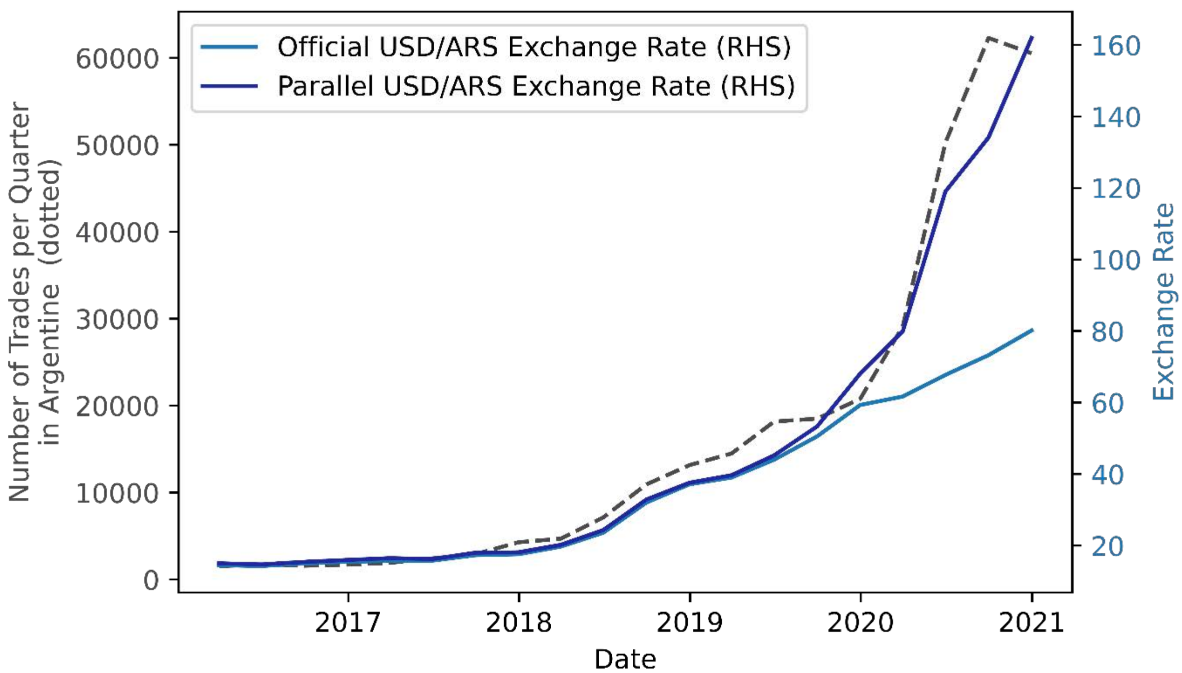

Source: LocalBitcoins.com API, BlueDollar.net, Authors' Calculations

Finally, Figure 7 expands the list to the 20 highest volume crypto vehicle trade currencies, breaking them down into the share where the two matched transactions are in the same currency (which we interpret as domestic transactions) and the share that are cross border. It is notable that the currencies with the highest share of cross-country transactions align well with countries that have had significant capital controls throughout the period, for example, Argentina and Nigeria. In Venezuela, although there are capital controls in place, that domestic transactions are as significant as international ones comes as little surprise, given the unfolding hyperinflation and strict bank-withdrawal limits, which make Bitcoin an attractive alternative to the local currency and domestic banking channels for domestic payments and transactions.

Figure 7 - Share of identified vehicle trades that are cross-border (i.e. where origin currency and destination currency are not the same) vs domestic.

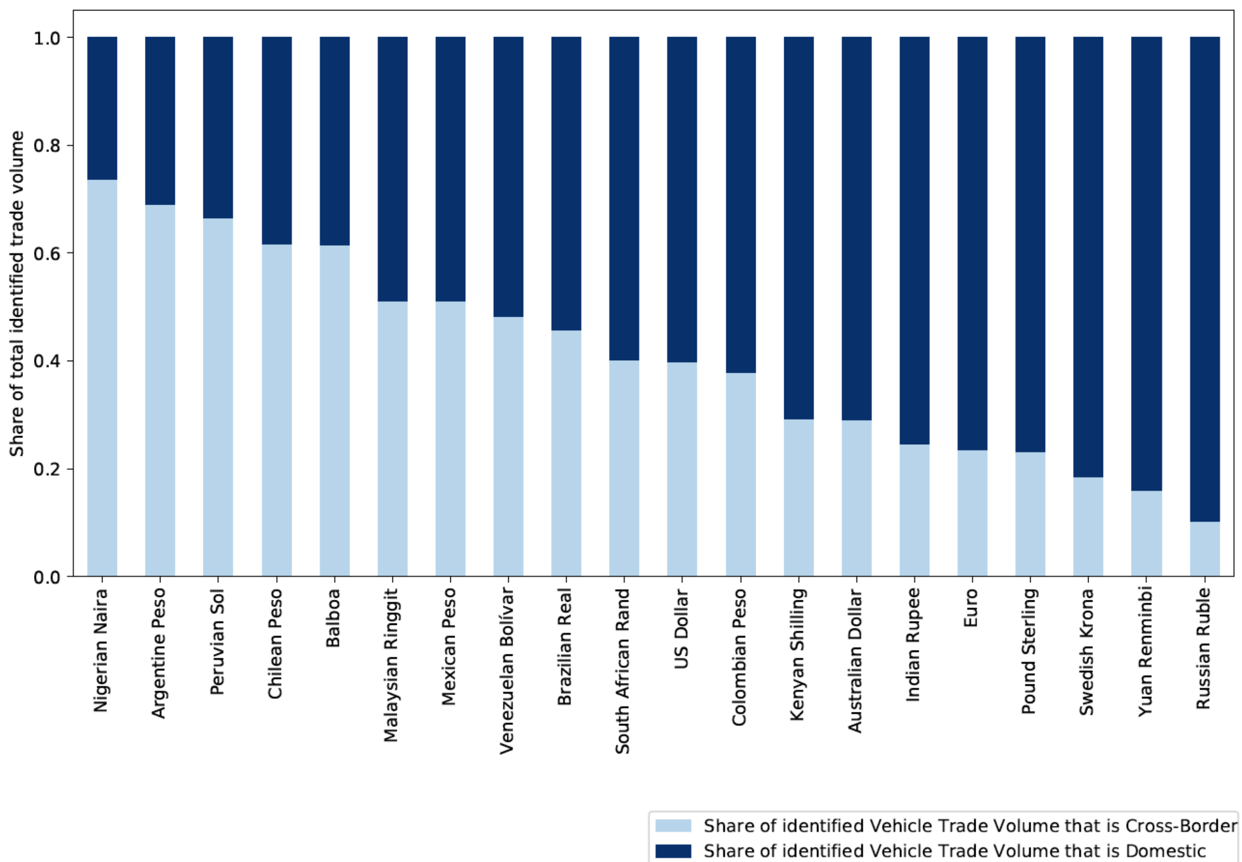

Source: LocalBitcoins.com API, Authors' Calculations

We reiterate that our estimates are **lower bound** estimates on crypto vehicle trades for several reasons, described in Table 7.

