recycle capital from one account to another, but in some exchanges this transfer can take up to several days. We discuss below some of the constraints to capital movements that reduce the speed of transferring capital across exchanges. While this strategy does not expose the arbitrageur to convergence risk, a drawback of this strategy is that the arbitrageur becomes exposed to bitcoin price fluctuations. To mitigate this risk she can establish short positions either on some of the exchanges described above or by borrowing bitcoin from people who hold large amounts of bitcoins without an interest to sell, the so-called *hodlers*.[17] Of course, these hodlers themselves would be in a great position to do the arbitrage in cryptocurrency markets. Starting from the end of December 2017, the arbitrageur can also use Chicago Board Options Exchange (CBOE) and Chicago Mercantile Exchange (CME) bitcoin futures contracts to hedge the price risk. The futures contracts track bitcoin price, on major US dollar exchanges and have an average daily open interest of about 10,000 bitcoins.[18]

8.2. Constraints to arbitrage

In practice, the arbitrageur has to incur a number of transaction costs, but their magnitudes are too small to prevent arbitrageurs from implementing the above trading strategies. To transfer bitcoins, the transaction has to be recorded on the Bitcoin blockchain; this is the work of the so-called miners that provide certification of transactions and add blocks to the blockchain if they win the hashing competition. The fees peaked around $40 in the end of December 2017 at the height of the bitcoin price but since February have come down to below $10. Since these are fixed cost, they are minuscule relative to the size of the potential arbitrage. In addition, exchanges have trading fees, which increase the cost of trading. In the appendix we show the magnitude of the fees for the exchanges used in this paper. These fees range from 0.25% of the amount traded to 0.1%. Most exchanges do not charge fees on a trade by trade basis but assign them based on the trading volume in a given month or week. Furthermore, most exchanges charge zero fees for trades that add to the liquidity of the order book. The exchange fees are comparable to the bid-ask spreads, which are, on average, between 1 and 10 basis points. Finally, many exchanges charge withdrawal

[17]The term hodler is a peculiarity of the bitcoin market since one investor in bitcoin wrote in a post on the bitcoin talk forum in 2013 while prices were dropping "*I AM HODLING*". This has become a meme for "Hold On for Dear Life".

[18]The CBOE contract settlement price is determined by results of the auction on the Gemini exchange. The CME contract settlement price is based on the CME CF bitcoin Reference Rate (BRR), which aggregates the price from major US dollar exchanges (see http://www.cmegroup.com/trading/equity-index/us-index/bitcoin.html for more details.)

In Fig. 11 we plot the ratios of the value weighted price of bitcoin to ethereum at the minute level across different regions. This is very similar to the calculations in Section 4.3. However, here we make two modifications to the process. First, we calculate the volume-weighted price at the five-second level since trading in ethereum has less volume than bitcoin. Second, since not all exchanges directly provide a platform to trade bitcoin to ethereum, but often trade bitcoin and ethereum only to the local fiat currency, we calculate the local exchange rate of bitcoin to ethereum as the cross-rate. For example, to calculate the exchange rate of ethereum to bitcoin on Bithumb, we take the ratio of the exchange rate of ethereum to Korean won to the exchange rate of bitcoin to Korean won.

[Fig. 11 About Here]

Fig. 11, Panel A shows the ratio of the two exchange rates of bitcoin to ethereum between the US and Japanese exchanges from November 1, 2017, to February 28, 2018. As mentioned before, if there were no frictions in the currency markets, then this ratio should be constant and equal to one all the time. We see that the price ratio indeed lies in a narrow band between 0.98 and 1.02.

