

An Experimental Study of Cryptocurrency Market Dynamics

Peter M. Krafft*
MIT CSAIL
Cambridge, MA, USA
pkrafft@csail.mit.edu

Nicolás Della Penna*
Australian National University
Canberra, Australia
me@nikete.com

Alex “Sandy” Pentland
MIT Media Lab
Cambridge, MA, USA
pentland@mit.edu

ABSTRACT

As cryptocurrencies gain popularity and credibility, marketplaces for cryptocurrencies are growing in importance. Understanding the dynamics of these markets can help to assess how viable the cryptocurrency ecosystem is and how design choices affect market behavior. One existential threat to cryptocurrencies is dramatic fluctuations in traders’ willingness to buy or sell. Using a novel experimental methodology, we conducted an online experiment to study how susceptible traders in these markets are to peer influence from trading behavior. We created bots that executed over one hundred thousand trades costing less than a penny each in 217 cryptocurrencies over the course of six months. We find that individual “buy” actions led to short-term increases in subsequent buy-side activity hundreds of times the size of our interventions. From a design perspective, we note that the design choices of the exchange we study may have promoted this and other peer influence effects, which highlights the potential social and economic impact of HCI in the design of digital institutions.

ACM Classification Keywords

J.4 Social and Behavioral Sciences: Economics; K.4.2 Social Issues; H.5.m. Information Interfaces and Presentation (e.g. HCI): Miscellaneous

Author Keywords

Cryptocurrencies; online markets; peer influence; computational social science; online field experiments; digital institutions; design; market design

INTRODUCTION

Cryptocurrencies, a.k.a. “cryptocoins”, are rapidly gaining popularity. The price and market cap of these assets are touching all-time highs, with billions of U.S. dollars of value per day currently being traded in cryptocoins. Financial institutions are investing in building digital currency technologies. Blockchain-based tech startups are thriving. As these changes in the cryptocurrency ecosystem occur, the need to understand the market dynamics of cryptocoins increases. At the same

time as cryptocurrencies have gained popularity, their rise has been punctuated by crises. From the collapse of Mt. Gox in 2014 to the 2016 hack of Ethereum, market crashes have been a regular occurrence. Understanding the dynamics of cryptocurrency markets may allow us to anticipate and avoid future disruptive events.

One potential threat to cryptocurrencies derives from the speculative nature of these assets. Many participants in these markets trade because they expect one or another cryptocurrency to increase in value. Such collective excitement can lead to bubbles and subsequent market crashes. The design choices of the online exchanges where cryptocurrencies are traded may also contribute to these effects if aspects of available functionality, graphical user interfaces (GUI), or application programming interfaces (API) promote collective excitement. Markets are human artifacts, not natural phenomena, and therefore a target of design [53]. In the present study we strive to better understand the factors that contribute to collective excitement in cryptocurrencies, and how the design of cryptocurrency market mechanisms and interfaces may affect these processes.

Central to these goals is understanding why people at a particular time decide to invest in a particular technology, product, or idea. If the asset is new, or information about it has just been released, investment might be a rational response to the present state of information [9]. Other factors could include authorities endorsing the investment, or big players making noticeably large bets on it [8]. Another hypothesized source of collective optimism is peer influence among small individual traders [12, 78, 41]. Understanding these endogenous peer influence effects is especially important. If the dynamics of financial markets are heavily affected by small trades, different solutions may be needed in order to stabilize the markets.

Peer influence may play a particularly large role in the cryptocurrency ecosystem due to the highly speculative nature of these assets. While in general the intrinsic value of currencies increases with greater levels of adoption, here we expect that much of the trading is speculative. As is characteristic of many new technologies, there is a great deal of uncertainty around which cryptocurrencies will eventually be successful, and there has been a general feeling that some altcoin will become a transformative financial technology. Many participants in these marketplaces therefore are likely hoping to be early investors in “the next Bitcoin”, and we can attempt to observe the extent to which the bets of these traders may be affected by the bets of their peers.

* Authors contributed equally. Corresponding author: NDP.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI 2018, April 21–26, 2018, Montreal, QC, Canada

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-5620-6/18/04...\$15.00

DOI: <https://doi.org/10.1145/3173574.3174179>

There are several challenges to identifying the effects of small individual trades in financial markets. Much of our present knowledge about financial markets is derived from analysis of observational data, but observational analyses are subject to confounding interpretations. For example, excess correlation in market prices is often cited as evidence against traders engaging in the purely rational behavior predicted by the efficient market hypothesis [57, 59], but these correlations could be due to delayed reaction or overreaction to news events, as well as perhaps due to peer influence. Experimental evidence is desirable because of these difficulties with observational data. Laboratory studies have been conducted in which causal inferences can be made (e.g., [77, 2, 22, 64]), but these studies are limited in their capacity for generalization to real financial markets. Since markets are noisy, and the effect of individual trades is likely to be small, a field experiment in this area requires a large sample size, which would be expensive to collect for scientific purposes in traditional financial markets.

We overcome these challenges using a field experiment in an online marketplace for cryptocurrencies. Cryptocurrency markets provide a unique opportunity for field experimentation due to their low transaction fees; low minimum orders; and free, readily accessible public APIs. We created bots to trade in 217 distinct altcoin markets in an online exchange called Cryptsy. Each bot monitored a market and randomly bought or sold the market's associated altcoin at randomly spaced intervals. By comparing our buy and sell interventions to control trials we can estimate the effects that our trades had on the dynamics of these markets. In total we conducted hundreds of thousands tiny trades in these markets over the course of six months, and this large sample size allows us to test the effects that our small individual trades have in these live markets. With an eye towards design implications, we also conduct an enumerative analysis of the importance in this context of Cryptsy's design choices. We identify the position of Cryptsy in a space of existing and potential exchange designs, and discuss the possible effects of the dimensions of its position. Our analysis reveals that the traders we study are susceptible to peer influence and highlights how Cryptsy's design choices might have exacerbated this effect.

CRYPTOCURRENCY MARKETS

Cryptocurrencies

Cryptocurrencies are a new type of digital asset that rely on distributed cryptographic protocols, rather than physical material and a centralized authority, to operate as currency. Bitcoin (BTC) was the first cryptocurrency to gain popularity, but hundreds of alternative cryptocoins (called "altcoins") have since been introduced. The current crop of altcoins has been directly inspired by Bitcoin, and the excitement about Bitcoin frames the hopes and desires of participants in the marketplace for altcoins. In discussions by those who create, promote, and scour such coins, a desire to not miss out again on being an early investor in the "next Bitcoin" is commonplace. However, many altcoins represent nothing more than minor changes to the source code of Bitcoin. While it might be tempting to dismiss all such coins as having essentially zero probability of success, some of them do innovate in non-technical ways. One case is Auroracoin, which was a trivial technical modification

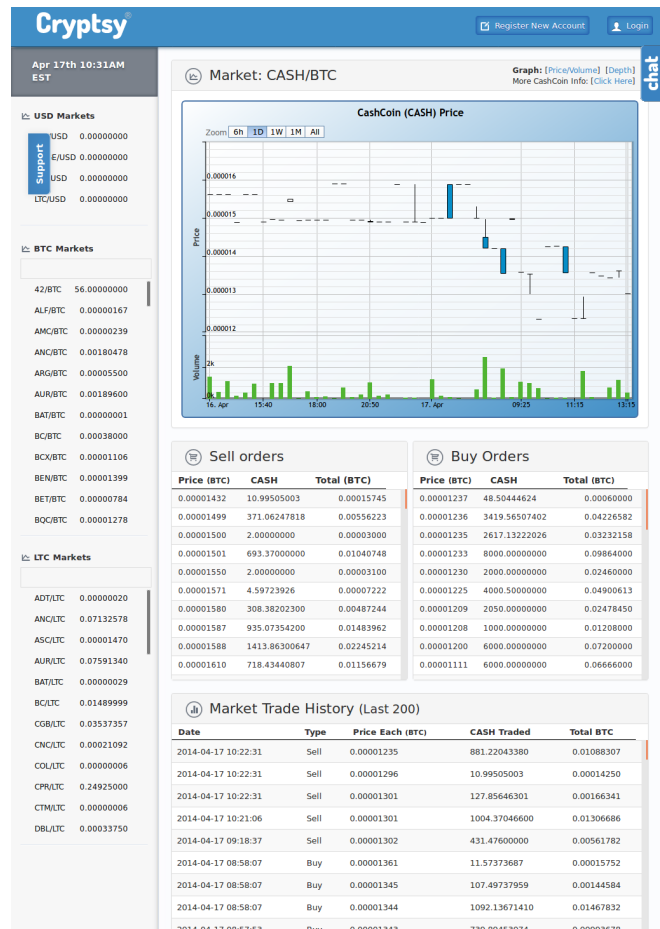


Figure 1: A screenshot of the Cryptsy trading platform interface. Source: <http://archive.is/tDF1w>.

