Bits and Bets

Information, Price Volatility, and Demand for Bitcoin

Martis Buchholz, Jess Delaney, and Joseph Warren

Jeff Parker

Economics 312

Spring 2012

I. Introduction

Bitcoin is an online, digital currency, operating on a peer-to-peer network. The goal of the system is to establish a viable private currency without the need for a third party guarantor of transactions. Because bitcoins exist as digital data, this leads to what is known as the “double-spending problem,” how can the system disallow individuals from copying the currency in their possession and using it multiple times? Bitcoin solves this problem by publicly recording transactions on “block chains” that cannot be undone. The records on block chains are created by CPU power given to the network by users, who receive a small number of bitcoins in return (Nakamoto 2008). As transactions become more frequent over time, bitcoin users donating CPU power, or “miners” as they are colloquially know, receive a diminishing number of bitcoins in return for each block recorded. Thus the total supply of bitcoins is increasing over time at a diminishing rate (as can be seen in Figure 1).

Bitcoin was born in the midst of the financial crisis of 2008-2009, and its ethos is aligned with much of the political sentiment most prominent in that period.

> When Nakamoto’s paper came out in 2008, trust in the ability of governments and banks to manage the economy and the money supply was at its nadir. … Bitcoin required no faith in the politicians or financiers who had wrecked the economy—just in Nakamoto’s elegant algorithms. (Wallace 2011)

Support for Bitcoin, and investment in bitcoins was a political statement about the role of government in finance and the economy, as well as the ability of denizens of the internet to manage their own affairs. Particularly in the early months of Bitcoin’s existence, its functioning as a currency was sustained by individuals who were willing to pay a greater price in exchange for the knowledge that they were using a new technology, more in line with their ideals.

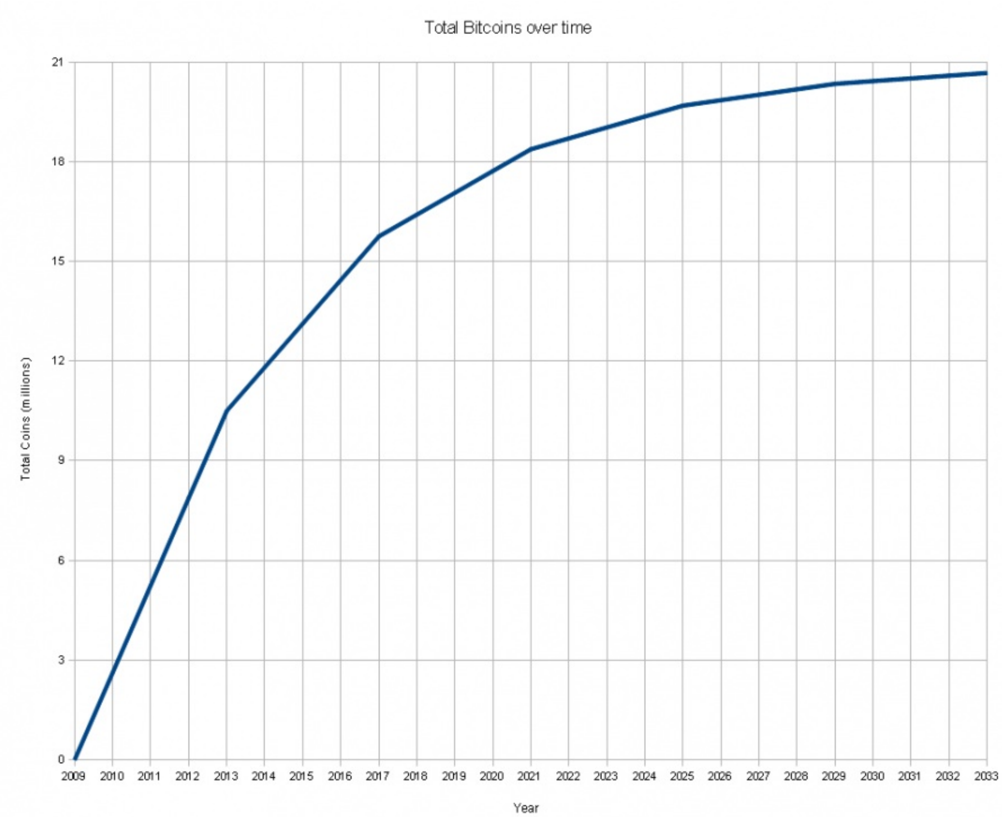

Figure 1. Expected total quantity of bitcoins over time (2009-2033), measured in millions.

These early adopters of bitcoin represented a variety of groups and motives, similar to users of many new technology or internet-related innovations,

> …including technology early adopters, privacy and cryptography enthusiasts, government-mistrusting "gold bugs," criminals, and speculators. A large number of online merchants accept bitcoins, catering to individuals with these interests, including web hosts, online casinos, illicit drug marketplaces, auction sites, technology consulting firms, and adult media and sex toy merchants. (Grinberg 2012, pp. 165)

Non-profit organizations such as Wikileaks, Freenet, Singularity Institute, Internet Archive, Free Software Foundation also accept donations in Bitcoins (wikipedia.org). One researcher took a poll of bitcoin enthusiasts (with 82 respondents) on an online forum, giving them a number of possible categories to explain their use of the product. The results (Figure 2), while neither scientific nor, probably, representative, are interesting.