Table 7 – Factors leading to algorithm understating the share of trades being crypto vehicle trades

Factor leading to algorithm understating	Explanation
The P2P platform (LocalBitcoins in this application) used only for one of the two legs	Imagine a remittance sending agent, working and living in a country with centralized exchanges and sending his remittances to a developing country, where the only platform available for Bitcoin trades is the P2P exchange. It might be cost efficient to purchase the Bitcoin on the lower cost centralized exchange of the sending country, transfer it to the e-wallet of the P2P platform, and then sell it on the P2P platform in the destination country. Our algorithm would not detect this vehicle trade, as only one of two trades is recorded in our P2P data.
The amount in Bitcoin of the two legs might differ.	The algorithm only identifies a vehicle trade where the amount of Bitcoin for both two legs is identical. Yet, some cases might exist, where an agent uses his wallet both for vehicle trades and speculative purposes, so the two amounts might differ. The same can happen when the best P2P price offer of the second leg is limited to a trade-size $Z < X$, when the agent to might split the second trade into two.
The two trades lie farther apart than our time window (five hours in the baseline case) allows.	Our algorithm does not match any two trades when the time difference between them exceeds 300 minutes (for a five-hour window). If the trader has access to a perpetual futures contract, that allows her to wait for better terms of trade on the destination's country P2P market, this can allow her to accept much longer periods of exposure to Bitcoin price volatility.
The payment technology in (at least) one country involved is slow.	Because the clearing time of the two trades is highly dependent on transfer technologies in the countries involved, our methodology might not capture trades from countries where the prevalent money transaction technology has a clearance time that exceeds the chosen window.

The trade is matched but wrongfully disregarded, because of the probability of it being matched by chance	The algorithm applies a fairly conservative methodology in identifying trades, skewed towards reducing Type 1 errors (False-Identification as vehicle trades) at the cost of larger Type 2 Errors (disregarding a trade that indeed was a vehicle trade).
The parties involved in a planned trade cancel the transaction to complete it in cash and avoid the escrow fees	Whereas a 1% fee for the transaction amount of 100 USD hardly creates an incentive to face the potential risk and nuisance of a arranging an in person transaction, anecdotal evidence suggests that agents seeking to make larger trades often seek to circumvent the fees by cancelling trades on the LocalBitcoins.com and instead use the messenger function on LocalBitcoins.com to organize an exchange of the crypto currency for cash and in person. Offers by sellers - who as market makers pay the 1% fees, proposing to share the amount saved by avoiding the fees are indeed common on the platform.
The sending party is a market maker in one of the two trade-legs	Being the market maker in a trade, in our application, means fees will be deducted from the nominal Bitcoin amount, and a division of the total nominal and multiple smaller trades is more likely.

Section IV: Event Study: República Bolivariana de Venezuela

To underscore that our crypto cross-country currency vehicle estimates are lower bounds, it is interesting to consider the spillover effects to global LocalBitcoins trade of a massive unanticipated three-day power outage that took place in the Venezuela, starting March 7th, 2019.

If a significant share of the trades are indeed paired to implement cross-border transfers, an exogenous shock that constrains the ability to trade Bitcoins in one currency should impact Bitcoin trade in other currencies that are a major destination for/origin for cross-border-flows from the affected country.

A major power-cut in the Venezuela that began on March 7, 2019, provides an interesting natural experiment.[31] The power-cut, caused by an incident at the country's major hydro-electric power plant at Guri Dam, left more than 30 million Venezuelans without electricity for more than 72 hours. Many Bitcoin trades were obviously halted. A question worth asking is how this affected Venezuela's main trading partners' (as identified by our algorithm) Bitcoin trade volume.

Figure 8 - Event Study: Nationwide Power Cut in Venezuela, March 2019

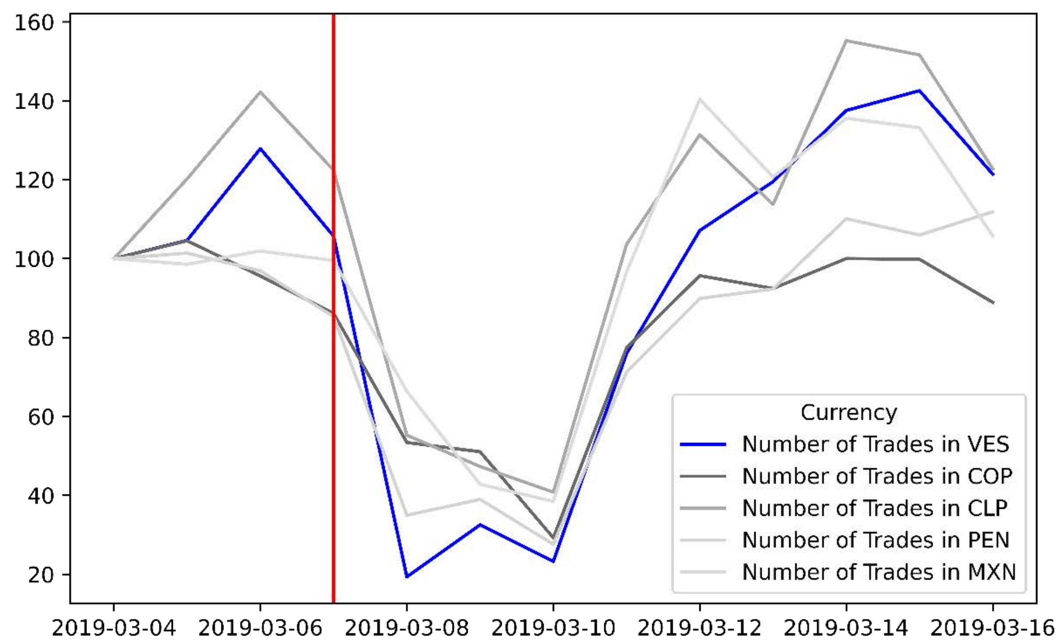

The number of trades is normalized to 100 on March 4th, three days before the power cut. See Figure A1 in the Appendix for the same graph, including a control group of countries not identified as engaged in Vehicle Trades with Venezuela. Source: LocalBitcoins.com API, Authors' Calculations.

Figure 8 shows the results of the event study, comparing P2P trade volumes in Venezuela around the time of the power cut with the number of trades effectuated in four other currencies: the Mexican peso, the Peruvian sol, the Chilean peso and the Colombian peso—these have been important destinations in the

[31] See The Guardian, "Venezuela: huge power outage leaves much of country in the dark", from March 8th, 2019 or BBC News

Venezuelan diaspora.[32] The event study highlights (a) the importance of cross-border crypto vehicle trades in the Bitcoin market and (b) that the share we identify as crypto vehicle trades likely only represents a lower bound of the true volume of crypto vehicle trades.

Indeed, the results from this natural experiment suggest that for the Venezuela and its main cash transfer partners, crypto vehicle trades seem to constitute a very large share of total trades made, over 50% for Mexico, Peru, Chile and Colombia. (For Venezuela, the drop also reflects a fall in internal payments.) It is instructive to compare Figure 7 and Figure 8. Figure 7 shows at the share of cross-currency trade pairs (within the five hour window) out of all trades identified as transactions related, and shares for Mexico, Peru, Chile and Colombia are also quite high, ranging from 40% to 70%, But in Figure 8 the drop in these same countries represents a fall as a share of *all* P2P trades, which is an order of magnitude larger.

Consider the case of Peru, where the algorithm produces a lower bound of 7.5% for the share of all trades involving the Peruvian sol identified as crypto vehicle trades.[33] Of the 7.5% percent, 38% represent trades where the other fiat currency involved is the Venezuelan bolivar.[34] If this lower-bound estimate captured most of the action, then we would expect the collapse of the Venezuelan market due to the electricity shutoff to lead to an approximately 3% reduction in LocalBitcoins trades in Peruvian sol during the course of the blackout (38% of 7.5%). Instead, as Figure 8 shows, the actual drop in trade volume was 65%.