to Litecoin but was branded as the official cryptocurrency of Iceland. Auroracoin at one point had a market cap of 500 million USD. Other coins are associated with real technical innovations. For example, Ether is used by the Ethereum protocol in order to implement a fully functional distributed global computer. Evaluating the prospective returns from investing in any of these coins is difficult and time-consuming, requiring expertise in both cryptography and economics.

Cryptsy

The platform we use for our experiments, Cryptsy, was a large cryptocurrency exchange that opened on May 20, 2013 and closed in early January 2016. At the outset of our experiment, Cryptsy claimed over 230,000 registered users. On the last recorded day of Cryptsy's trading, its daily trading volume was 106,950 USD (248 BTC), which placed it as the tenth largest cryptocurrency exchanges by trading volume (of 675 listed) at the time. By this time Cryptsy had 541 trading pairs (including Bitcoin, Litecoin, and fiat markets), which placed it as the third largest exchange in terms of the total size of its marketplace. Cryptsy was a popular exchange because of the large number of altcoins it made available for trade, which is also the reason it is uniquely appropriate for our experiment.




Figure 2: Descriptive statistics of the Cryptsy marketplace.

We require a popular exchange with many coins available. By the end of 2015, Cryptsy had begun having well-publicized issues involving users not being able to withdraw money they had deposited on the site. On January 14, 2016 Cryptsy halted all trading. Our experiment—spanning April 12, 2015 until October 19, 2015—preceded the beginning of the final decline of Cryptsy.

Market Mechanism

Cryptsy implemented a continuous double auction with an open order book as its trading mechanism. In a continuous double auction there is no centralized market maker. Current asset prices are determined by the best available prices being offered by any of the traders on the platform. A market for a particular cryptocurrency consists of a set of open “buy orders” and a set of open “sell orders”, all placed by peers on the site. A buy order is a request to buy a quantity of a coin at a price specified in the order. A sell order is a request to sell. The current price to buy, i.e., the “buy price”, is given by the lowest priced open sell order. The current “sell price” is given by the highest priced open buy order. A transaction occurs when a new buy order is placed with a price at or above the lowest sell price, or when a new sell order is placed with a price at or below the highest buy price. The minimum denomination on Cryptsy was $1e-8$. At a typical exchange rate of approximately 200 USD to 1 BTC during the time of our experiment, $1e-8$ BTC corresponded to roughly 0.000002 USD. The minimum total order size for a single trade varied across coins and over time, but was typically $1e-7$ BTC, and the transaction fees were negligible.

Interfaces

Cryptsy had both a graphical user interface (GUI) and an application programming interface (API). The GUI, pictured in Figure 1, allowed users to browse through USD, BTC, and Litecoin (LTC) markets. The site allowed users to view a standalone list of all the markets, which could be sorted by volume or price (recent price, 24-hour high price, or 24-hour low price), and the site allowed users to view details of the specific coin markets. All current coin prices for each of those markets were displayed in panels on the left-hand side of the

screen in alphabetical order by the coin name. For the BTC and LTC markets, these panels also displayed green or red marks when the prices of each coin had recently moved up or down. These marks accumulated during an idle web session. The single most recent BTC and LTC market price changes were displayed at the tops of these panels.

In addition to showing all of the open buy and sell orders, Cryptsy also made prior transactions visible to users. Prior transactions were shown both in a list and as a chart. The list included the sizes, prices, and times of the last 200 transactions, reverse ordered by time. The chart displayed an interactive summary of prior transactions. Each “tick” on the chart visualized the highest transacted price, the lowest transacted price, the open price, and the close price over the duration of a certain time interval. The total traded volume over the time duration of each tick was also shown. The time interval used for display depended on the time granularity at which the chart was being viewed. The chart could display the price trends over the last 6 hours, day, week, or month, or over all time.

Descriptive Statistics

Figure 2 summarizes various descriptive statistics of the Cryptsy marketplace. We observe a mean daily trading volume of approximately 400 BTC per day on Cryptsy. This quantity places the mean daily trading volume on Cryptsy to be at least in the tens of thousands of USD per day during our experiment. The mean daily trading volume remained relatively constant over the course of our experiment, with some periods of higher volume. There was substantial heterogeneity in the volume of each coin. Most coins have only on the order of 10 or 100 USD being traded in their markets per day, while a few have tens of thousands. The average size of observed trades also varies widely across coins. Across all coins, average trade sizes tend to be in the range of tens of cents to a few dollars. The maximum trade size we use in our experiments of $5e-6$ BTC is in the bottom 8% of the distribution of observed trade sizes, while our minimum trade size of $5e-7$ is in the bottom 1%. Using the average BTC size of trades and the average hourly BTC volume per coin, we can also estimate the average number of trades per hour per coin. These estimates indicate that most coins tend have a handful of trades per hour at the times we execute our interventions.

EXPERIMENT

Procedure

We conducted 310,222 randomly spaced trials in the Cryptsy exchange over the course of six months. Before our experiment, we purchased 0.002 BTC worth each of 217 different altcoins available on Cryptsy. We then created bots that monitored and periodically traded in each of these altcoins’ markets in parallel. Each bot waited a random amount of time, between one and two hours, at the beginning of the experiment before conducting a first trade. Then, assuming there had been at least one trade in the last hour on a coin, the bot for that coin recorded the current market state and randomly chose to either buy a random small amount of the coin, sell the coin, or do nothing (as a control condition). The trade sizes were chosen uniformly at random between $5e-7$ and $5e-6$ BTC. Each bot then observed its coin’s market state at 15 minutes, 30 minutes,

Time	Condition	Dependent Var.	n Control	n Treat	Control Mean	Mean Effect	t -stat	p -value
15 Min.	Buy	Buy Prob.	25483	25602	0.279	0.019	4.79	2.96e-05
15 Min.	Buy	% Buy Vol.	24321	24313	0.290	0.017	4.40	1.97e-04
15 Min.	Buy	Trade Prob.	52050	51314	0.490	0.009	3.00	4.81e-02
15 Min.	Sell	Sell Prob.	25483	25987	0.721	-0.006	-1.44	1.00e+00
15 Min.	Sell	% Sell Vol.	24321	24660	0.710	-0.005	-1.31	1.00e+00
15 Min.	Sell	Trade Prob.	52050	51727	0.490	0.013	4.12	6.71e-04
30 Min.	Buy	Buy Prob.	23647	23871	0.278	0.003	0.83	1.00e+00
30 Min.	Buy	% Buy Vol.	23583	23802	0.291	0.005	1.33	1.00e+00
30 Min.	Buy	Trade Prob.	52049	51312	0.454	0.011	3.51	7.98e-03
30 Min.	Sell	Sell Prob.	23647	23809	0.722	0.001	0.15	1.00e+00
30 Min.	Sell	% Sell Vol.	23583	23735	0.709	0.002	0.58	1.00e+00
30 Min.	Sell	Trade Prob.	52049	51724	0.454	0.006	1.94	9.53e-01
60 Min.	Buy	Buy Prob.	31065	31118	0.274	0.003	0.76	1.00e+00
60 Min.	Buy	% Buy Vol.	30984	31044	0.289	0.001	0.33	1.00e+00
60 Min.	Buy	Trade Prob.	52030	51288	0.597	0.010	3.18	2.70e-02
60 Min.	Sell	Sell Prob.	31065	31351	0.726	0.000	0.14	1.00e+00
60 Min.	Sell	% Sell Vol.	30984	31275	0.711	-0.003	-1.00	1.00e+00
60 Min.	Sell	Trade Prob.	52030	51713	0.597	0.009	3.02	4.50e-02

Table 1: t -tests for our three main dependent variables at each monitor time associated with our interventions. All statistics were computed on our confirmatory dataset. p -values are two-sided and Bonferroni corrected for 18 tests. We have differing numbers of observations between our probability and percentage statistics for like conditions at identical monitor events due to differences in sensitivity to trade execution latency in how we implemented these statistics.