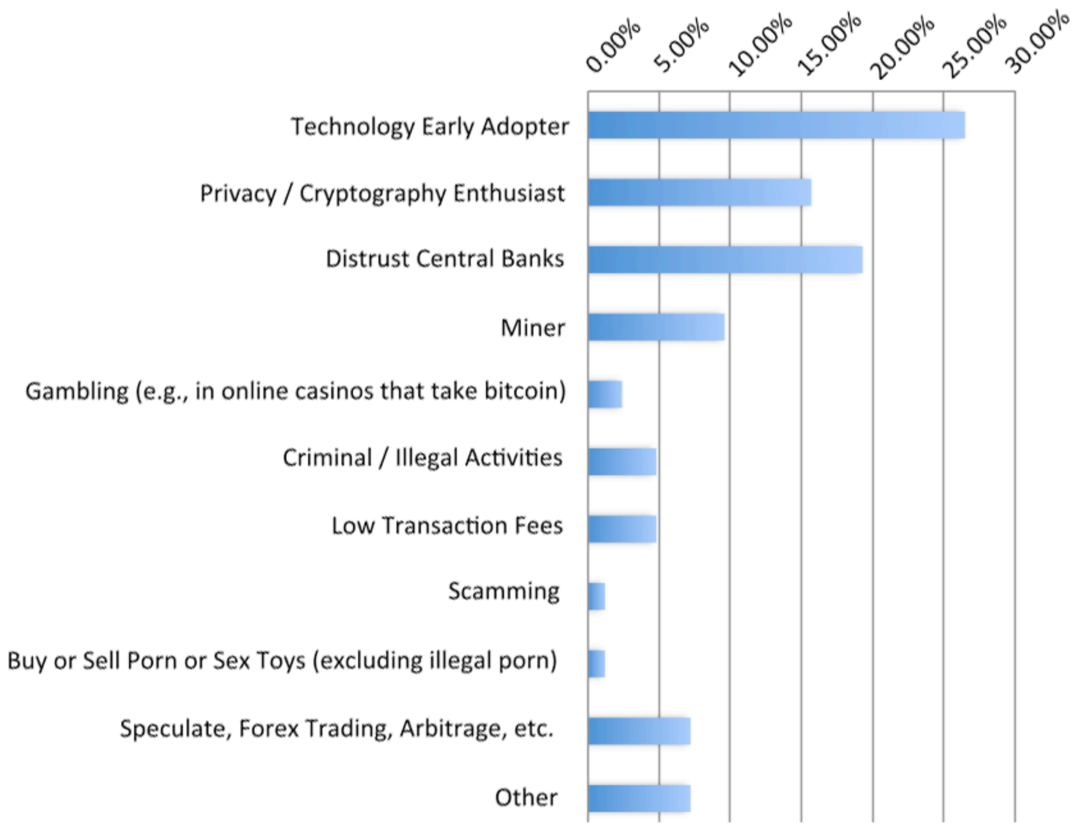

Figure 2. Reasons for Bitcoin adoption in poll https://bitcointalk.org/index.php?topic=4465.0

The most significant feature of Bitcoin's history is a sharp increase in price and users in the summer of 2011. Price increased exponentially, growing by several hundred thousand percent in several weeks, after which it fell by thirty percent in one day (Jeffries 2011). This growth and fall can be observed in Figure 3. The decline in interest in Bitcoin is emphasized by information from Sourceforge.com. Sourgeforge is the site where the Bitcoin client software used to store bitcoins on a user's desktop computer is obtained. Downloading this software might generally indicate an individual's intention to become a Bitcoin user. This data is monthly, and thus of limited value for analysis, but telling in terms of the spike in enthusiasm in June 2011, and the subsequent decline (Figure 4).

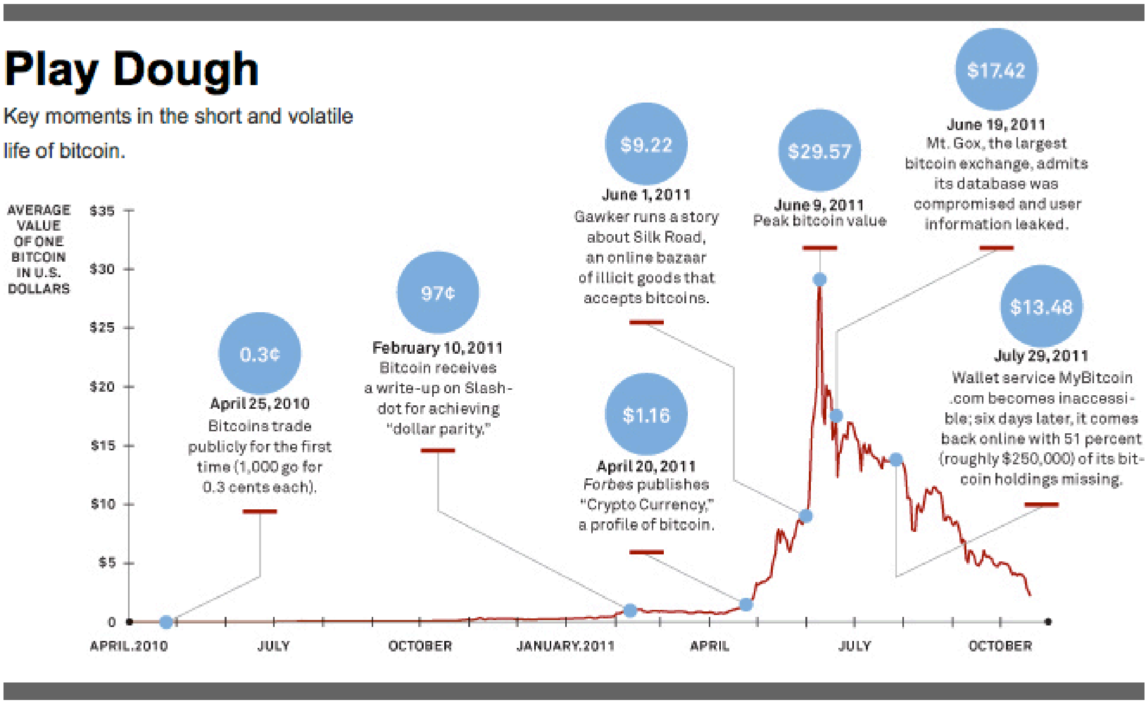

Figure 3. Price of a bitcoin over time in dollars alongside important events (from Wallace 2011).