We repeat the same calculations in Fig. 11, Panel B for the price of bitcoin to ethereum, but for the US and Korea. Again, we see that the deviations from the ratio of one are relatively small, and even in December and January it hovers around 1.03 to 0.97. In comparison, these were the months when the Kimchi premium (price of bitcoin in fiat currency in Korea versus US) was highest, almost 50% for several days. This again confirms that the arbitrage opportunities are much less pronounced and persistent between different cryptocurrency markets than between cryptocurrency and fiat currency markets. We finally repeat the same exercise in Fig. 11, Panel C for the US and Europe and find that the differences in the price of BTC to ethereum across all three months are small—less than 1% for the average day—and there are only several days in mid-December where the ratio is around 1.03. However, this lack of price dispersion should not be too surprising since we have previously shown that, even in the bitcoin to fiat currency market, the difference between US and European exchanges is smaller than in other regions.

7. Order flow and prices

To provide an estimate of how much capital is required to close the arbitrage spreads we show above, we develop and estimate a model of order flows and prices. The existing literature shows the importance of net order flows for price formation in traditional

6. Arbitrage in other cryptocurrency markets

To further investigate if controls on fiat currencies play an important role in explaining arbitrage spreads, in this section we analyze the price of ethereum and ripple relative to fiat currencies and then relative to bitcoin. If capital controls on fiat currencies play an important role, we would expect to find large arbitrage spreads between ethereum to fiat currencies but much smaller arbitrage deviations between two cryptocurrencies.

Fig. 10 plots the arbitrage index for ethereum and ripple. We can see that, similar to the bitcoin arbitrage index, there is significant variability during the year. Periods of relatively low levels of the arbitrage index alternate with prolonged spikes. Similar to bitcoin, at the height of its peak December and January, the ethereum arbitrage index stays at about 1.5. The ripple index displays similar behavior, but the series only starts from August 2017 due to the data availability.

[Fig. 10 About Here]

As in the case of bitcoin, a significant part of the arbitrage spread in ethereum and ripple prices is driven by price deviations across geographic regions. We do not report this analysis in the paper since they are very similar to the reported arbitrage indices for bitcoin, but they can be obtained from the authors on request. Furthermore, by comparing the three arbitrage indices, one can notice the high degree of correlation between them. All three arbitrage indices usually spike at about the same time and take similar levels.

6.1. Arbitrage between cryptocurrencies

To analyze if the same arbitrage spreads exist between cryptocurrencies, we focus on ethereum as the second-most traded cryptocurrency after bitcoin.[15] We only look at the months of November 2017 to February 2018 since these are the time periods when the BTC price has the strongest price dispersion relative to fiat currencies. It is also the time period when trading in ethereum and other coins become more liquid. If constraints in the movement of capital contributes to the arbitrage profits between BTC and the local fiat currencies, then these price deviations should be much smaller across cryptocurrencies, which, by design, do not obey the same restrictions.

[15]We run the same analysis for ripple and obtain qualitatively similar results, which are available upon request from the authors.

and transaction mechanism. Ciaian, Rajcaniova, and Kancs (2016), Harvey (2016), Bohme et al. (2015), and Raskin and Yermack (2017) provide a broad perspective on the economics of cryptocurrencies and the blockchain technology they are built upon. Athey et al. (2016) and Pagnotta and Buraschi (2018) propose models of the valuation of digital currencies. Cong, He, and Li (2019), Easley, O'Hara, and Basu (2017), and Huberman, Leshno and Moallemi (2017) study Bitcoin mining fees and the incentives of miners in equilibrium. We view our paper as complementary to this literature. To our knowledge, we are the first to provide a systematic empirical study of trading and price formation in cryptocurrency markets using transaction-level data.

Our paper is also linked to the limits of arbitrage, which argues that prices can deviate from law of one price even in the presence of arbitrageurs; see, e.g., DeLong et al. (1990), Gromb and Vayanos (2002), and Gromb and Vayanos (2018). On the empirical side, our paper is closest to the studies that analyze deviations from one price in different markets. In particular, Rosenthal and Young (1990) and Froot and Dabora (1999) study "Siamese twin" companies. They show that prices of two types of shares, which are traded in different markets but have identical claims on the cash flows and assets of the same company, can nevertheless substantially deviate from each other. Similar to Rosenthal and Young (1990) and Froot and Dabora (1999), we show that bitcoin and other cryptocurrencies can be traded at vastly different prices on different exchanges. The deviation from the law of one price is even more striking in the case of cryptocurrencies, since unlike shares that are traded within specific countries, bitcoins can be transferred to any market. As a result, typical explanations such as tax-induced investor heterogenity or index membership do not apply in this case. On a broader level, our paper is also linked to the market segmentation literature; see, e.g., Bekaert et al. (2011). Similar to this literature, our results suggest that capital controls and the development of financial markets can be important at explaining the differences in the marginal valuation of investors across countries.