Dep. Var.	Independent Var.	Coef.	t -stat	p -value
Buy Prob.	Buy Treat.	0.016	4.47	1.4e-04
Buy Prob.	Buy Treat.*Time 2	-0.016	-3.02	4.48e-02
Buy Prob.	Buy Treat.*Time 3	-0.015	-3.15	2.92e-02
Buy Prob.	Sell Treat.	0.004	1.01	1.00e+00
Buy Prob.	Sell Treat.*Time 2	-0.004	-0.72	1.00e+00
Buy Prob.	Sell Treat.*Time 3	-0.003	-0.61	1.00e+00
% Buy Vol.	Buy Treat.	0.014	4.14	6.39e-04
% Buy Vol.	Buy Treat.*Time 2	-0.012	-2.46	2.53e-01
% Buy Vol.	Buy Treat.*Time 3	-0.015	-3.30	1.75e-02
% Buy Vol.	Sell Treat.	0.003	0.98	1.00e+00
% Buy Vol.	Sell Treat.*Time 2	-0.005	-1.03	1.00e+00
% Buy Vol.	Sell Treat.*Time 3	0.001	0.21	1.00e+00
Trade Prob.	Buy Treat.	0.006	2.28	4.05e-01
Trade Prob.	Buy Treat.*Time 2	0.002	0.41	1.00e+00
Trade Prob.	Buy Treat.*Time 3	0.000	0.08	1.00e+00
Trade Prob.	Sell Treat.	0.011	4.11	7.1e-04
Trade Prob.	Sell Treat.*Time 2	-0.007	-1.79	1.00e+00
Trade Prob.	Sell Treat.*Time 3	-0.004	-0.98	1.00e+00

Table 2: Linear regressions for our three main dependent variables, including fixed effects for each coin, control variables for market state, and cluster-robust standard errors. All statistics were computed on our confirmatory set. p -values are two-sided and Bonferroni corrected for 18 tests. The significant negative effects at Time 2 and Time 3 indicate that the positive peer influence effect of our buy interventions diminishes over time.

and 60 minutes following the trial, after which the bot waited a random amount of time between 0 and 60 minutes before engaging in another trial and repeating. As a part of monitoring the market state, each bot recorded the details of the most recent trade in the market, as well as the total buy-side and sell-side volume since the bot's intervention. We conducted a power analysis based on a pilot study to determine the length of time to run our experiment.

Dependent Variables

We examine three dependent variables in our main analyses. In order to measure peer influence on trade direction, we use two statistics. The first statistic is the probability that the last trade observed on a coin is a buy as opposed to a sell, conditional on there being a trade in the observation period of the statistic. This statistic uses the transaction type of the last transaction that occurred in the time period associated with a given monitor event. The statistic aggregates binary indicator variables that check whether the last observed transaction types were buys or sells. For this statistic, we omit monitor events where we have not observed a trade after our own initial trade or after the last monitor event. The second statistic is the proportion of trading volume on the buy-side or sell-side after our interventions. To compute this statistic, we calculate the total volume of trades of a given coin that occurred in the time period associated with a given monitor event, and we calculate the fraction of that volume associated with buy or sell trades. Once again, this statistic is undefined when there has been no new trading activity at a given monitor event. Our statistical tests compare the values of these statistics across the treatment and control trials in each monitor period. Peer influence in buying or selling would be indicated by buy treatments leading to higher buy-side activity as measured by these statistics, or by sell treatments leading to higher sell-side activity. A third statistic, the probability of observing any trade at all in each monitor period, allows us to test for peer influence in overall trading activity in addition to trade direction. We compute each of these statistics in three mutually exclusive periods: in the 15 minutes after our interventions, the following 15 minutes, and the following 30 minutes.

RESULTS

Statistical Analysis

For a basic analysis, we use t -tests to compare the average values of our dependent variables in our buy and sell trials

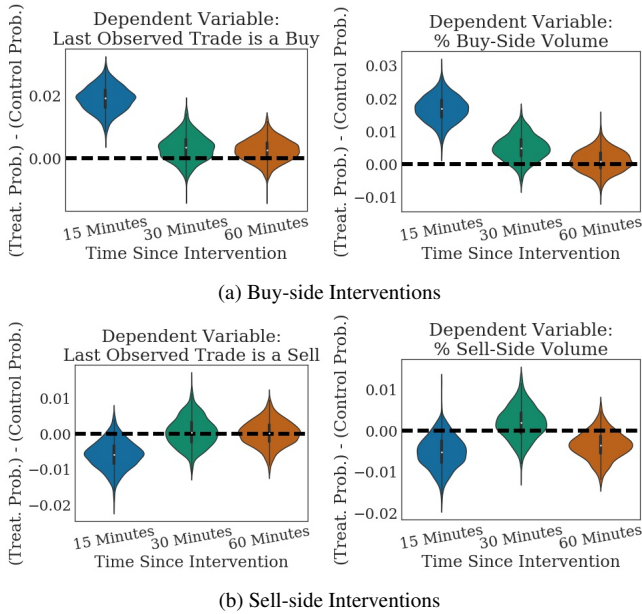


Figure 3: Bootstrap test statistic distributions for two dependent variables as a function of time after our interventions.

to the average values in our control trials. The results of these tests are given in Table 1 and visualized in Figures 3 and 4. We find significant effects of our buy interventions on buy probability and percent buy-side volume at the 15-minute monitor event ($p < 0.001$), and significant effects of buy and sell interventions on overall trading activity at all monitor events ($p < 0.05$) except the 30-minute monitor for sell interventions.

In addition to these t-tests, we perform a more robust analysis that combines the analysis of our buy and sell treatments across monitor times, while also controlling for potential dependence between observations on identical coins and potential dependence due to correlations in time. We perform linear fixed-effects regressions that include intercepts for each coin, and include regressor variables summarizing the market state to control for temporal dependence in the form of spillover effects. The variables we use to control for a coin’s market state are: six indicator variables for whether the best buy or sell order in the coin’s order book are above, below, or at the coin’s last transaction price—these variables indicate how a new trade would affect the price trend of that coin; an indicator variable for whether the last trade before our intervention was a buy or a sell; a continuous variable giving the percentage of buy-side versus sell-side volume in the hour before the intervention; and a continuous variable giving the logarithm of the coin market volume in the hour before our intervention divided by the average hourly volume on that coin. These regressions also include two indicator variables for whether an observation is associated with a buy trial or a sell trial (versus a control trial, which is absorbed into the intercept terms), and the regressions include interaction terms between these trial indicators and indicator variables for each observation’s associated monitor number. We use cluster-robust White standard

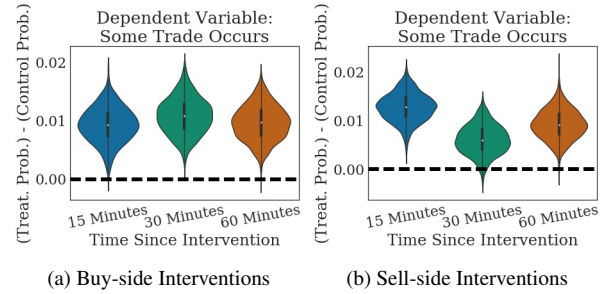


Figure 4: Bootstrap test statistic distributions for our final dependent variable as a function of time after our interventions.

errors [84], which adjust the model standard errors to account for heteroskedasticity and correlation in error terms within each altcoin’s observations. The results of these regression are given in Table 2. These results confirm the results of our t-tests. The effect of our buy treatments on buy trade probability and buy-side volume percentage remain significant at the 15-minute monitor event ($p < 0.001$). The effect of our buy interventions on the probability of observing a trade at all is marginally significant, being significant at the 0.05 level before we correct for multiple comparison but not afterwards. The effect of our sell interventions on the probability of observing a trade at all remains significant ($p < 0.001$). In the appendix we examine the randomization validity of our experiment, the extent of selection bias we might have suffered, and heterogeneity in the effects of our interventions across coins.