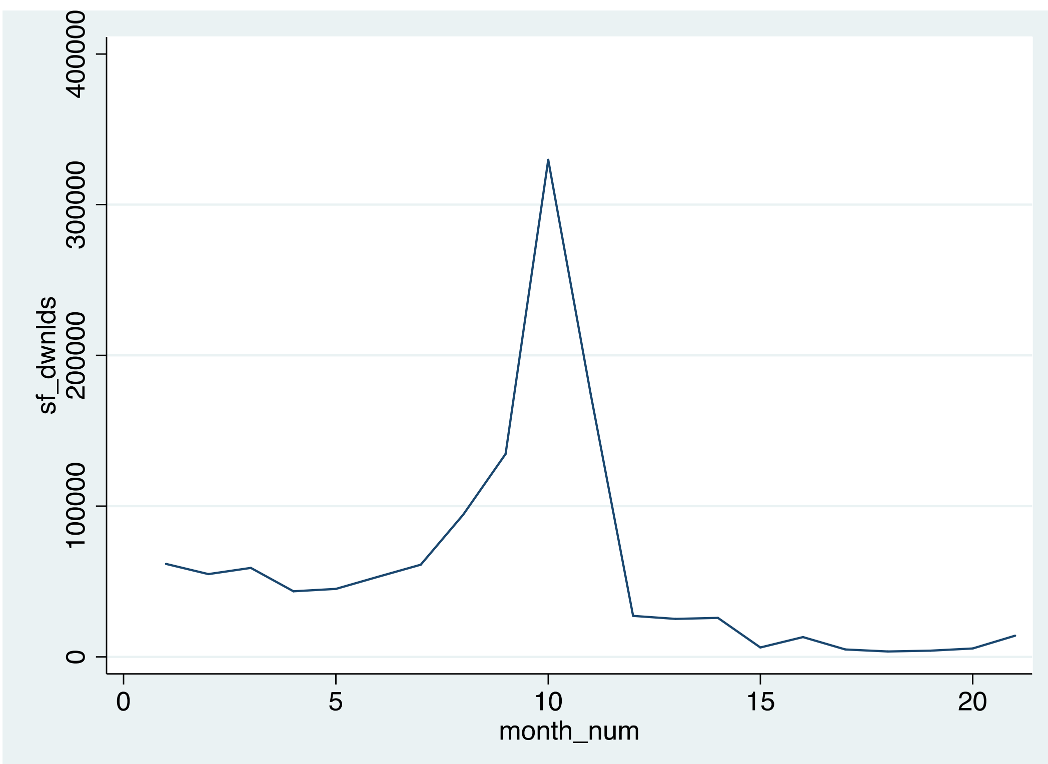

Figure 4. Downloads of the bitcoin software system from Sourceforge.com (6/2010 - 3/2012).

We have found no published economic literature on Bitcoin. A few law review articles explore the legal aspects of Bitcoin and other digital currencies, and in the process touch on technical and economic features of the systems. However, they do not systematically investigate any particular element of the economics of Bitcoin. Thus, we are left to apply more general models to this specific case.

II. Theory

Money exists to solve the problem of the “double coincidence of wants,” and it does so by fulfilling three functions: medium of exchange, unit of account, and store of value. The first of these is its fundamental and unique aspect, as other goods can fulfill the functions of unit of account and store of value, but the purpose of money is for exchange. McCallum (1989) presents an intuitive model explaining the role of money, in which transactions are costly but necessary for consumption. Thus, consumers seek to minimize their shopping time by holding positive amounts of money. Consumption leads to money demand because money lowers transaction costs. The basic money demand function states that the quantity of money demanded, divided by the price level, depends on consumption, divided by the interest rate.

$$M/P = \text{constant}*C/R$$

We are interested in individuals’ choices between competing currencies. Consumers can substitute between Bitcoin and other currencies in order to fulfill transactions for consumption: the users are variable. Hence, we are not interested in factors that affect both the dollar and bitcoins equally, but we are interested in features of bitcoin that influence an individual’s choice to hold the currency as opposed to dollars.

$$M(B)/P(B) = \text{constant}*C$$

$$C = (\#\text{users})*(\#\text{individual's transactions})*(\text{magnitude of ind. trans.})$$

C = f(qualities of bitcoins vs. dollars)

Bitcoin is both a product with the purpose to service transactions, and a currency that competes with the dollar. As a currency, it can be categorized as commodity-based, fiat, or somewhere in between. Commodity money is based on the value of a real good (such as gold). Because Bitcoin is composed of data that is of much lower value than the bitcoins themselves, and is not tied to any commodity or multiple commodities, Bitcoin is clearly not a commodity-based currency.