The same can be shown for other currencies our algorithm identifies as equally important senders or receivers of crypto vehicle trades from or to Venezuela. In addition to the Peruvian sol, Mexican peso and the Chilean peso, these also include Argentine peso and the Panamanian balboa. We note, however, that the gap between our algorithm's lower bound estimate and the true measure of crypto cross-country

[32] See also Matt Ahlborg's medium post from March 24t, 2020 for a related discussion (https://medium.com/open-money-initiative/latin-american-Bitcoin-trading-follows-the-heartbeat-ofvenezuela-71a28cb86ba0).
[33] In Mexico the equivalent amounts to 6.6%, in Colombia 6.5% and in Chile 7.2%.
[34] In Mexico the share of identified trade volume that moves to or originates fromVenezuela amounts to 42%. In Colombia and Chile the same amounts to 30% and 35% respectively.

vehicle trades is likely to be particularly large for transactions involving Venezuela's currency, where lags in fiat money payments are likely to take much longer than for most other countries, given the country's economic dysfunction during this period. Thus, our maximum 5-hour window (and even the 24-hour window results included in the appendix) is likely to miss the majority of crypto-vehicle transactions involving Venezuela.

V: Extensions and applications

We have used comprehensive off-chain transactions data from what has been the world's largest peer-to-peer crypto exchange platform over the past five years. The analysis provides evidence that strongly suggests that Bitcoin is used actively as a vehicle currency in international transactions; in most countries it is also used extensively as vehicle for domestic currency transactions. Thus, it runs counter the oft-expressed view of crypto currencies as a purely speculative asset class. As we have emphasized, the nature of the data and, in particular, our ability to trace transactions price and currency is completely different than for research analyzing on-chain transactions.

Nevertheless, as computationally intensive as our exercise may be, the LocalBitcoins.com data set represents only a small share of the universe of off-chain crypto transactions when one includes centralized exchanges (where the exchange acts as financial intermediary).

In principle, our methodology can be applied for any more targeted investigation in any particular country/region, as well as to data from any exchange that identifies trades in terms of the fiat currency used to purchase crypto, as long as a number of minimum conditions are fulfilled (see Appendix A.6 for the precise description). Of course, when applying our methodology to data from other exchanges, one must account of their individual structure and features, including the average speed of clearing fiat money payments, and how to factor in pro-rata transactions fees. Also, in some centralized exchanges, the fee structure makes it more feasible to engage in very high-frequency speculative trading, although these almost invariably involve buying and selling in the same currency and one can thus exclude the impact of such trades on aggregates by exclusively analyzing international crypto vehicle currency trades.[35]

Although data from many exchanges is private, regulators can typically access data from centralized exchanges in their own local jurisdiction, or potentially beyond that given sufficient international cooperation. Regulators might, for example, use this algorithm to identify suspicious cases for which they can make targeted requests for IP addresses from exchanges. Note that our methodology can also be applied to assess the probability of a single pair of trades constituting a Crypto Vehicle Transaction for any time window (which could be one minute or one week) and for any probability threshold (which could be, say 80% instead of 95%). Our approach also allows researchers to show how cryptocurrencies are used for off-chain capital flows and transactions, without requiring knowledge of private data.

In principle, the methodology can also be applied to identify/estimate international capital transfers where an agent buys Bitcoin on an exchange then transfers the funds out of the exchange's wallet to their own personal wallet, that is a combination of off-chain and on-chain transaction. The on-chain data is already public, so one needs only access to the data from the exchange where the off-chain transaction was executed.

Moreover, rather than estimating aggregates of (cross-border) capital flow volumes, the methodology can also be applied to assess the individual probability of two trades constituting a capital transfer. While arguably less important for macroeconomic research, such an application could be very valuable for regulators and law enforcement. Appendix A.6 describes such a procedure in detail.

Looking forward, Bitcoin remains by far the largest cryptocurrency, but the same general approach can be applied to other high volatility cryptocurrencies. It may even be possible to adapt the approach to

[35] Similarly Wash trades, a fairly common and unregulated phenomenon in crypto markets (Cong et al, 2020; Le Pennec et al., 2021), can also be controlled, by assessing international capital flows exclusively.

certain classes of stablecoins, whose characteristics render them particularly well-suited for use as vehicles in capital transfers. For example, the leading stablecoin Tether (USDT) is measured to 8 decimal places, and although its price is far less volatile than Bitcoin, it still has experienced significant periodic price fluctuations. True, the lower volatility reduces the incentive to trade in and out quickly, but presumably most vehicle trades would still take place fairly quickly, given the risks. The *p* values might be higher for the matching algorithm, but it would not necessarily render the methodology inapplicable.[36] In any event, our analysis suggests that this approach could prove useful for both regulatory authorities and researchers.

VI: Conclusions

The results of this paper challenge the dominant view that Bitcoin is little used for transactions purposes (other than buying other cryptocurrencies), and that its value is almost entirely based on speculation. Exploiting extraordinarily rich and detailed public data on trades in 135 fiat currencies, we have presented what we believe to be the most systematic and detailed evidence to date on how crypto is used as a vehicle for international capital transfers, as well as for domestic transactions. Our main identification technique takes advantage of the fact that trades are reported to eight decimal places, and the data set gives a timestamp, price, and currency for each trade. To bolster our analysis, we have also looked an interesting natural experiment of a three-day power outage in the Venezuela, where crypto has been a popular vehicle for private citizens to evade draconian capital controls, and perhaps also for official entities to avoid US sanctions. We find that the power cut had a massive footprint not only in the Venezuelan Bitcoin market, but also in the Bitcoin markets of its major destination and remittance-sending countries. Indeed, although we do not know whether the Venezuelan experience can be extrapolated to other markets, it suggests that our estimate of share of trades that are crypto vehicle trades is very much a lower bound. As such, our work complements and considerably extends previous studies that indirectly attempt to

[36] A similar comment applies to those exchanges that report holdings only to six digits instead of eight digits.

identify transactions use based on network analysis of on-chain transactions, which do not include payment currency and are only a small share of the Bitcoin trade universe.

Even though crypto exchanges are subject to regulation in the country where they are based, off-chain transactions, particularly smaller ones, are difficult to systematically audit, especially where the regulator in one country needs information on a centralized exchange in another, Finland in the case of LocalBitcoins. We do not comment here on the future of Bitcoin regulation, but one can certainly infer from these results that any country aiming to institute or maintain capital controls will also need to find a way to prevent these from being circumvented via crypto (in addition to the plethora of "traditional" methods), and that regulation will be much more effective if there is widespread international cooperation.

If there are controls on crypto, these need to be included in any "new age" measure of controls on international capital movements. For example, the International Monetary Fund has recently argued for explicitly incorporating controls into fully defining a country's exchange rate regime (see also Ilzetzki, Reinhart and Rogoff (2019), Erten et al (2021) and Basu et al (2020)). If so, then differences in restrictions on transfers via cryptocurrencies increasingly need to be accounted for.

To do so, it will be important to extend the kind of analysis illustrated here to a wider range of crypto markets. Doing so will also give researchers, investors and policy makers a much better handle on "new age" international capital flows, both for purposes of determining the crisis vulnerability of individual countries (particularly EMDEs), as well as developing a deeper understanding of how international capital flows are impacting broader cross-country macroeconomic transmission.

Much more work will be needed to establish a metric for whether current crypto prices are too high or too low, given current regulation and expectations of future regulation. However, in parallel to Garber's classic (1989) analysis of the tulip mania, the transaction usage documented here should give one pause from drawing any strong conclusions yet on whether underlying payments (illicit or not) or international capital transfer use (in violation of capital controls or not) can eventually underpin prices.[37]

[37] Although our analysis suggests that crypto prices are not a pure speculative bubble, analyzing just one exchange does not allow us to infer anything about the level of Bitcoin prices, which at a bare minimum would require applying our algorithm to a much broader set of crypto markets. Obviously, as active as the developing transactional use of Bitcoin may be, the potential value can be quite high if regulators permit extensive transactions use, or cannot stop it. Suppose, for example, we use a back of the envelope estimate of the world underground economy to be 20% of global GDP, or roughly 18 trillion dollars (see Medina and Schneider (2018), Rogoff (2016) for further discussion). Then if, say Bitcoin, established itself as a major player in transactions alongside paper currency, and its transactions value were equal to just 0.5%, that would amount to 90 billion dollars per year.