Peer Influence

Our results indicate strong evidence for peer influence on our buy interventions in these markets. The probability the last observed trade is a buy 15 minutes after a buy intervention is 30%—an increase of 2% above the control probability of 28%. The average percentage buy-side BTC volume is 2% higher than the 29% we observe in control trials during these 15 minute windows. Our interventions also led to overall increases in trading activity. The probability of observing a trade at all within 15 minutes after our interventions increased from 49% to approximately 50% after buy or sell interventions.

The aggregate effects of slightly higher percentages of buy-side activity and slightly higher overall levels of trading, in combination with a fat-tailed distribution of trading volume, accumulate into a large 7% average increase in total buy-side BTC volume after our buy interventions as compared to control trials (MWU-test $p = 0.0011$; we use an MWU-test since the distributions are fat-tailed, $MAE/MSE = 0.273$ as compared to 0.799 of the normal distribution). In total, we observe approximately 16,000 USD additional buy-side trading in our buy intervention trials as compared to the control trials. Even though we conducted slightly fewer buy interventions than there were control trials, the sum of all buy-side volume immediately following our buy interventions was 1513 BTC, as compared to 1430 BTC after the control trials (and 1399 after sell trials). This difference of 83 BTC is large compared to the total size of all of our interventions, which was approxi-




Figure 5: Plots of average hourly and daily volume.

mately 0.14 BTC. The total effects of our buy interventions were approximately 500 times larger than their cost.

Asymmetric Null Effect

We do not observe a symmetric peer influence effect from sell interventions. Sell interventions had no detectable effect on the proportion of future sell-side trading. There are multiple potential causes of this asymmetry. Since we executed roughly the same number of buy and sell interventions, and since the dependent variables in each case have the same variance, the asymmetry cannot be due to a difference in statistical power. Therefore the difference must be due to differences in how the market behaves following buy versus sell events. Notably, asymmetric peer influence effects of this sort have been observed in online social recommendation systems in which upvotes lead to peer influence but not downvotes [62]. However, other contextual factors could be causing the difference in our case. One possibility could be the fact that sell actions do not decrease prices in these markets as frequently as buy actions raise prices (sell actions only decrease prices 38% of the time, compared to the 78% of buy actions), and hence sell treatments may not lead to the same momentum effects. Another potential factor is that conducting a sell-side trade requires having holdings in the asset since short-selling was not implemented on Cryptsy. Therefore the sell effects could be weaker due to a smaller population of peers with the capability to sell.

Temporal Trends

The peer influence effects we observe after buy interventions do not lead to detectable permanent shifts in market dynamics. By 30 minutes after our interventions there is no longer a detectable peer influence effect on trade direction. This diminishing effect over time could be due to the fact that natural variation in subsequent trading after our interventions consists of a mix of buying and selling, which is variance that could dampen the effects of our buy trades. The overall excitation effects we observe are persistent over time. The probability of observing a trade between a half hour and an hour after either a buy or a sell intervention remains over 1% higher than the baseline probability in the control condition, which in this time period is 60%.

Market Composition

The interpretation of our results depends strongly on what sorts of participants compose the Cryptsy marketplace. If there are many active bots, or if our interventions disproportionately affect bots, then the implications of our experiment for the

potential effects of exchange design choices would differ. In order to check if our results are due mainly to human activity, we conduct a regression analysis including an independent variable that indicates whether each intervention was executed during the hours of the NY Stock Exchange (9am to 4pm Eastern, Monday-Friday) and an interaction term between this variable and the experimental condition. Here we compare buy versus sell interventions directly, rather than each to the control condition, to maximize our statistical power. We do observe cyclical trends overall in average volume during U.S. work hours, as shown in Figure 5, which indicates that the markets likely have substantial human trading. However, our analysis is inconclusive since we detect no significant interaction in our regression. In exploratory analysis, we obtained similarly inconclusive results using a variety of independent variables, including holidays, weekdays, and hour of the day. We therefore take our null result in this analysis to be weak evidence that bots are contributing meaningfully to our results.

DISCUSSION

Mechanisms

We now discuss potential mechanisms underlying the results of our experiment. We have limited ability to distinguish between potential causes due to the scope our data and experiment. In lieu of being able to provide evidence to favor one hypothesis over another, we attempt to enumerate possible explanations.

Potential Ultimate Causes

In the analysis of our experiment we attempted to identify to what extent our results might have been attributable to bot trading. We have a limited ability to do so since we cannot observe the identities of traders or which trades are executed via the GUI versus via the API, and we know of no major changes to the trading page interface during our experiment that could afford quasi-experimental analysis. In any case, our results are ultimately due to human behavior, whether:

- Direct human decision-making in GUI trading.
- Human planning, as manifest in what bots are programmed.

These two possibilities each implicate different mediating mechanisms and moderating contextual factors.

Potential Behavioral Mechanisms

There are a few plausible simple behavioral mechanisms that could underlie our results:

- Explicitly copying buy trades.
- Buying on an increasing price trend, i.e. momentum.
- Buying salient coins, such as coins with recent activity.

Since trade history is openly available on Cryptsy, a simple explanation is that some fraction of traders decide to buy if there has recently been a buy from another trader. Another possibility is that traders are paying attention to price increases, rather than explicitly copying buy actions. Buy actions raised the last transacted price 78% of the time during our experiment. Another possibility is that people buy salient coins. The frequent price increases from buy actions were reflected

	GUI			API		
	Bitfinex	Bitstamp	Cryptsy	Bitfinex	Bitstamp	Cryptsy
Ticker chart	Blue	Blue	Blue	White	White	White
Prominent ticker chart	Blue	Blue	Blue	White	White	White
Interactive ticker chart	Blue	Blue	Orange	White	White	White
Order book visualization	Blue	Blue	Orange	White	White	White
Customizable chart colors	Blue	Orange	Orange	White	White	White
Audio representations of activity	Orange	Blue	Orange	White	White	White
Quick buy/sell buttons	Orange	Blue	Orange	White	White	White
Page with coin statistics compared	Orange	Orange	Blue	White	White	White
Active markets/price changes highlighted	Orange	Orange	Blue	White	White	White
Moderated forums	Orange	Orange	Orange	White	White	White
Coin information pages	Orange	Orange	Orange	White	White	White
Your trades and orders	Blue	Blue	Blue	Blue	Blue	Blue
Recent trades displayed	Blue	Blue	Blue	Blue	Blue	Blue
Live updating	Blue	Blue	Orange	Blue	Blue	Blue
Basic market summary statistics displayed	Blue	Blue	Blue	Blue	Blue	Blue
Summaries of all markets displayed on page	Blue	Orange	Blue	Orange	Orange	Blue
Market indicators available	Blue	Blue	Orange	Orange	Orange	Orange
Communication channel	Blue	Orange	Blue	White	White	White
Trader identities available	Orange	Orange	Orange	White	White	White
“Recommended coins”	Orange	Orange	Orange	White	White	White
Published rate limits	White	White	White	Blue	Blue	Orange

Table 3: Comparison of the graphical user interfaces (GUIs) and application programming interfaces (APIs) of three cryptocurrency exchanges. Blue indicates the feature is present. Orange indicates the feature is absent. White indicates the feature is not applicable.