Selgin (2012) considers Bitcoin a “quasi-commodity” currency, which he defines as an asset in finite supply that does not have non-monetary value. However, Selgin does not provide strong reasons for distinguishing between quasi-commodity currencies and rule-based fiat currencies. Quasi-commodity money is simply at the extreme end of the continuum of possible restrictions on discretionary policy of the currency issuer. But even in the case of Bitcoin, the developers of the software could, in theory, offer an updated version altering the supply growth rule. In fact, this has already been suggested (Barber et. al. 2012). Thus, because Bitcoin is neither a commodity nor quasi-commodity-based currency, it is best classified as a fiat currency.

Private fiat currencies are predicted to suffer from at least two fundamental problems. The first of these objections regards network externality effects of holding currency. In the potential case of competition, one consumer’s decision to hold a particular brand of currency increased the returns to other consumers’ holding the same currency. This creates economies of scale in currency production.

> [T]he proliferation of notes, each convertible into different commodities-assets and issued by banks with differing portfolios, assessed riskiness, etc. would severely impair the information and transactional advantages that gives *(sic)* money its main functional role. Natural incentives would arise to standardise on a single commodity set as a

> base and/or to make the liabilities of smaller banks convertible at par into those of some dominant bank. (Goodhart 1989, pp. 48)

The second flaw considered to bar the sustainable implementation of a system of competitive fiat currencies is the time-inconsistency problem. Private issuers of fiat currencies do not have suitable incentives to avoid hyperinflation, in the absence of legal restraints. Fiat currencies are founded on faith, and thus consumers must trust private issuers to maintain a stable money supply. This is the distinguishing feature of fiat currency (White 1999).

But as the currency producer can increase revenue through hyperinflation, potential customers will not hold the private currency. Thus the system disintegrates due to the, “failure to show that the issuer will not break its promise of stable purchasing power,” (White 1999). In order for a private currency producer to establish itself, it must convince consumers to trust its product.

> That a profit-maximizing private issuer of inconvertible money would hyperinflate means that the time-inconsistency problem bedevils private fiat-type money production.... The presence of “brand name capital” does not solve the problem. (White 1999, pp. 238)

How does Bitcoin address these issues? In the first case, Bitcoin has undergone a process of diffusion similar to other innovations. The features of Bitcoin are most advantageous to a subset among all possible users. These individuals are the early adopters, and their own (potentially idiosyncratic) reasons for using Bitcoin have been discussed above. What is significant for understanding diffusion, as Nelson et. al. (2002) explain, is that adopters face sunk costs and flow benefits. Bitcoin has fairly low initial sunk costs for programmers and advanced computer users, but higher costs for other consumers.

The main categories of factors impacting the diffusion of innovation (Hall 2005):

- benefit received (constant + increasing with number of users)
- costs of adoption (increasing for less tech-savvy later adopters)

- industry or social environment (network-based, favoring early adopters)
- uncertainty and information problems (variable over time)

As Hall (2005) observes, the cost of adoption, “includes not only the price of acquisition, but more importantly the cost of the complementary investment and learning required to make use of the technology,” (pp. 473). This is likely to be of relevance to the diffusion of Bitcoin past early adopters. “Nontechnical newcomers to the currency, expecting it to be easy to use, were disappointed to find that an extraordinary amount of effort was required to obtain, hold, and spend bitcoins,” (Wallace 2011)

$$C = f(\text{benefits} - \text{costs} +/- \text{env.} - \text{uncertainty})$$

$$\text{Users} = f(\text{news, information})$$

What appears to have occurred in mid-2011 is the increasing costs of adoption for later, less tech-savvy customers overwhelmed the increasing benefits due to the expanding network of Bitcoin users. Because demand stopped shifting out, the price stopped rising. This lead many individuals, who had been hoarding bitcoins, to sell them for profit, causing the price to crash. Demand for bitcoins, currently, appears to have stabilized at a lower level.

The second problem, trust, is more serious for private currency issuers, and Bitcoins solution correspondingly more central to the system as a whole. The developers of Bitcoin encourage trust through a fixed money supply growth rule, supported by several mechanisms. The software is open-source and easily inspectable by any user. The developers do not gain revenue through supply increases, and therefore do not have any incentive to hyperinflate. Instead, profits from the increase in supply are distributed to users, with the fixed supply growth rate ensuring that selfish users to not drastically depreciate the value of the currency. As the growth rate of supply is currently fixed, bitcoin users know they must be alerted to any changes in policy because the only means for such a change to occur is through a new version of the software.

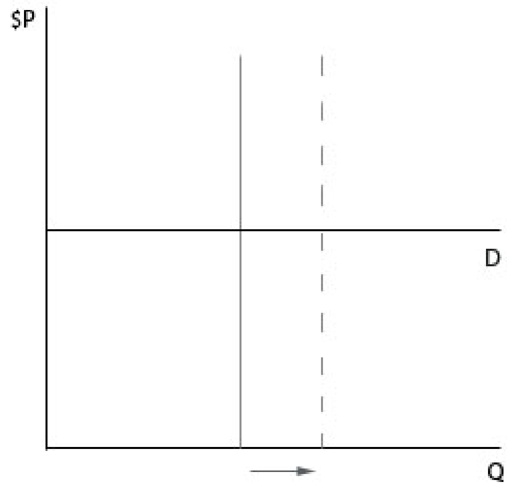

Figure 5. Supply is inelastic with respect to dollar price, and increases over time at about a constant rate during the sample period under investigation.

Supply is exogenous; it has not relationship to demand or price. Because supply does not change in response to price, we know that observed price fluctuations are due to shifts in demand. Because the quantity of bitcoins is increasing over time, the intersection between demand and supply is still moving down the demand curve.