References

Aizenman, J. (2008). On the hidden links between financial and trade opening. journal of International Money and Finance, 27(3):372 - 386.

Basu, M. S. S., Boz, M. E., Gopinath, M. G., Roch, M. F., & Unsal, M. F. D. (2020). *A conceptual model for the integrated policy framework*. International Monetary Fund.

Baur, D. G. and Dimp, T. (2021). The volatility of Bitcoin and its role as a medium of exchange and a store of value. Empirical Economics, pp. 1-21.

Chung, Y. (2019). Cracking the code: How the US government tracks Bitcoin transactions. *Anal. Appl. Math*, 152-168.

Cong, L. W., Li, X., Tang, K., and Yang, Y. (2020). Crypto wash trading. Available at SSRN 3530220.

Coppola, A., Maggiori, M., Neiman, B., and Schreger, J. (2020). Redrawing the map of global capital flows: The role of cross-border financing and tax havens. Technical report, National Bureau of Economic Research.

Erten, B., Korinek, A., & Ocampo, J. A. (2021). Capital controls: Theory and evidence. *Journal of Economic Literature*, *59*(1), 45-89.

Efron, B. (2007). Size, power and false discovery rates. *The Annals of Statistics*, *35*(4), 1351-1377.

Foley, S., Karlsen, J. R., & Putniņš, T. J. (2019). Sex, drugs, and Bitcoin: How much illegal activity is financed through cryptocurrencies?. The Review of Financial Studies, 32(5), 1798-1853.

Framewala, A., Harale, S., Khatal, S., Patel, D., Busnel, Y., & Rajarajan, M. (2020, April). Blockchain Analysis Tool for Monitoring Coin Flow. In *2020 Seventh International Conference on Software Defined Systems (SDS)* (pp. 196-201). IEEE.

Garber, Peter M. (1989, June). Tulipmania. Journal of Political Economy 97 (3), 535-60.

Haibo, H. J. L. (2008). An empirical study on determinants of hidden capital owing into china. World Economy Study, 6.

Ilzetzki, E., Reinhart, C. M., & Rogoff, K. S. (2019). Exchange arrangements entering the twenty-first century: Which anchor will hold?. *The Quarterly Journal of Economics*, *134*(2), 599-646.

Le Pennec, G., Fiedler, I., and Ante, L. (2021). Wash trading at cryptocurrency exchanges. Finance Research Letters, page 101982.

Medina, Leandro and Friedrich Schneider (2018). Shadow economies around the world: What did we learn over the last 20 years? *International Monetary Fund Working Paper* WP/18/17.

Rogoff, Kenneth (2016). *The Curse of Cash*. Princeton: Princeton University Press.

Ron, D., & Shamir, A. (2013, April). Quantitative analysis of the full Bitcoin transaction graph. In *International Conference on Financial Cryptography and Data Security* (pp. 6-24). Springer, Berlin, Heidelberg.

Schneider, B. (2003). Measuring capital flight: estimates and interpretations. Overseas Development Institute.

Wakefield, J. (2007). A Bayesian measure of the probability of false discovery in genetic epidemiology studies. The American Journal of Human Genetics, 81(2), 208-227.

Yang, L., Dong, X., Xing, S., Zheng, J., Gu, X., & Song, X. (2019, October). An abnormal transaction detection mechanism on Bitcoin. In *2019 International Conference on Networking and Network Applications (NaNA)* (pp. 452-457). IEEE.

Zhao, C., & Guan, Y. (2015, January). A graph-based investigation of Bitcoin transactions. In *IFIP International Conference on Digital Forensics* (pp. 79-95). Springer, Cham.

Appendix

Table A.1: Largest Cross-Border Capital Flows through Crypto Vehicle Trades

Origin Currency	Destination Currency	Trade Volume as share of total identified crypto vehicle trade volume Origin Currency (in %)
US Dollar	Nigerian Naira	24.3
Colombian Peso	Venezuelan Bolivar	24.8
Peruvian Sol	Venezuelan Bolivar	56.8
Euro	Venezuelan Bolivar	18.7
US Dollar	Venezuelan Bolivar	5.9
Chilean Peso	Venezuelan Bolivar	44.6
Argentine Peso	Venezuelan Bolivar	46.2
Mexican Peso	Venezuelan Bolivar	36.2
Malaysian Ringgit	Nigerian Naira	35.5
Balboa	Venezuelan Bolivar	37.2

Figure A 1 - Crypto trades in Currencies without significant vehicle trade volume to/from Venezuela around the power cut in Venezuela, compared to one currency (PEN) with significant vehicle trade volume with Venezuela.

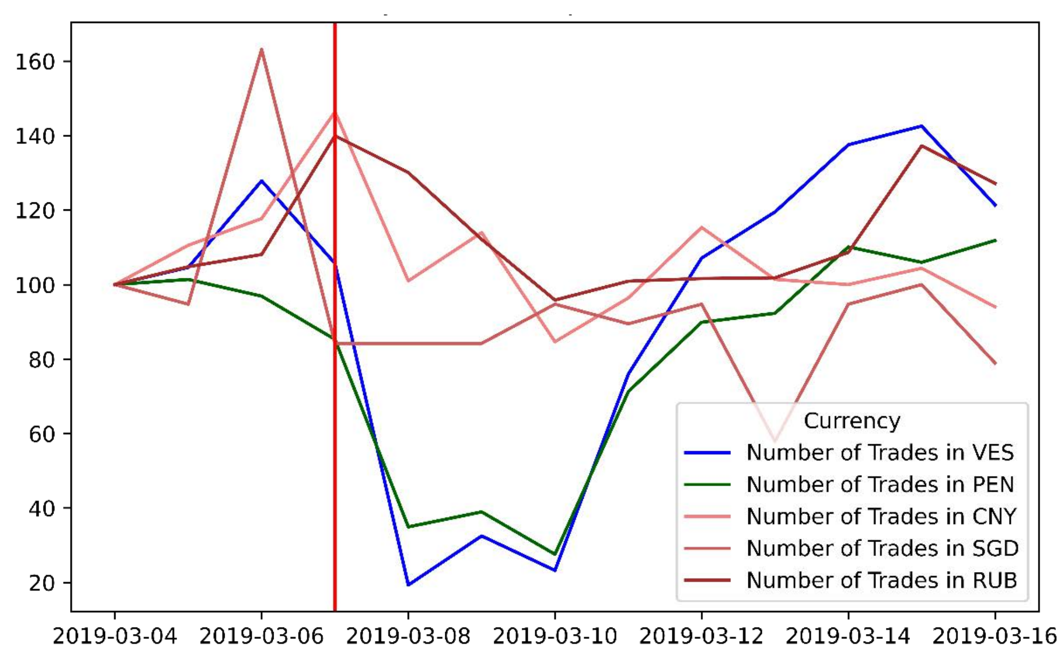

Sources: LocalBitcoins.com API, Authors' Calculations

Table A.2: Algorithm Output Example

Timestamp 1st trade	Currency 1st trade	Trade size x_i (1st and 2nd trade)	Timestamp 2nd trade	Currency 2nd trade	p_i	N_i
2020-11-01 01:12:43	USD	0.00202160	2020-11-01 02:03:431	VES	0.0000763	4083

Note: The matched trade presented in this table would not be considered a Crypto Vehicle trade, as the probability of the amount of 0.0020216 Bitcoin occurring, p_i in conjunction with the number of trades that occurred within five hours after the first trade, N_i , leads to a probability of this match being random that is greater than the 0.05 threshold.