	Bitfinex	Bitstamp	Cryptsy
Coins	25	5	217
Continuous double auction	Blue	Blue	Blue
Open order book	Blue	Blue	Blue
Open trade history	Blue	Blue	Blue
Spot trading	Blue	Blue	Blue
Margin trading	Blue	Orange	Orange
Short-selling	Blue	Orange	Orange
Market-making	Orange	Orange	Orange

Table 4: Comparison of cryptocurrency exchange market mechanisms. The color code is the same as in Table 3.

in the charts that display summaries of the prices of each coin. Peer influence from buy actions may therefore be due to buying these salient coins. The first two mechanisms could easily be implemented either by people directly or through bots. The last could be implemented by bots, but would more plausibly be implemented by people since the API does not highlight price changes. More sophisticated mechanisms are also a possibility if traders conduct more complex calculations that result in individual buy trades having an overall marginal positive peer influence effect.

Potential Moderating Contextual Factors

There are also several plausible contextual factors that might moderate the effects we observe:

- Interface and market features
- Uncertainty of cryptocurrencies
- Aspects of the community disposition
- Size of Cryptsy

We discuss interface and market features in more detail below. Other contextual factors that might affect the generalizability

of our results are the highly speculative nature of cryptocurrencies, peculiarities of the composition of traders on Cryptsy, or the size of the Cryptsy marketplace. We expect that speculation, amateur trading, and a relatively small marketplace in which our interventions are more visible might have promoted peer influence in our study.

Cryptocurrency Exchange Comparison

The website CryptoCoinCharts currently lists 125 online exchanges where bitcoins and various altcoins can be traded. In this section we discuss common features of these exchanges as well as other features that could plausibly be implemented in order to explore a design space of cryptocurrency exchanges and to provide context for our experiment. This analysis illustrates that the structure of Cryptsy is relatively representative of the structure of currently popular cryptocurrency exchanges, and provides a point of reference for a discussion of plausible interface effects in our experiment.

We compare Cryptsy to the two highest-volume cryptocurrency exchanges at the time of writing: Bitfinex and Bitstamp. Tables 3 and 4 summarize key features of these exchanges. The three exchanges are similar in many ways, but also have some important differences. Common to all the exchanges is the use of a continuous double auction market mechanism, open order books and trading history, spot trading functionality, and a prominent ticker chart visualizing recent price changes. One notable difference in market mechanisms is that Bitfinex implements margin trading, which allows for leveraged trades and short-selling. Notable differences in the interfaces are that Cryptsy did not include charts for market indicators; Bitfinex and Bitstamp do not highlight outside market activity on an individual coin’s trading page; and Bitstamp does not have a communication channel for traders—Cryptsy had its own chat room, while Bitfinex links directly to a separate site,

TradingView, from its platform. None of the exchanges implemented market makers, and none of the exchanges implemented common features of other types of online platforms, such as moderated forums, information pages, deanonymized user activity, or recommendations.

Interface Effects

Given the design space we outlined in our comparison of cryptocurrency exchanges, we can hypothesize about which interface and market features might enable or promote peer influence. The design choices of the GUI clearly could impact human trading, and the affordances of the API clearly impact what bots can do. GUI design choices also could influence what types of bots people decide to implement. Every feature we observe in the Cryptsy interfaces plausibly has either a neutral or positive effect on peer influence, and many of these features are shared by other exchanges. For both human and bot traders, the features of the market mechanism and the ready availability of recent trading activity, common to all the exchanges we compared, is what enables peer influence in the first place. For human traders, the prominent display of trends in price history in the ticker chart, again a feature common to all exchanges we compare, plausibly encourages peer influence. More uniquely, Cryptsy's side-panel display of upward and downward price movements could have led to a saliency bias [42]. The chat functionality on Cryptsy also likely promotes peer influence, although of a different sort than the type we studied. An important area of future work is to more rigorously explore the causal impact of these design features through laboratory experiments, field experiments, or quasi-experimental observational analysis.

Generalizability

With the considerations of potential causes in mind, we expect our results to likely generalize at a minimum to other cryptocurrency exchanges, and possibly also at least to other online trading platforms (ZuluTrade, eToro, etc.) and other small markets (e.g., penny stocks or pink sheets stocks). The fact that Cryptsy highlighted price change direction in the side-panel of its GUI is relatively unique, but other exchanges have their own unique features that could also promote peer influence, such as the audible beeps that occur with all trades on Bitstamp. Generalization to larger and more professional financial markets is an important topic. Similar effects may be observable in higher-volume markets, but perhaps may require larger or more sustained interventions to be detectable. Obtaining a more rigorous understanding of the interface effects in these larger markets would be especially interesting.

Design Implications

In our discussion we have directed attention towards interface features in particular as potential moderating contextual factors. We highlight the point that many of Cryptsy's design choices could have plausibly promoted peer influence, and that Cryptsy could have potentially made alternative choices that might inhibit it. There is an interesting moral hazard implicated in these considerations. Cryptsy and other cryptocurrency exchanges make more money when more people use the platform, so they are incentivized to optimize their

sites to stimulate trading, including via peer influence. At the same time, there are multiple factors that go into creating a stable marketplace [53]. Maintaining an awareness of incentives and a focus on design goals could help to balance the positive and negative systemic effects of design choices.

RELATED WORK

There are several bodies of work within computational social science, economics and finance, and human-computer interaction related to the present study. We briefly review work on online field experiments in peer influence, peer influence in financial markets, online collective behavior, and the design of digital institutions.

Online Field Experiments in Peer Influence

A growing area across communities studying computational social science is digital and online experimentation [5, 69]. Our work was directly inspired by earlier online experiments studying peer influence in a variety of different types of online social systems [37, 70, 13, 62, 80]. These experiments have provided compelling evidence for the ubiquity of peer influence across an array of domains, and in some cases the importance of these effects on collective outcomes. Ours is the first study to apply the large-scale online field experimentation techniques innovated in these early works to online financial markets.

Peer Influence in Financial Markets

There have been a number of observational studies, laboratory experiments, and small-scale field experiments in finance and economics studying peer influence in financial markets. For instance, the widely recognized empirical phenomenon of momentum in price dynamics is related to our work. A number of researchers have identified evidence for momentum through observational analysis of real financial markets [75, 57, 59] and through the implementation of momentum-based trading strategies [46, 68, 47]. In another line of work, models of "herding" formalize the behavior of traders copying decisions to invest [7, 11, 4, 19]. These models have been tested directly in stylized laboratory experiments [2, 22, 21, 25]. Herding has been argued to occur in real markets through observational analyses of coarse-grained market data and individual institutional investor data [12, 78], but some analyses with market data have yielded negative results [78]. The empirical work in finance and economics most relevant to our own comes from the literature on attempting to manipulate asset prices in laboratory asset markets [36, 81, 82, 16] and certain types of real markets [17, 67]. The general form of these existing studies has been to execute large trades in the markets being studied and observe the effects on market prices over a short time period. We focus on the effects of small individual trades rather than attempts at market manipulation through abnormally large actions.

Cryptocurrency Market Dynamics

A growing body of work studies the dynamics of cryptocurrency markets specifically. Observational data analyses, with sometimes conflicting results, have been used to examine competition between cryptocurrencies for market volume [28]; to

compute the fundamental value of cryptocurrencies [40]; to investigate the efficiency of bitcoin markets [79, 63]; to confirm the presence of speculative behavior in bitcoin markets [58, 20]; to provide evidence for non-fundamentals-driven trading behavior in altcoin markets [26]; and to document how factors such as fundamental value [52, 14] or online search and discussion activity [51, 30] are related to cryptocurrency prices. A recent study closely related in spirit to our own provided observational evidence of market manipulation in USD-BTC markets [29].