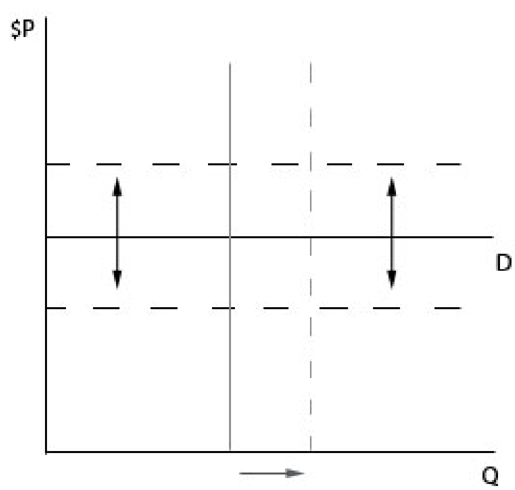

Figure 6. All observed price fluctuations occur due to shifts in demand.

The demand curve should be horizontal because any change in quantity is fully expected, implying that rational merchants will raise their prices in expectation of supply increases so as to avoid the effects of inflation. As supply over time should have no effect on the price of bitcoins in dollars, all observed price fluctuations should be due to demand shifts.

III. Questions

Bitcoin has unique solutions to the two problems faced by competitive fiat currencies. In the following sections, we will explore the effects of these solutions on Bitcoin's effectiveness as an innovative online currency. Numerous aspects of the Bitcoin system revolve around its solution to the problem of trust: a fixed money supply growth rule. Much of our attention will be focused on the effects of this feature, particularly the sensitivity of the dollar price of bitcoins to demand fluctuations.

1) How does information and online attention to Bitcoin diffuse and interact with changes in demand?
2) How does the transaction behavior of users respond to changes in the dollar price of bitcoins?
3) To what extent does price volatility affect demand?

IV. Data

Because Bitcoin exists exclusively online, every aspect of the system is, in theory, recordable. However, data does not exist (or is not readily available), for all the variables in which an economist might be interested. For example, we found no direct measure of the number of users. We obtained information of variables where data exists from a variety of sources.

Our time variable, date, covers the period from July 2010 through March 2012 for most of our variables. The online sources we utilized were accessed on 1 April 2012.

From http://www.blockchain.info/charts we accessed data on supply, number of transactions, total transaction value, and a price estimate.

total_bit	Supply of bitcoins in existence (exogenous).
transactions	Total number of bitcoin transactions per day.
transact_val	Total value of bitcoin transactions (measured in bitcoins) per day. (C)
price_est	Estimate of price($) of bitcoins from MtGox and Tradehill per day. This is more accurate than just MtGox data, as MtGox lost trust and market share after it was hacked in mid-2011, largely to Tradehill. On the other hand, we do not know precisely how this estimate was calculated.

From these variables, we were also able to develop an estimate of the average price in bitcoins for each day in our sample period.

ave_transact	A measure of the average price: (total transaction value)/(total number of transactions) = average value of transactions in bitcoins per day. The use of this variable assumes the bundle of goods is not changing, but this assumption is also made in the standard CPI.

Mt. Gox dominates the bitcoin exchange market, currently (as of May 2012, http://bitcoincharts.com/charts/volumepie/) taking up 72% of trade volume whereas the second largest exchange services only 6%. The data on Mt. Gox comes from http://bitcoincharts.com/charts/mtgoxUSD.

mtgox_price	price of bitcoins in dollars at MtGox per day.
mtgox_vol	trade volume at MtGox exchange per day (measured in dollars).

Data on historical google searches comes from http://www.google.com/insights/search/.

Google	google searches by week (normalized to 100).

Data on historical news articles and blogs comes from LexisNexis http://www.lexisnexis.com/hottopics/lnacademic/.

LexisNexus Mentions in news articles and blogs by week.

We obtained daily data for Twitter and the RSS-feed through the the TopicWatch application currently in beta-testing mode by LuckySort, a Portland startup.

Twitter Daily mentions of “bitcoin” on Twitter (Nov. 2011 – March 2012)
RSS Daily mentions of “bitcoin on the RSS-News feed in same period

V. Analysis

A. Diffusion of Bitcoin Information

Since our sample contains a period when bitcoin was relatively unknown, we were interested in estimating the effects of publicity on the market for bitcoins. We decided that the number of weekly Google hits (normalized to 100) would be a good proxy variable because it should be correlated with people hearing about bitcoin. Looking at the number of weekly Google hits, we see that for about the first 25 weeks, bitcoin was receiving almost no hits (shown in Figure 2). During this time, bitcoin was relatively unknown to the general public. Around the 30th week, Google hits begin to skyrocket, which corresponds to a sharp upward trend in the number of transactions, creating a bubble in the market. We wanted to estimate the relationship between Google hits and the number of transactions, since changes in the number of transactions should be a good indicator of people entering the market after learning about bitcoin. Since we only had weekly observations for Google hits, we used Stata’s collapse command to turn our daily data into Weekly averages. This generated 82 observations of weekly averages over the sample period. The summary statistics are shown in the following table.

. sum transactions Google

Variable	Obs	Mean	Std. Dev.	Min	Max
transactions	**82**	**4571.66**	**3067.11**	**301.4**	**12166.71**
Google	**82**	**14.91463**	**19.12674**	**0**	**100**

To avoid running spurious regressions, we proceeded to determine the appropriate time series model.

We decided not to transform the data by taking logs because this would drop all of our observations of Google that are zero. Since we're particularly interested in the rise of bitcoin's popularity, these observations are crucial. We begin by looking at plots of average weekly transactions and weekly Google hits to check for any evidence of stationarity.