A.1 Robustness Check - Applying different Time Windows

Table A.3: Robustness Check: Time Windows Compared

	(1) 2hr Window	(2) 5hr Window	(3) 10hr Window	(4) 24hr Window
Number of trades with one match	4 421 056	4 782 374	5 050 014	5 534 366
Number trades identified as vehicle-trade with (P(Match is Random) < 0.05), d_i	3 875 622	3 895 362	3 871 564	3 759 518
Number trades identified as vehicle-trade with (P(Match is Random) < 0.05), d_i net of False Discoveries from multiple hypothesis test, c_i	3 539 629	3 373 096	3 152 381	2 649 806
Share of trades identified as vehicle-trades, φ	7.7%	7.41%	6.9%	5.8%
Cross-border trades / identified vehicle-trades	15%	17%	20%	28%
Cross-border trades / identified vehicle-trades	25%	27%	30%	36%

when excluding RMB and RUB				
Cross-border trades / identified vehicle-trades (in USD Volume)	18%	20%	22%	30%

Figure A.2 - Time Windows Compared

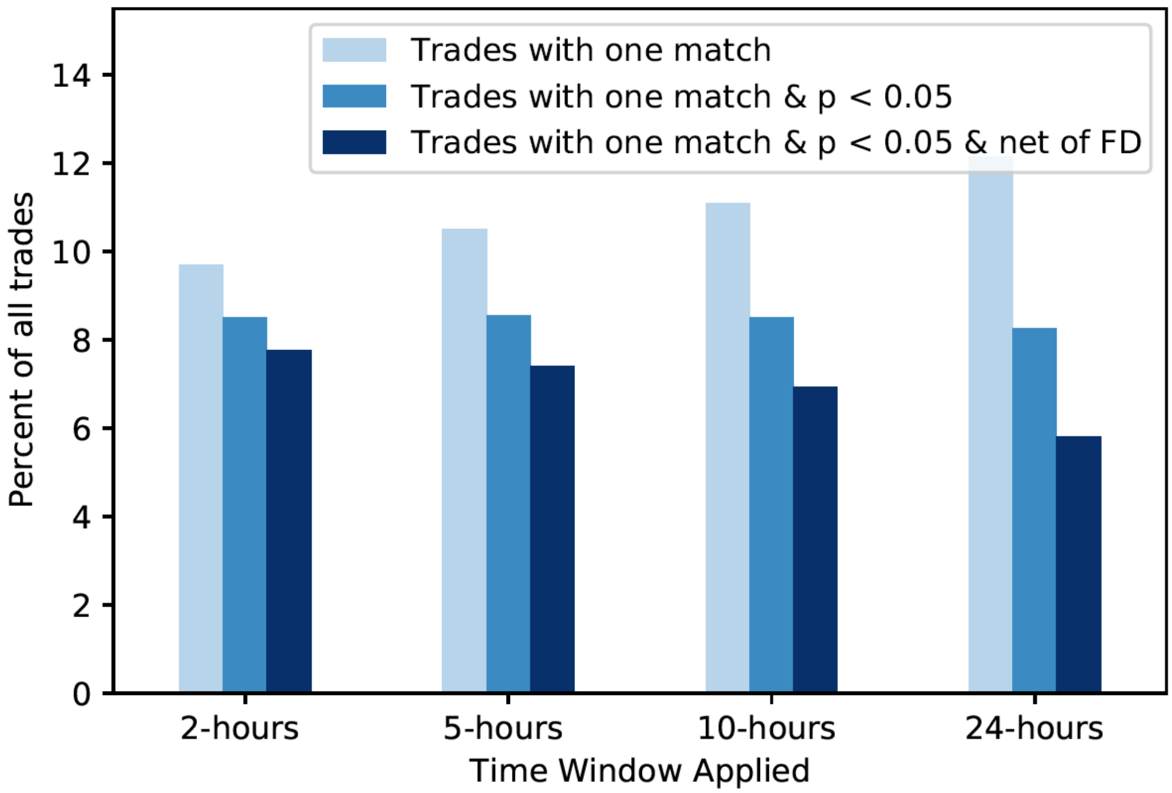

As Table A.3 and Figure A.2 show, our main result, that Bitcoin is used for crypto vehicle trades is robust to the selection of the time-window. However, the two illustrations also show that whereas the number of matched trades increases with longer time-windows considered to identify matching trades, the number of trades that happen within the time window considered, N_i, must increase also, meaning that ceteris paribus, θ_i^* increases, so that the number of trades that we identify as matched vehicle trades with at least a 95% confidence level eventually decreases. To choose a time window from the selection of time-windows applied, we applied a decision rule that imposed moving to the next longer time-window whenever the impact of matching new trades that were missed using the shorter time window is greater than the number of trades that are no longer considered in a longer time-window, because of the greater number of trades, N_i.

Formally, let $y(t)$ be the number of trades identified with a 95% confidence interval, as a function of the time-window, t, and define $z(t)$ as the number of trades that are disregarded, (although they are matched), because the individual test's $\theta_i^* > 0.05$. Then, our decision rule imposes choosing the next longer time window as long as:

$$\frac{dy}{dt} > \left|\frac{dz}{dt}\right|$$

A.2 Monte Carlo Simulation of trade matching in randomly drawn samples

Our vehicle trade identification algorithm relies on the assumption that identifying the same trade size twice within a short time window is unlikely to happen purely by chance, given the fact that trade sizes are specified to the 8th decimal and given the historical distribution of different trade sizes being sufficiently widely spread (i.e. the algorithm would not prove efficient in a world, where all trades were to be of the same trade-size, regardless how many decimals that number has). We support this assumption by running a Monte Carlo simulation, using a pair of randomly matched trade size and number of trades within five hours following the trade and equally as many trades randomly drawn from the real historical distribution, to analyze how many vehicle trades are identified as such purely by chance. This simulation proceeds as follows:

1. Based on the data since 2017 we record in our sample; we derive two quasi-random variables for the trade-size and number of trades occurring within five hours. We therefore define the trade-size

$$x_i \sim X(),$$

and the number of trades within five hours,

$$N_i \sim \Theta(),$$

where $X()$ and $\Theta()$ are probability distribution functions based on the historical distribution of trade sizes and number of trades occurring within five hours, respectively.

2. We then draw 1,000,000 random pairs of x_i and N_i from $X()$ and $\Theta()$.
3. For each pair (x_i, N_i), we further draw a random multiset, S_i with N_i elements from $X()$.

4. We count the number of instances, where x_i occurs in S_i exactly once,

$$\sum_{s \in S} 1_{s=x} \geq 2$$

and the probability

$$P\left(\sum_{s \in S} 1_{s=x} \geq 2 \right) = 1\ (-(\ 1 - p_i)^{N_i} < 0.05$$

Where p_i is the unconditional probability of x_i being drawn from X().

In short, we replicate the trade vehicle identification algorithm on randomly drawn subsets of the data, whereby the trade-sizes have been shuffled and randomly matched with numbers of trades within five hours. If it were true that rather than true instances of vehicle trades, our algorithm identified trades that randomly happen to have the same trade size, the share of trades identified as vehicle trades should be approximately the same in the Monte Carlo simulation and our algorithm's output. Instead, in the Monte Carlo simulation finds 1.1% of trades that find a match are whose individual hypothesis test leads to the conclusion that they are indeed crypto vehicle trades. This number is significantly lower than the 8.6% of the trades identified with the same methodology applied to the real data.

Of course, because the sample is randomly drawn, all of the vehicle trades identified within such a data set must be false positives. Applying the False-Discovery control we introduced in Section II should thus control for these false-positives and lead to a share of crypto-vehicle trades identified close to or equal to zero. And indeed, when deducting the share of expected false positives given the data structure in the Monte Carlo simulations, we arrive at an estimate of crypto vehicle trades equal to 0%.[38] Again, this result stands in stark contrast to the 7.4% of trades being identified as crypto vehicle trade (net of False Discoveries) found in the real data.