Finally, our paper is also related to research that shows a strong positive relation between asset prices and net order flow in "traditional" financial markets. For example, Evans and Lyons (2002), Berger et al. (2008), and Fourel et al. (2015) look at foreign exchange markets; Brandt and Kavajecz (2004) at US Treasury markets; Deuskar and Johnson (2011) at the S&P 500 futures market; and Chordia, Roll, and Subrahmanyam (2002), Goyenko, Holden, and Trzcinka (2009), and Hendershott and Menkveld (2014) for NYSE stocks. These papers suggest that order flow imbalances typically explain about 15%-30% of the day-to-day variation of stock returns or treasury yields and up to 50% of foreign exchange returns. We show that a very strong positive relation exists in cryptocurrency markets as well. But the R-squared that we show for cryptocurrency

4022---dogecoin--.pdf

Thus, from January 1, 2016, to February 28, 2018, the return on bitcoin is about 900%.

Table 3 shows the higher moments of bitcoin returns at the daily, hourly, and five-minute level from January 1, 2017, to February 28, 2018. These statistics are calculated by averaging the corresponding moments across all available exchanges. For each frequency we report the annualized standard deviation, skewness, and kurtosis of returns, as well as the autocorrelation and cross-correlation across exchanges.

Column (1) of Table 3 reports the standard deviation. We see that the volatility of returns is very high. Even at the daily frequency, the annualized standard deviation is 107%. In comparison, the annualized standard deviation of Nasdaq, from 1985 to 2017, is 18%. However, the kurtosis at the daily frequency is 3.86, which is not too far from that of the normal distribution. The daily returns are positively skewed, which is perhaps not very surprising given the steep increase in the price of bitcoin over the considered time period. Columns (4) through (6) show the autocorrelation in returns for one, two, and three lags. We can see that even at the five-minute frequency the autocorrelations are small, which shows that there is little predictability in the market.

Finally, in column (7) we report the average cross-correlation of returns. We average across all the pairwise correlations but take out the diagonal (i.e., the autocorrelation of an exchange with itself). We see that at the five-minute level the correlation between exchanges is quite low, only 57%, while at higher frequencies the correlation increases: it is 83% at the hourly level and 95% for daily returns. These results are similar to what is observed in other well-established markets; see, for example, Budish, Cramton, and Shim (2015). However, if in equity markets the break of correlations happens at millisecond frequencies, then here it is already present at the minute levels.

The lower correlations at higher frequencies point to the existence of price deviations between exchanges. In the following section we investigate the existence of arbitrage opportunities in more detail.

4. Arbitrage

4.1. Arbitrage index

The low cross-correlation in returns across exchanges, which we computed in the previous section, already suggests that the crypto market is far from being efficient. To show the amount of price dispersion between exchanges at a given point in time, we form an arbitrage index that compares the maximum difference in prices between exchanges. We start by calculating this arbitrage index at the minute level. For this purpose, for a given minute we first compute the volume-weighted average price in that liquid exchanges, such as the US, Japan, Korea, and, to a lesser extent, Europe. We calculate that the daily average price ratio between the US and Korea from December 2017 until the beginning of February 2018 was more than 15% and reached 40% for several days. This has been noted in the popular press as the "Kimchi premium". Similarly, the average price difference between Japan and the US was around 10%, and between US and Europe about 3%. To provide a sense of the magnitude of the money left on the table, we calculate the daily profits that could have been achieved in this market. The daily amount of potential arbitrage profits was often more than $75 million, and in the period from December 2017 to February 2018 we estimate a minimum of $2 billion of potential total arbitrage profit.[2] In contrast, the price deviations between exchanges in the same country typically do not exceed 1%, on average.