Studying Online Collective Behavior

Our work also builds on a growing area within the human-computer interaction (HCI) and computer-supported cooperative work (CSCW) communities involving the study of human collective behavior in online platforms. Some of this work has focused on financial markets [31, 85]. Peer influence is related to the study of popularity dynamics in follower behavior (e.g., [43]) and online voting behavior (e.g., [56, 76, 55]). Others have studied how design can help ameliorate the effects of peer influence in online social recommendation systems (e.g., [50, 1]). Our work also employs the experimental technique of using bots in online field experiments that has recently been developed in these communities [72, 49].

Design of Digital Institutions

A final related area of work is the design of digital institutions. Many traditional institutions are shifting to having digital components, and as this shift occurs, the study of institutions becomes more relevant to researchers in human-computer interaction and computer-supported cooperative work. Researchers have studied a variety of different types of institutions, for example: payment systems [86]; economics [6] and marriage [27] in online worlds; organizational behavior [33]; activism [72]; knowledge markets [74]; labor markets [73, 3, 48, 32], large-scale collaboration [10, 61]; money (including cryptocurrencies) [18, 71]; personal finance [35, 34]; supply chains [66]; entrepreneurship [45]; hospitality [54]; and environmental sustainability [24]. Broadly, these works bring a design lens to bear on patterns of repeated digitally mediated large-scale social interaction, and particularly a lens for how the structure of technological artifacts affects our interaction patterns. Others in the community have considered ethical [44, 83, 15, 23, 38] and conceptual [60, 39] frameworks applicable to this type of design. Specific interest in HCI-oriented financial market design has existed since the early days of online markets. An early study investigated how market interfaces can impact market liquidity compared to physical “trading pits” [65]. Recent work has emphasized a view of markets as technologically-mediated human systems, and therefore a potentially fruitful target of HCI design and critique [53].

CONCLUSION

Institutional design is a major area of study in the social sciences. The HCI community has an opportunity to contribute to this conversation as many farflung institutions—from banks to marriage in Second Life—migrate to digital spaces. The methodologies for studying digital systems; the awareness of interface effects; and the keen eyes for bias, ethics, and inclusion in the HCI community could add unique perspectives to

the design of digital institutions. Markets are an example of an enormously important institution that is becoming increasingly digitized, and market irrationality may be a problem in markets that design-thinking could help address. Our specific application to cryptocurrencies is timely and urgent as new platforms are growing and potentially encouraging users to adopt risky trading strategies.

Our work provides an example of how peers in an online system can audit the system dynamics through experimentation with typical behavior. Bots that randomly execute actions of normal users could provide a way to understand peer influence and other phenomena in a variety of online systems. These bots allow us to understand the causal impact of individual actions that can be taken by users in these systems. In studying the dynamics of cryptocurrency markets with this technique, we show that even trades worth just fractions of a penny can influence the nature of other much larger trades in the cryptocurrency markets we study. We observe an approximately two percentage point increase in buying activity after our buy interventions, and a cumulative monetary effect of 500 times the size of our interventions. While an increase of two percentage points might seem small for an individual action, in a large marketplace this amount is non-trivial. For example, at the time of writing Apple stock on the NASDAQ exchange had an average daily trading volume of 30 million USD, 2% of which would amount to hundreds of thousands of dollars over the course of a day on that stock alone. Designers of online markets should be aware of how minor changes in their systems that affect individual and collective behavior could have major social and economic impact.

APPENDIX

Supplementary Analysis

Randomization Validity

We assess the validity of our randomization procedure to check that our treatment groups and control groups are not systematically different. We observe no detectable systematic differences between treatment and control groups on the dependent variables we measure before our interventions, but we do see significantly fewer treatment observations than expected by chance (binomial test, $n = 2 \cdot (52056)$, $p < 1e - 5$). We attribute this difference to a failure to treat, likely caused by a bug in Cryptsy that occasionally prevented us from executing trades. However, we find that our results are robust to simulating these failures to treat in the control condition. Plots of the numbers of control and treatment observations reveal that no single coin is much more unbalanced than others in terms of the number of control versus treatment observations on that coin.

Selection Bias

There are several potential sources of bias in our experimental design that could have influenced our estimated effect sizes. We only conducted interventions when the trade history and order book were accessible via the Cryptsy API, when we had sufficient funds to buy and sell at least $5e-6$ of the coins they were monitoring, and when at least one trade on the coin had occurred in the last hour. The fact that we condition on having access to the Cryptsy API, having enough coins to trade, and

having observed a trade at least an hour before our interventions means that we are always conditioning on a particular, albeit likely fairly common, market context in our experiments. Another source of bias in two of our dependent variables is missing trials due to having no trades observed in our monitor windows. Both the probability that the last observed trade is a buy and the percentage of buy-side volume are undefined when no trades are observed. We achieved only approximately 50% overall probability of observing any trades within 15 minutes after our interventions, and this probability varies widely across coins. The combined effects of these two sources of selection bias can be observed by looking at the number of observations we have per coin. Since higher volume coins are likely to meet both our condition for intervention and our condition for measurement, any bias on the marginal effects we examine will be towards the effect on higher volume coins, which are of more general interest anyway. Regardless, the fixed effects regressions we performed helps to control for the selection bias due to coin-level effects. Since our treatments affect the observability of our outcome variables, we might also be concerned that observability alone is driving our effects. However, we see that there is a significant difference in effect between buy and sell treatments, which do not differ strongly in terms of observability. In this analysis, the treatment type does not significantly interact with pre-treatment market state variables when predicting observability and controlling for multiple comparison in a linear regression.

Heterogeneous Effects

We observe substantial variation across coins in our dependent variables and in the effect sizes of our interventions. However, we did not identify any descriptive statistics of the coins that were reliably predictive of effect size (with effect size measured by the difference between the treatment mean and the control mean). A multivariate linear regression including median values of coin attributes (price, volume, spread, and best open sell/buy order size) yielded no significant relationships with effect size and low R-squared values of approximately 0.01. Since much of the variability occurs on coins with fewer observations, the major variations across coins are therefore likely largely due to noise. We needed a large sample size in order to be able to detect aggregate peer influence effects in these markets at all. We appear to have too small of a sample size per individual altcoin, and too few coins, to examine heterogeneous treatment effects.

Code, Data, Preregistration, and Institutional Review

Our code and data, including a list of cryptocurrencies we traded, is available online: <https://github.com/pkrafft/An-Experimental-Study-of-Cryptocurrency-Market-Dynamics>. Code for our experiment specifying the details of our experimental design was preregistered.¹ Our final statistical analysis differed from our preregistered one. Before conducting non-preregistered data analysis, we split our entire dataset of trials uniformly at random into two halves, one for exploratory statistical analysis and one for confirmatory statistical analysis. The results we show are from our confirmatory validation set. This confirmatory set was held-out from analysis as much

¹<https://osf.io/djezp/>

possible. Our experiment was approved by the human subjects review boards of MIT and ANU (MIT Protocol #1409006623, ANU Protocol: 2015/652). Because our experiments were conducted in the field and consisted of no more than regular trading activity in the markets we used, our study posed minimal risk to our participants and therefore was granted a waiver of informed consent.

ACKNOWLEDGMENTS

This material is based upon work supported by the NSF Graduate Research Fellowship under Grant No. 1122374, the MIT Media Lab Members Consortium, the Australian National University, Harper Reed, and Erik Garrison. Any opinion, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of our sponsors. Statistical support was provided by data science specialists Simo Goshev and Steven Worthington at the Institute for Quantitative Social Science, Harvard University. Special thanks to Iyad Rahwan for suggesting the moral implications of our results, and to the participants of the Second EC Workshop on Crowdsourcing and Behavioral Experiments for their feedback on an early version of this work.