[38] More precisely, -0.06% of trades are identified as vehicle trades in the randomly shuffled data set. The estimate being negative thereby stems from the fact that the Monte Carlo simulation only considers instances with exactly one match, whereas the False-Discovery-Control per se also considers instances, where there is more than one match, meaning the latter is slightly greater than the former, meaning subtracting the latter from the former results in a number smaller than 0. Because it is extremely rare for trades to find two matches within a five-hour-time window, the impact of this difference on the results is negligible.

Figure A.3 - Monte Carlo Simulation result compared to real results

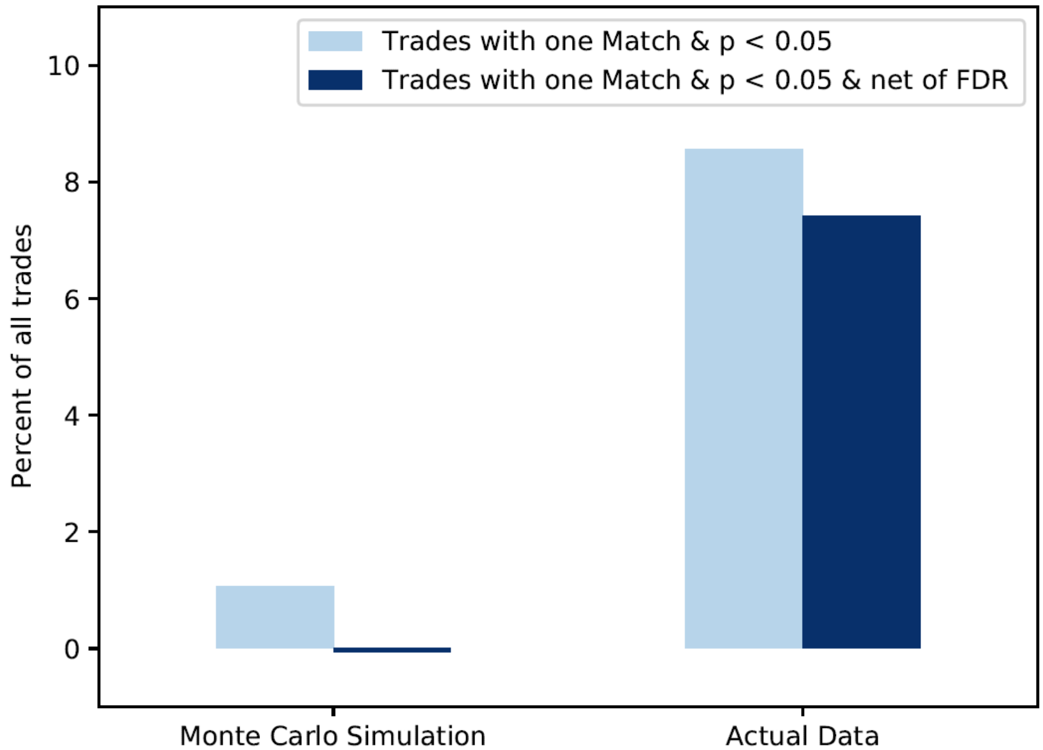

Figure A 4 - Monte Carlo Simulation using different time windows

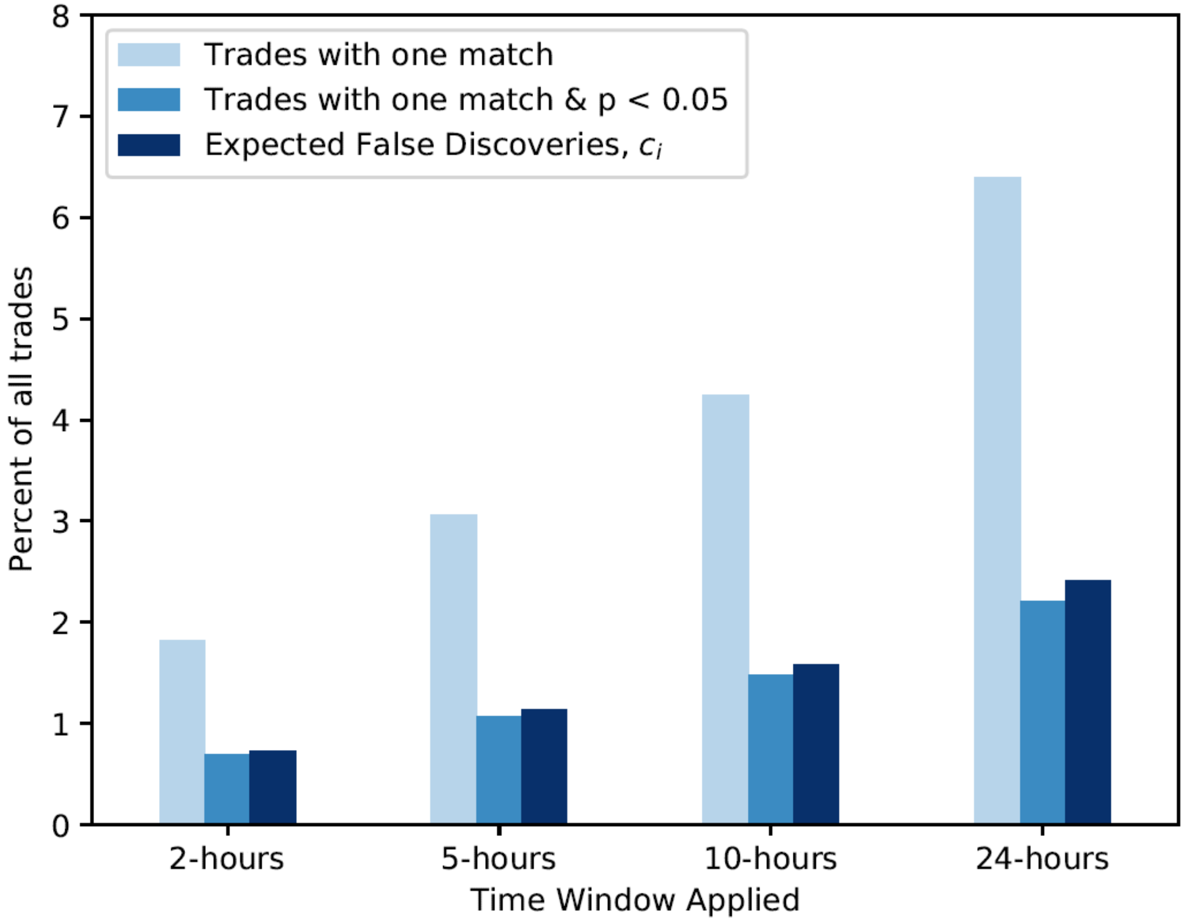

A.3 Time-series dynamic in trade-size distribution function

In our baseline crypto-vehicle-trade identification algorithm, we derive the probability of two trade-sizes matching from the historical distribution of trade-sizes. This is based on the observation that whereas certain trade sizes are common (e.g. 1 Bitcoin, 0.1 Bitcoin, etc.), many occur only once or twice within the whole sample. Observing the latter kind of trade twice within a short period of time is much stronger evidence for a vehicle trade than observing common trade-sizes that reoccur. For the historical distribution function, we apply one uniform distribution of all trade-sizes between 2017 and 2021. This relies on the assumption that the distribution remained fairly constant over time. A more precise approach would be to apply a time-series function to the trade-size distribution that is based on a rolling time window. We abstain from this approach here, because a preliminary and less computationally intensive analysis using fixed-annual windows for the historic distribution arrives at comparable estimates. This indicates that the computationally significantly more efficient approach applying a uniform trade-size distribution presents a close enough approximation.