Third, we find that deviations in bitcoin prices across countries are highly asymmetric. In countries outside the US and Europe, bitcoin typically trades at a premium relative to the US and almost never at a price below the US. In addition, there is significant co-movement in price deviations across countries: arbitrage spreads open up and close at the same time across countries.

Fourth, our analysis shows that price deviations occur during periods of a particularly quick appreciation of bitcoin prices. Since we show later that bitcoin prices react strongly to order flows, these periods also coincide with the times when there is a particularly strong increase in demand for bitcoin worldwide. To construct a measure of "buying pressure" in bitcoin markets, we take the difference between the actual log price of bitcoin in the US and its trend component, which we estimate using the Hodrick-Prescott filter. The bitcoin price in the US is a good proxy for the world market price of bitcoin. We then regress the deviations of a country's bitcoin price relative to the US on our measure of "buying pressure". This gives us a measure of the sensitivity of a country's bitcoin price to changes in the world market price of bitcoin. We call this the bitcoin beta of a country. We show that the countries that, on average, have a higher premium over the US bitcoin price are also those with a higher bitcoin beta. So these countries respond more strongly in widening arbitrage deviations in times when buying pressure goes up in the US.

Our results thus show that the marginal investor outside the US and Europe is willing to pay more for bitcoin in response to positive news or sentiment. How can one explain this differences in valuation? We conjecture that they might reflect tighter

[2]Our approach allows us to abstract from any assumptions about price impact of additional arbitrage trades or the speed of convergence. Since we rely on trades that were executed on the exchanges, it also eliminates concerns about stale prices or illiquid exchanges. See Section 4.4 for details.

References

Athey, S., Parashkevov, I., Sarukkai, V., and Xia, J., 2016. Bitcoin pricing, adoption, and usage: theory and evidence. Unpublished working paper. Stanford University.

Baba, C., Kokenyne, A., 2011. Effectiveness of capital controls in selected emerging markets in the 2000s. Unpublished working paper. IMF.

Bekaert, G., Harvey, C. R., Lundblad, C. T., Siegel, S., 2011. What segments equity markets? Review of Financial Studies 24, 3841–3890.

Berger, D., Chaboud, A., Chernenko, S., Howorka, E., Wright, J., 2008. Order flow and exchange rate dynamics in electronic brokerage system data. Journal of International Economics 75, 93–109.

Bohme, R., Christin, N., Edelman, B., Moore, T., 2015. Bitcoin: economics, technology, and governance. Journal of Economic Perspectives 29, 213–238.

Brandt, M. W., Kavajecz, K. A., 2004. Price discovery in the US Treasury market: the impact of order flow and liquidity on the yield curve. Journal of Finance 59, 2623–2654.

Budish, E., Cramton, P., Shim, J., 2015. The high-frequency trading arms race: frequent batch auctions as a market design response. Quarterly Journal of Economics 130, 1547–1621.

Chordia, T., Roll, R., Subrahmanyam, A., 2002. Order imbalance, liquidity, and market returns. Journal of Financial Economics 65, 111–130.

Ciaian, P., Rajcaniova, M., Kancs, D., 2016. The economics of bitcoin price formation. Applied Economics 48, 1799–1815.

Cong, L. W., He, Z., Li, J., 2019. Decentralized mining in centralized pools. Unpublished working paper. NBER.

DeLong, J. B., Shleifer, A., Summers, L. H., Waldmann, R. J., 1990. Noise trader risk in financial markets. Journal of Political Economy 98, 703–738.

Deuskar, P., Johnson, T. C., 2011. Market liquidity and flow-driven risk. Review of Financial Studies 24, 721–753.

Easley, D., O'Hara, M., Basu, S., 2017. From mining to markets: the evolution of bitcoin transaction fees. Unpublished working paper. Cornell University.