Figure 1:

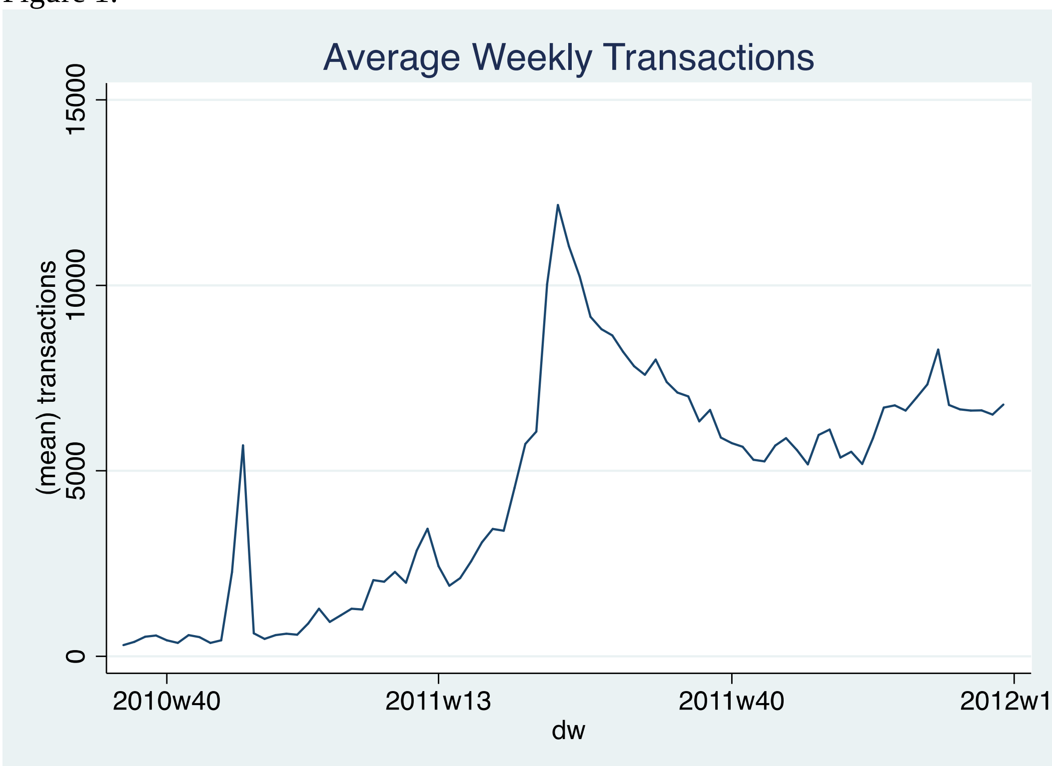

Figure 2:

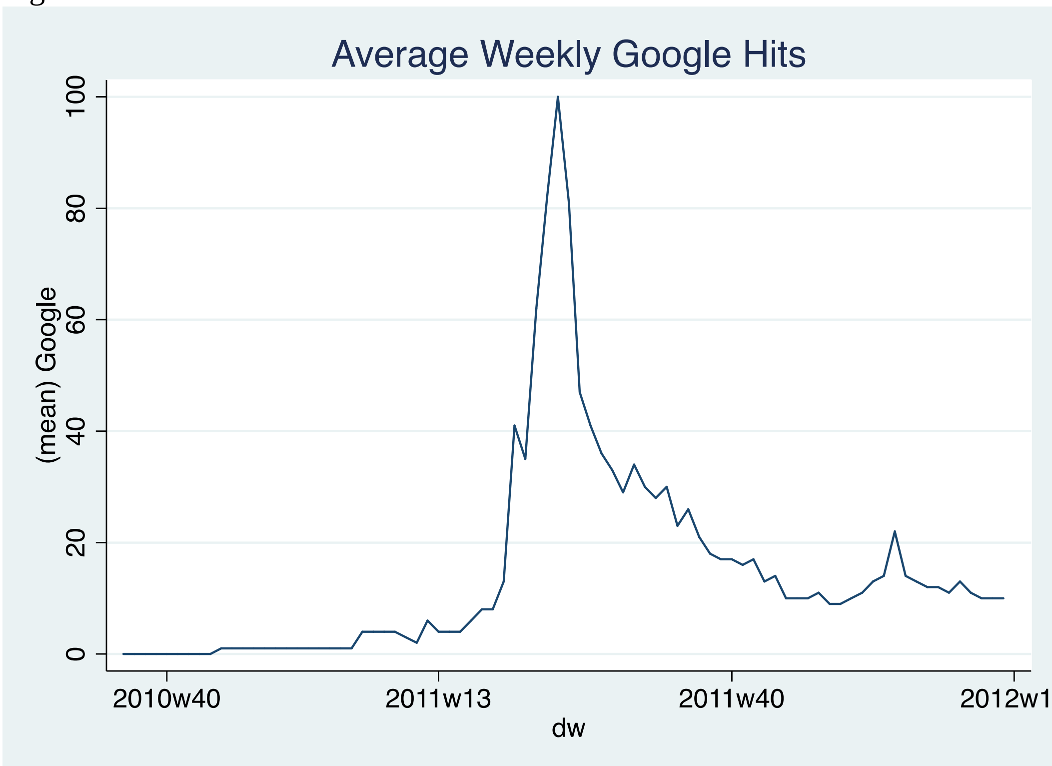

It appears that average weekly transactons may be trend stationary. We see that the number of Google hits has a large spike but seems to return to around zero. Now we will perform Augmented Dickey Fuller Tests to formally test for stationarity.