Trade Size Distribution

2017

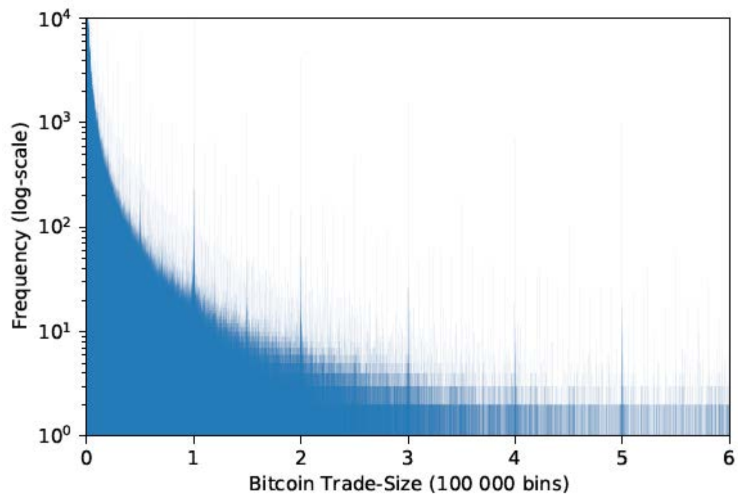

2018

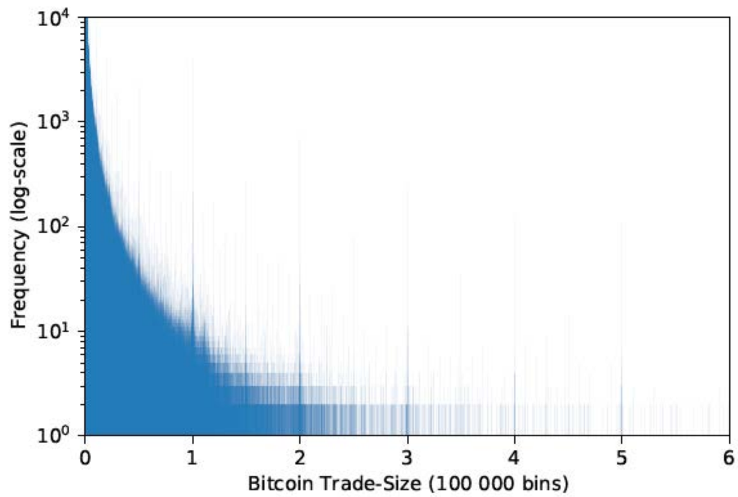

2019

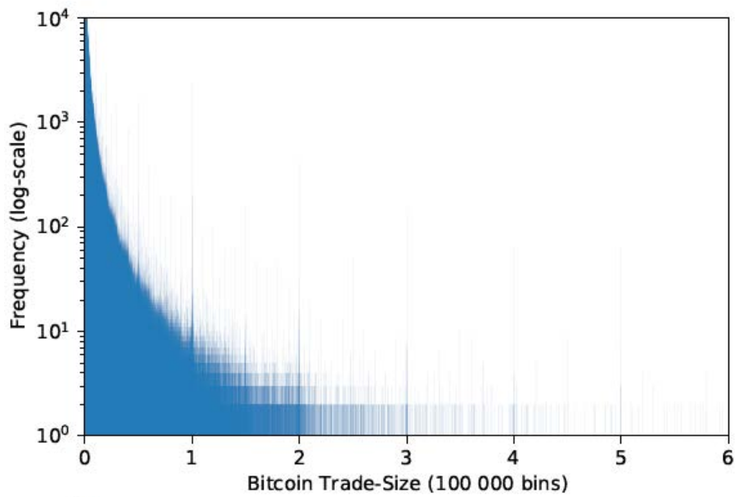

2020

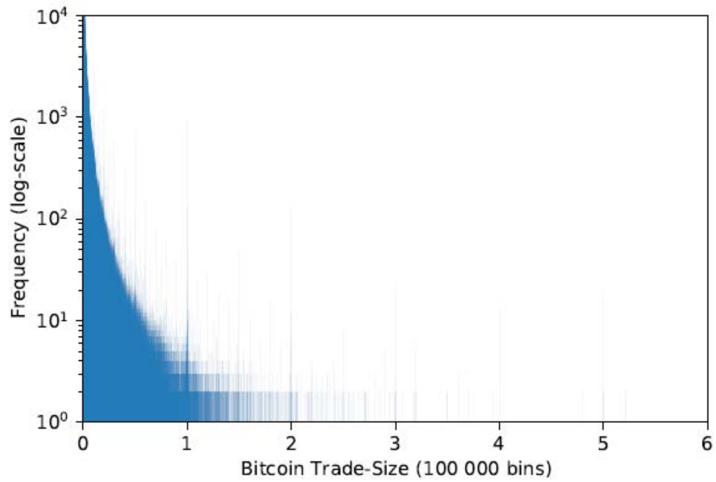

A.4 Crypto Vehicle Trade Volume Estimate

For the largest part of the paper, we ask what share of trades in our dataset are likely to be Crypto Vehicle Trades. An equally important question, especially for the capital flow related literature, would be the size of the transfer volume that is associated with crypto-vehicle-trading. The following extension of the methodology allows us to arrive at unbiased estimate of crypto vehicle trade volume, γ, while controlling for false discoveries.

Whereas Assumption 1 and Definitions 2 and 3 remain the same as in the original methodology, the "vehicle trades volume estimand" is now equal to

$$\gamma = \sum_{i=1}^{I} x_i \, \boldsymbol{\alpha}_i \; (\theta_i - \theta_i^*) \; with \; \boldsymbol{\alpha}_i = \begin{cases} 1 \; if \; \theta_i^* < \Theta_\theta \\ 0 \; if \; \theta_i^* \geq \Theta_\theta \end{cases}$$

We thus arrive at what can be thought of as a thresholding device, with the estimated volume of vehicle trades discovered being equal to:

$$\hat{\gamma} = \sum_{i=1}^{I} x_i \, (d_i - c_i)$$

Theorem 4 still holds, because both the estimator and the estimand were adjusted in the same fashion:

<u>Theorem 4</u>b

$$\mathrm{E}[\, \hat{\gamma} \mid N_{i,\dots,I} \,] = \varphi$$

Proof.

$$\begin{aligned} \mathrm{E}[\, \hat{\gamma} \mid N_{i,\dots,I} \,] &= \sum_{i=1}^{I} x_i \, (E\,[\, d_i \mid N_i] - c_i) \\ &= \sum_{i=1}^{I} x_i \, \boldsymbol{\alpha}_i \; (E\,[\, \boldsymbol{\phi}_i \mid N_i] - \theta_i^*) \\ &= \gamma \end{aligned}$$

A.5 Applying non-overlapping time windows

Applying overlapping time-windows, meaning a single time window for each trade, allows us to update the information for every trade, however it also means a departure from the strict independence between individual hypothesis tests, inter alia because a matching trade in the previous five-hour window would de facto reduce the trades of the following window by one. Although the large number of trades in each window renders the impact of the departure negligible, as a robustness check we provide an alternative approach, where we apply distinct and non-overlapping time-windows, thus guaranteeing the independence between distinct hypothesis tests.[39] The alternative algorithm differs from the original approach, in that we concentrate on separate, non-overlapping five-hour windows, and analyze the number of times each trade size occurs in these five-hour windows.

The definition of the null model, the estimand of the share of trades that are vehicle trades and finally the estimator thereof, are as those introduced in our original methodology, only that each arriving trade is defined by both its individual characteristics, and those of the five-hour block it occurs in. The proof of unbiasedness of the estimator in this adjusted methodology is analogous. And the proof of unbiasedness of the variance is equivalently straight forward.

As expected, given that vehicle trades that span over two time-windows are simply ignored, we find a slightly lower share of trades to be vehicle trade. Yet, just as in the approach applied to arrive at the main results of the paper, we find a crypto vehicle trade share that is significantly greater than zero.

[39] The authors are extremely grateful for Neil Shephard for suggesting this robustness test and suggesting the proof of unbiasedness.