REFERENCES

1. Andrés Abeliuk, Gerardo Berbeglia, Pascal Van Hentenryck, Tad Hogg, and Kristina Lerman. 2017. Taming the Unpredictability of Cultural Markets with Social Influence. In *Proceedings of the 26th International Conference on World Wide Web (WWW)*. International World Wide Web Conferences Steering Committee, 745–754.
2. Lisa R Anderson and Charles A Holt. 1997. Information Cascades in the Laboratory. *The American Economic Review* (1997), 847–862.
3. Judd Antin and Aaron D. Shaw. 2012. Social Desirability Bias and Self-reports of Motivation: A Study of Amazon Mechanical Turk in the US and India. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
4. Christopher Avery and Peter Zemsky. 1998. Multidimensional Uncertainty and Herd Behavior in Financial Markets. *American Economic Review* (1998), 724–748.
5. Eytan Bakshy, Dean Eckles, and Michael S Bernstein. 2014. Designing and Deploying Online Field Experiments. In *Proceedings of the 23rd International Conference on World Wide Web (WWW)*. ACM, 283–292.
6. Eytan Bakshy, Matthew P. Simmons, David A. Huffaker, Chun-Yuen Cheng, and Lada A. Adamic. 2010. The Social Dynamics of Economic Activity in a Virtual World. In *International Conference on Weblogs and Social Media (ICWSM)*.
7. Abhijit V Banerjee. 1992. A Simple Model of Herd Behavior. *The Quarterly Journal of Economics* (1992), 797–817.

8. Brad M Barber and Terrance Odean. 2008. All that Glitters: The Effect of Attention and News on the Buying Behavior of Individual and Institutional Investors. *Review of Financial Studies* 21, 2 (2008), 785–818.
9. Nicholas Barberis and Richard Thaler. 2003. A Survey of Behavioral Finance. *Handbook of the Economics of Finance* 1 (2003), 1053–1128.
10. Yochai Benkler. 2012. The Penguin and the Leviathan: Towards Cooperative Human Systems Design. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
11. Sushil Bikhchandani, David Hirshleifer, and Ivo Welch. 1992. A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades. *Journal of Political Economy* 100, 5 (1992), 992–1026.
12. Sushil Bikhchandani and Sunil Sharma. 2000. Herd Behavior in Financial Markets. *IMF Staff Papers* (2000), 279–310.
13. Robert M Bond, Christopher J Fariss, Jason J Jones, Adam DI Kramer, Cameron Marlow, Jaime E Settle, and James H Fowler. 2012. A 61-Million-person Experiment in Social Influence and Political Mobilization. *Nature* 489, 7415 (2012), 295–298.
14. Jamal Bouoiyour, Refk Selmi, Aviral Kumar Tiwari, Olaolu Richard Olayeni, and others. 2016. What Drives Bitcoin Price? *Economics Bulletin* 36, 2 (2016), 843–850.
15. Barry A. T. Brown, Alexandra Weilenmann, Donald McMillan, and Airi Lampinen. 2016. Five Provocations for Ethical HCI Research. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
16. Patrick Buckley and Fergal O’Brien. 2015. The Effect of Malicious Manipulations on Prediction Market Accuracy. *Information Systems Frontiers* (2015), 1–13.
17. Colin F Camerer. 1998. Can Asset Markets be Manipulated? A Field Experiment with Racetrack Betting. *Journal of Political Economy* 106, 3 (1998), 457–482.
18. John M. Carroll and Victoria Bellotti. 2015. Creating Value Together: The Emerging Design Space of Peer-to-Peer Currency and Exchange. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
19. Christophe Chamley. 2004. *Rational Herds: Economic Models of Social Learning*. Cambridge University Press.
20. Eng-Tuck Cheah and John Fry. 2015. Speculative Bubbles in Bitcoin Markets? An Empirical Investigation into the Fundamental Value of Bitcoin. *Economics Letters* 130 (2015), 32–36.
21. Marco Cipriani and Antonio Guarino. 2005. Herd Behavior in a Laboratory Financial Market. *American Economic Review* (2005), 1427–1443.
22. Marco Cipriani and Antonio Guarino. 2009. Herd Behavior in Financial Markets: An Experiment with Financial Market Professionals. *Journal of the European Economic Association* 7, 1 (2009), 206–233.
23. Nicola Dell and Neha Kumar. 2016. The Ins and Outs of HCI for Development. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
24. Paul Dourish. 2010. HCI and Environmental Sustainability: The Politics of Design and the Design of Politics. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*.
25. Mathias Drehmann, Jorg Oechssler, and Andreas Roeder. 2005. Herding and Contrarian Behavior in Financial Markets: An Internet Experiment. *American Economic Review* (2005).
26. Abeer ElBahrawy, Laura Alessandretti, Anne Kandler, Romualdo Pastor-Satorras, and Andrea Baronchelli. 2017. Evolutionary dynamics of the cryptocurrency market. *Royal Society Open Science* 4, 11 (2017), 170623.
27. Guo Freeman, Jeffrey Bardzell, Shaowen Bardzell, and Susan C. Herring. 2015. Simulating Marriage: Gender Roles and Emerging Intimacy in an Online Game. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
28. Neil Gandal and Hanna Halaburda. 2014. Competition in the Cryptocurrency Market. *Bank of Canada Working Paper 2014-33* (2014).
29. Neil Gandal, JT Hamrick, Tyler Moore, and Tali Oberman. 2017. Price Manipulation in the Bitcoin Ecosystem. *Journal of Monetary Economics* (2017).
30. David Garcia and Frank Schweitzer. 2015. Social Signals and Algorithmic Trading of Bitcoin. *Royal Society Open Science* 2, 9 (2015), 150288.
31. Eric Gilbert and Karrie Karahalios. 2010. Widespread Worry and the Stock Market. In *International Conference on Weblogs and Social Media (ICWSM)*. 59–65.
32. Mareike Glöss, Moira McGregor, and Barry A. T. Brown. 2016. Designing for Labour: Uber and the On-Demand Mobile Workforce. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
33. Jonathan Grudin. 1988. Why CSCW Applications Fail: Problems in the Design and Evaluation of Organization of Organizational Interfaces. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
34. Junius Gunaratne, Jeremy Burke, and Oded Nov. 2017. Empowering Investors with Social Annotation When Saving for Retirement.. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW)*.