We begin by looking at the transactions variable. In order to find the appropriate number of lags to use in the ADF test we created a do-file to run several ADF tests with varying lags and perform Breusch-Godfrey tests to check for serial correlation. An example of that code is shown in Table 1. Since transactions appears to have an upward trend, we generate a variable t to capture that trend in our ADF tests.
. gen t=_n

Table 1: Example Do-File

```
forvalues p = 1/3 {
 qui reg L(0/`p').D.transactions L.transactions t
 di "Lags =" `p'
 estat bgodfrey, lags(1/10)
}
```

Table 2 shows the results Breusch-Godfrey test run after one ADF test for transactions. We can see that adding one lag eliminates serial correlation of the first 10 orders. So, we determined that using one lag in our ADF test would be appropriate. Table 3 shows the results of this ADF test.

Table 2:
Lags =1

Breusch–Godfrey LM test for autocorrelation

lags(*p*)	chi2	df	Prob > chi2
1	**0.502**	**1**	**0.4788**
2	**0.671**	**2**	**0.7150**
3	**0.904**	**3**	**0.8243**
4	**0.976**	**4**	**0.9134**
5	**1.354**	**5**	**0.9293**
6	**1.356**	**6**	**0.9685**
7	**1.490**	**7**	**0.9827**
8	**1.490**	**8**	**0.9929**
9	**1.512**	**9**	**0.9971**
10	**1.552**	**10**	**0.9988**

H0: no serial correlation

Table 3:
. dfuller transactions, trend lags(1)

Augmented Dickey–Fuller test for unit root   Number of obs   =   **80**

		Interpolated Dickey–Fuller		
	Test Statistic	1% Critical Value	5% Critical Value	10% Critical Value
Z(t)	**−2.165**	**−4.084**	**−3.470**	**−3.162**

MacKinnon approximate p–value for Z(t) = **0.5095**

We see that we cannot reject the null hypothesis that transactions has a unit root. Therefore, we conclude that the series is nonstationary. Next we repeat these steps for Google. Table 4 displays our results

Table 4:

Lags =1

Breusch-Godfrey LM test for autocorrelation

lags(*p*)	chi2	df	Prob > chi2
1	**1.209**	**1**	**0.2714**
2	**1.604**	**2**	**0.4484**
3	**2.241**	**3**	**0.5239**
4	**2.293**	**4**	**0.6821**
5	**2.425**	**5**	**0.7877**
6	**6.594**	**6**	**0.3601**
7	**6.769**	**7**	**0.4533**
8	**7.222**	**8**	**0.5129**
9	**7.413**	**9**	**0.5942**
10	**7.570**	**10**	**0.6707**

H0: no serial correlation

. dfuller Google, lags(1)

Augmented Dickey-Fuller test for unit root Number of obs = **80**

	Test Statistic	Interpolated Dickey-Fuller 1% Critical Value	5% Critical Value	10% Critical Value
Z(t)	**-2.380**	**-3.538**	**-2.906**	**-2.588**

MacKinnon approximate p-value for Z(t) = **0.1475**

Following the same steps that we did for transactions, we determined that 1 lag was appropriate in the ADF test and that Google is nonstationary. Next, we check to see if Google and transactions are I(1) variables. First, we take their first differences. The new variables are labeled with a “d” as their first letter. Figures 3 and 4 display the first difference series for Google and transactions, respectively.