Table A.4

	(1) Overlapping time windows	(2) Discrete time-windows
5-Hour Time Window	7.4 %	6.6 %
10-Hour Time Window	6.9 %	6.2 %
24-Hour Time Window	5.8 %	4.9 %

Given that matching trades which happen to fall in different time windows per se missed by the methodology using discrete time windows, the share of trades identified by this alternative application must be lower. This is in line with what is presented in Table A.4. More importantly however, the approach using discrete time windows allows us to assess whether the violation of the independence assumption and thus possible serial correlation of p-values has any significant impact on our results. Indeed, the results suggest that the impact, if any, is negligible.

A.6 Applying the methodology to data from different exchanges or the analysis of individual trade pairs

In principle, our methodology can be applied for any more targeted investigation in any particular country/region, as well as to any exchange that identifies trades in terms of the fiat currency used to purchase crypto, as long as a number of minimum conditions are fulfilled:

1. The exchange makes the trade data available in a non-aggregated fashion (either voluntarily making the data public or confidentially available to a researcher, or if the exchange's regulator requisitions the data).

2. The exchange allows trade sizes to be sufficiently precise. Some exchanges round the trade-size to six decimal points. This will structurally increase θ_i^* values, but by no means obviates the approach.[40]
3. The exchange has a transparent fee structure that either does not deduct fees from the nominal trade reported or does so in a transparent fashion allowing for an adjustment of the algorithm, e.g. to match trades that have the same size, net of fees.

Depending on further exchange specific factors, such as the fiat payment technology used, the algorithm might require other adjustments. For instance, as the execution of trades on centralized exchanges is usually almost instantaneous, examining data from these would require applying significantly shorter time windows.[41]

Whereas for economists and researchers in general, the aggregate, or the share of trades that constitute vehicle trades is most relevant, our methodology might also prove useful for regulators or law enforcement agencies, when wanting to assess, whether it is worth it to further investigate certain transactions. For such a use, the universal application of one given time window is likely to be less useful. We thus present an example of such an investigation of a single trade pair for illustrative purposes. Suppose a law enforcement agency is interested in capital movements between Malaysia and Panama. It finds in the data, that on September 16, 2017 at 03:12:30 am (GMT) there occurred a purchase of 2.5974026 Bitcoin, in Malaysian ringgit. Exactly 10 days later, on September 26th, 2017 at 01:19:44 pm (GMT), the exact same amount was sold for Panamanian balboa. Before requesting IP addresses from the exchanges involved, the law enforcement agency would like to assess the probability of this trading pattern occurring exclusively by chance. Over the 10 days between these two trades, 175,118 trades occurred. The probability of any given trade having the size of 2.5974026 Bitcoin is very low. Empirically speaking, only one in around 5.7

[40] How precise is sufficiently precise depends on the depth of the market and the time-windows one needs to consider (depending on the fiat payment mechanism for example), as the P-values are crucially dependent on the interaction of these three factors, which does not allow for general thresholds.

[41] The relevant minimum time required to complete a vehicle trade is easily assessed in an experimental fashion.

million trades would have this size. Together, the number of trades that occurred over this period and the probability of the given trade allow the calculation of the match being purely due to chance, using:

$$\theta_i^* \cong 1 - (1 - p_i)^{N_i}$$

In this case, the probability of this trade being random is around 3%. Let's imagine the law enforcement agency works with a rule to investigate any transfer between currency pairs of interest, when the probability of these being vehicle trades exceeds 80%. In such a case, the agency would likely proceed and request the IP addresses of all parties involved in the described transaction.

A.7 Data Appendix

The core data set makes use of data from the Application Programming Interface (API) published by LocalBitcoins.com. In principle, the data, made available in JSON format, goes back to the year 2013. However, the standards have seen some changes over the period prior to March 2017. Which is why, for our analyses, we concentrate on the 45.528.193 trades that occurred between March 15, 2017– July 23, 2021. [42] The trades are grouped and thus retrieved by fiat currency. For each of the observations we retrieve, there exists a unique trade id, the timestamp (converted to a UTC ISO format), the trade size, expressed in Bitcoin, and the price paid (expressed as the price of one Bitcoin in the given fiat currency). Because the trades are retrieved by fiat currency, we add that information to each observation. The fiat currencies included are: UAE Dirham, Afghani, Albanian Lek, Armenian Dram, Netherlands Antillean Guilder, Kenyan Kwanza, Argentine Peso, Australian Dollar, Aruban Florin, Azerbaijan Manat, Bosnia & Herzegovina's Convertible Mark, Barbados Dollar, Bangladeshi Taka, Bulgarian Lev, Bahraini Dinar, Burundi Franc, Bermudian Dollar, Brunei Dollar, Boliviano, Brazilian Real, Bahamian Dollar, Botswana

[42] LocalBitcoins has existed since 2012, however we limit our analysis to the period since March 2017, when the website revamped the exchange's back-end, guaranteeing consistency in the format of the data.

Pula, Belarusian Ruble, Belize Dollar, Canadian Dollar, Congolese Franc, Swiss Franc, Chilean Peso Chinese Offshore Renminbi, Chinese Yuan Renminbi, Colombian Peso, Costa Rican Colon, Peso Convertible, Czech Koruna, Danish Krone, Dominican Peso, Algerian Dinar, Egyptian Pound, Eritrea Nakfa, Ethiopian Birr, Euro, Fiji Dollar, Pound Sterling, Georgian Lari, Ghana Cedi, Gambian Dalasi, Guinean Franc, Guatemala Quetzal, Guyana Dollar, Hong Kong Dollar, Honduras Lempira, Croatian Kuna, Haiti Gourde, Hungarian Forint, Indian Rupiah, New Israeli Sheqel, Indian Rupee, Iraqi Dinar, Iranian Rial, Iceland Krona, Jamaican Dollar, Jordanian Dinar, Japanese Yen, Kenyan Shilling, Kyrgyz Som, Cambodian Riel, Korean Won, Kuwaiti Dinar, Cayman Islands Dollar, Kazakhstan Tenge, Lebanese Pound, Sri Lanka Rupee, Liberian Dollar, Lesotho Loti, Moroccan Dirham, Moldovan Leu, Malagasy Ariary, North Macedonian Denar, Myanmar Kyat, Macao Pataca, Mauritius Rupee, Maldives Rufiyaa, Malawi Kwacha, Mexican Peso, Malaysian Ringgit, Mozambique Metical, Namibia Dollar, Nigerian Naira, Cordoba Oro, Norwegian Krone, Nepalese Rupee, New Zealand Dollar, Rial Omani, Panama Balboa, Peruvian Sol, Papua New Guinea Kina, Philippine Peso, Pakistan Rupee, Polish Zloty, Paraguayan Guarani, Qatari Rial, Romanian Leu, Serbian Dinar, Russian Ruble, Rwanda Franc, Saudi Riyal, Seychelles Rupee, Sudanese Pound, Swedish Krona, Singapore Dollar, Saint Helena Pound, Surinam Dollar, South Sudanese Pound, Syrian Pound, Eswatini Lilangeni, Thai Baht, Tunisian Dinar, Turkish Lira, Trinidad and Tobago Dollar, New Taiwan Dollar, Tanzanian Shilling, Ukrainian Hryvnia, Uganda Shilling, US Dollar, Peso Uruguayo, Uzbekistan Sum, Venezuelan Bolívar Soberano, Vietnamese Dong, CFA Franc BEAC, East Caribbean Dollar, CFA Franc BCEAO, Yemeni Rial, South African Rand, Zambian Kwacha, Zimbabwe Dollar.[43] It is important to note that the transaction volume in each of these currencies can differ widely, with some currencies seeing but a handful of trades over the period we study.

[43] Additional to the 135 fiat currencies, the dataset includes three further non-traditional-fiat-currency means of payments: silver (11 trades), gold (three trades) and ethereum (9,127 trades), but as these represent less than 0.1% of all trades, we disregard in them in the further analysis.