35. Junius Gunaratne and Oded Nov. 2015. Informing and Improving Retirement Saving Performance Using Behavioral Economics Theory-driven User Interfaces. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
36. Robin Hanson, Ryan Oprea, and David Porter. 2006. Information Aggregation and Manipulation in an Experimental Market. *Journal of Economic Behavior & Organization* 60, 4 (2006), 449–459.
37. Ward A Hanson and Daniel S Putler. 1996. Hits and Misses: Herd Behavior and Online Product Popularity. *Marketing Letters* 7, 4 (1996), 297–305.
38. Ellie Harmon, Chris Bopp, and Amy Volda. 2017. The Design Fictions of Philanthropic IT: Stuck Between an Imperfect Present and an Impossible Future. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
39. John Harvey, David Golightly, and Andrew Smith. 2014. HCI as a Means to Prosociality in the Economy. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
40. Adam S Hayes. 2016. Cryptocurrency Value Formation: An Empirical Study Leading to a Cost of Production Model for Valuing Bitcoin. *Telematics and Informatics* (2016).
41. David Hirshleifer and Siew Hong Teoh. 2003. Herd Behaviour and Cascading in Capital Markets: A Review and Synthesis. *European Financial Management* 9, 1 (2003), 25–66.
42. Nathan Oken Hodas and Kristina Lerman. 2012. How Visibility and Divided Attention Constrain Social Contagion. In *Privacy, Security, Risk and Trust (PASSAT), 2012 International Conference on and 2012 International Confernece on Social Computing (SocialCom)*. IEEE, 249–257.
43. Clayton J Hutto, Sarita Yardi, and Eric Gilbert. 2013. A Longitudinal Study of Follow Predictors on Twitter. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*. ACM, 821–830.
44. Lilly Irani, Janet Vertesi, Paul Dourish, Kavita Philip, and Rebecca E. Grinter. 2010. Postcolonial Computing: A Lens on Design and Development. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
45. Karim Jabbar and Pernille Bjørn. 2017. Growing the Blockchain Information Infrastructure. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
46. Narasimhan Jegadeesh and Sheridan Titman. 1993. Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance* 48, 1 (1993), 65–91.
47. Narasimhan Jegadeesh and Sheridan Titman. 2001. Profitability of Momentum Strategies: An Evaluation of Alternative Explanations. *The Journal of Finance* 56, 2 (2001), 699–720.
48. Aniket Kittur, Jeffrey V Nickerson, Michael Bernstein, Elizabeth Gerber, Aaron Shaw, John Zimmerman, Matt Lease, and John Horton. 2013. The Future of Crowd Work. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
49. Peter M Krafft, Michael Macy, and Alex Pentland. 2017. Bots as Virtual Confederates: Design and Ethics. *The 20th ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)* (2017).
50. Sanjay Krishnan, Jay Patel, Michael J Franklin, and Ken Goldberg. 2014. A Methodology for Learning, Analyzing, and Mitigating Social Influence Bias in Recommender Systems. In *Proceedings of the 8th ACM Conference on Recommender Systems*. ACM, 137–144.
51. Ladislav Kristoufek. 2013. BitCoin meets Google Trends and Wikipedia: Quantifying the Relationship between Phenomena of the Internet Era. *Scientific Reports* 3 (2013), 3415.
52. Ladislav Kristoufek. 2015. What are the Main Drivers of the Bitcoin Price? Evidence from Wavelet Coherence Analysis. *PLoS ONE* 10, 4 (2015), e0123923.
53. Airi Lampinen and Barry Brown. 2017. Market Design for HCI: Successes and Failures of Peer-to-Peer Exchange Platforms. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, 4331–4343.
54. Airi Lampinen and Coye Cheshire. 2016. Hosting via Airbnb: Motivations and Financial Assurances in Monetized Network Hospitality. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
55. Alex Leavitt and Joshua A Clark. 2014. Upvoting Hurricane Sandy: Event-based News Production Processes on a Social News Site. In *Proceedings of the 32nd Annual ACM Conference on Human Factors in Computing Systems (CHI)*. ACM, 1495–1504.
56. Kristina Lerman and Tad Hogg. 2010. Using a Model of Social Dynamics to Predict Popularity of News. In *Proceedings of the 19th International Conference on World Wide Web*. ACM, 621–630.
57. Andrew W Lo and A Craig MacKinlay. 2002. *A Non-random Walk Down Wall Street*. Princeton University Press.
58. Alec MacDonell. 2014. Popping the Bitcoin Bubble: An Application of Log-periodic Power Law Modeling to Digital Currency. *University of Notre Dame Working Paper* (2014).
59. Burton G Malkiel. 2003. The Efficient Market Hypothesis and its Critics. *The Journal of Economic Perspectives* 17, 1 (2003), 59–82.

60. Thomas W Malone and Kevin Crowston. 1990. What is Coordination Theory and How Can it Help Design Cooperative Work Systems?. In *Proceedings of the ACM Conference on Computer-Supported Cooperative Work (CSCW)*.
61. Thomas W. Malone, Jeffrey V. Nickerson, Rob Laubacher, Laur Hesse Fisher, Patrick M. De Boer, Yue Han, and W. Ben Towne. 2017. Putting the Pieces Back Together Again: Contest Webs for Large-Scale Problem Solving. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
62. Lev Muchnik, Sinan Aral, and Sean J Taylor. 2013. Social Influence Bias: A Randomized Experiment. *Science* 341, 6146 (2013), 647–651.
63. Saralees Nadarajah and Jeffrey Chu. 2017. On the Inefficiency of Bitcoin. *Economics Letters* 150 (2017), 6–9.
64. Stefan Palan. 2013. A Review of Bubbles and Crashes in Experimental Asset Markets. *Journal of Economic Surveys* 27, 3 (2013), 570–588.
65. Satu S Parikh and Gerald L Lohse. 1995. Electronic Futures Markets Versus Floor Trading: Implications for Interface Design. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI)*.
66. Larissa Pschetz, Ella Tallyn, Rory Gianni, and Chris Speed. 2017. Bitbarista: Exploring Perceptions of Data Transactions in the Internet of Things. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
67. PW Rhode and KS Strumpf. 2006. Manipulating Political Stock Markets: A Field Experiment and a Century of Observational Data. (2006).
68. K Geert Rouwenhorst. 1998. International Momentum Strategies. *The Journal of Finance* 53, 1 (1998), 267–284.
69. Matthew J Salganik. 2017. *Bit by Bit: Social Research in the Digital Age*. Princeton University Press.
70. Matthew J Salganik, Peter Sheridan Dodds, and Duncan J Watts. 2006. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* 311, 5762 (2006), 854–856.
71. Corina Sas and Irni Eliana Khairuddin. 2017. Design for Trust: An Exploration of the Challenges and Opportunities of Bitcoin Users. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
72. Saiph Savage, Andres Monroy-Hernandez, and Tobias Höllerer. 2016. Botivist: Calling Volunteers to Action Using Online Bots. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing (CSCW)*. ACM, 813–822.
73. Aaron D. Shaw, John J. Horton, and Daniel L. Chen. 2011. Designing Incentives for Inexpert Human Raters. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
74. Dawei Shen, Marshall W. van Alstyne, Andrew Lippman, and Hind Benbya. 2012. Barter: Mechanism Design for a Market Incented Wisdom Exchange. In *The ACM Conference on Computer-Supported Cooperative Work and Social Computing (CSCW)*.
75. Robert J Shiller. 1981. Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends? *The American Economic Review* (1981).
76. Ruben Sipos, Arpita Ghosh, and Thorsten Joachims. 2014. Was this Review Helpful to You?: It Depends! Context and Voting Patterns in Online Content. In *Proceedings of the 23rd International Conference on World Wide Web (WWW)*. ACM, 337–348.
77. Vernon L Smith, Gerry L Suchanek, and Arlington W Williams. 1988. Bubbles, Crashes, and Endogenous Expectations in Experimental Spot Asset Markets. *Econometrica: Journal of the Econometric Society* (1988), 1119–1151.
78. Spyros Spyrou. 2013. Herding in Financial Markets: A Review of the Literature. *Review of Behavioral Finance* 5, 2 (2013), 175–194.
79. Andrew Urquhart. 2016. The Inefficiency of Bitcoin. *Economics Letters* 148 (2016), 80–82.
80. Arnout van de Rijt, Soong Moon Kang, Michael Restivo, and Akshay Patil. 2014. Field Experiments of Success-breeds-success Dynamics. *Proceedings of the National Academy of Sciences* 111, 19 (2014), 6934–6939.
81. Helena Veiga and Marc Vorsatz. 2009. Price Manipulation in an Experimental Asset Market. *European Economic Review* 53, 3 (2009), 327–342.
82. Helena Veiga and Marc Vorsatz. 2010. Information Aggregation in Experimental Asset Markets in the Presence of a Manipulator. *Experimental Economics* 13, 4 (2010), 379–398.
83. John Vines, Rachel Clarke, Peter C. Wright, John C. McCarthy, and Patrick Olivier. 2013. Configuring Participation: On How We Involve People in Design. In *Proceedings of the ACM Conference on Human Factors in Computing Systems (CHI)*.
84. Halbert White. 1980. A Heteroskedasticity-consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity. *Econometrica: Journal of the Econometric Society* (1980), 817–838.
85. Wenbin Zhang and Steven Skiena. 2010. Trading Strategies to Exploit Blog and News Sentiment. In *International Conference on Weblogs and Social Media (ICWSM)*.
86. Xinyi Zhang, Shiliang Tang, Yun Zhao, Gang Wang, Haitao Zheng, and Ben Y. Zhao. 2017. Cold Hard e-Cash: Friends and Vendors in the Venmo Digital Payments System. In *International Conference on Weblogs and Social Media (ICWSM)*.