Figure 3: dgoogle

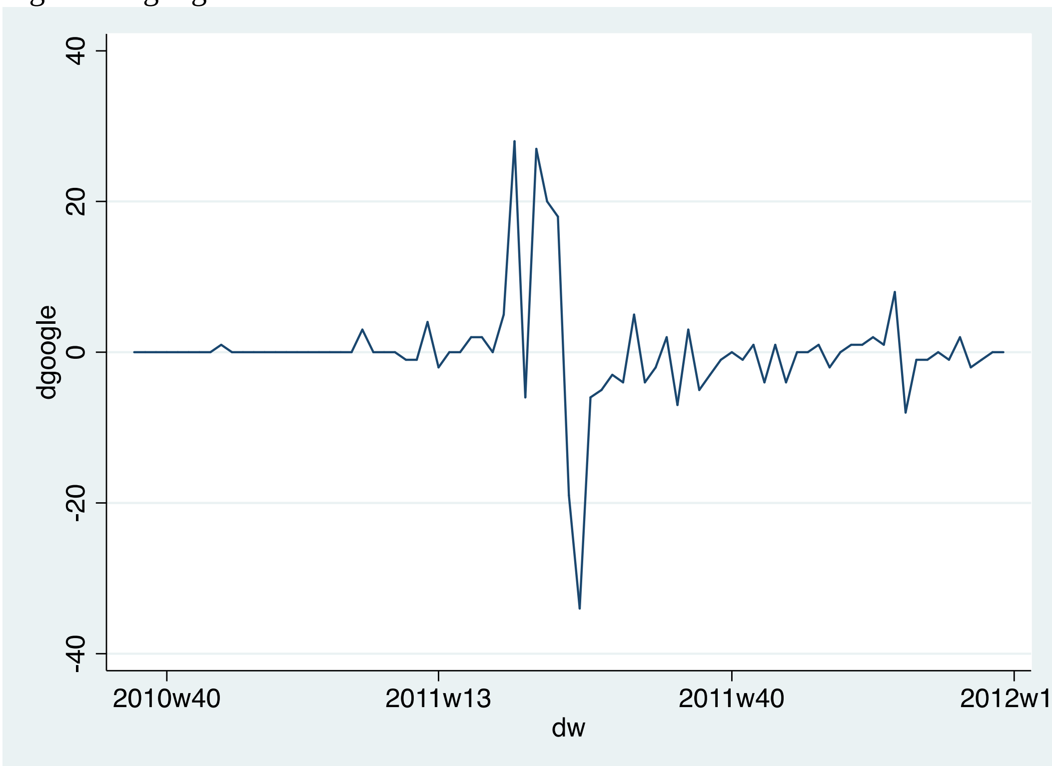

Figure 4: dtransactions

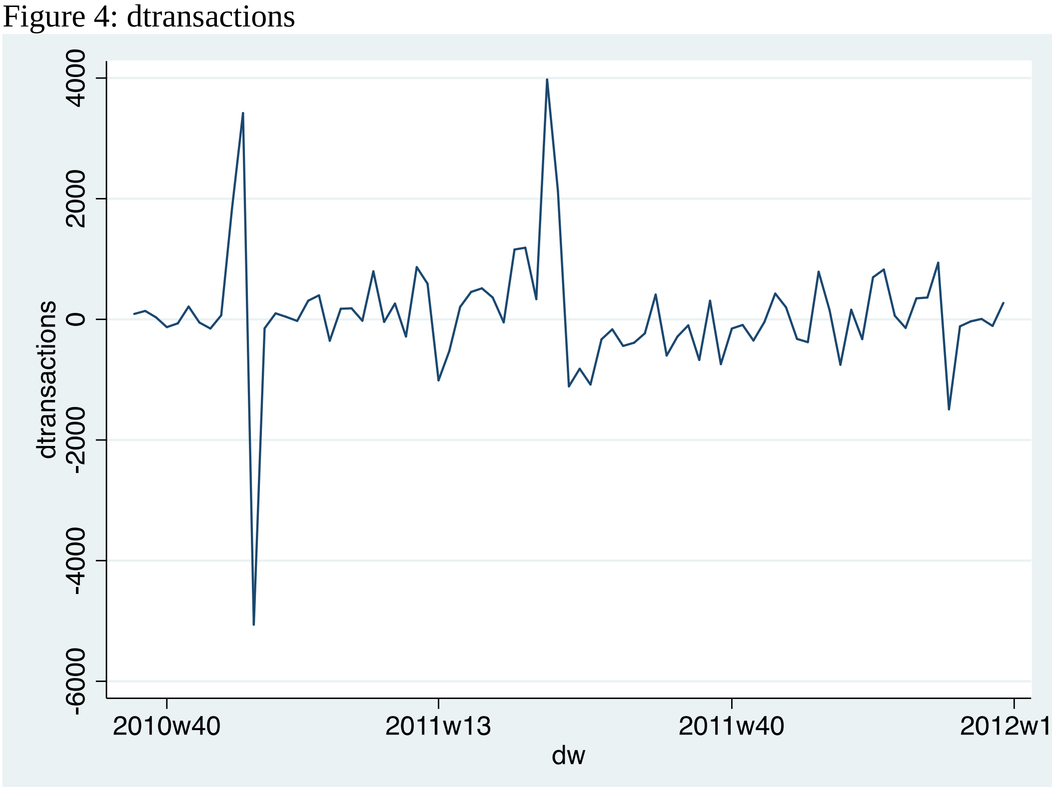

We can see in Figures 3 and 4 that these series appear to be stationary about zero. Now we formally test for stationarity using ADF tests. We repeat the same process as we did before. Table 5 and 6 show the results for dgoogle and dtransactions, respectively.

Table 5:

Lags =1

Breusch–Godfrey LM test for autocorrelation

lags(*p*)	chi2	df	Prob > chi2
1	**1.814**	**1**	**0.1781**
2	**1.820**	**2**	**0.4025**
3	**2.057**	**3**	**0.5607**
4	**3.220**	**4**	**0.5217**
5	**6.307**	**5**	**0.2774**
6	**7.690**	**6**	**0.2617**
7	**7.778**	**7**	**0.3526**
8	**7.892**	**8**	**0.4441**
9	**8.518**	**9**	**0.4829**
10	**8.757**	**10**	**0.5553**

H0: no serial correlation

. dfuller dgoogle,lags(1)

Augmented Dickey–Fuller test for unit root Number of obs = **79**

	Test Statistic	Interpolated Dickey–Fuller 1% Critical Value	5% Critical Value	10% Critical Value
Z(t)	**−5.030**	**−3.539**	**−2.907**	**−2.588**

MacKinnon approximate p–value for Z(t) = **0.0000**

Table 6:

Lags =1

Breusch–Godfrey LM test for autocorrelation

lags(*p*)	chi2	df	Prob > chi2
1	**0.027**	**1**	**0.8705**
2	**0.060**	**2**	**0.9705**
3	**0.069**	**3**	**0.9953**
4	**0.395**	**4**	**0.9829**
5	**0.395**	**5**	**0.9955**
6	**0.410**	**6**	**0.9988**
7	**0.428**	**7**	**0.9997**
8	**0.666**	**8**	**0.9996**
9	**0.828**	**9**	**0.9997**
10	**0.845**	**10**	**0.9999**

H0: no serial correlation

. dfuller dtransactions,lags(1)

Augmented Dickey–Fuller test for unit root   Number of obs = **79**

	Test Statistic	Interpolated Dickey–Fuller 1% Critical Value	5% Critical Value	10% Critical Value
Z(t)	**−7.064**	**−3.539**	**−2.907**	**−2.588**

MacKinnon approximate p–value for Z(t) = **0.0000**

We can see in Tables 5 and 6 that one lag was appropriate for the ADF test for both variables, and we reject the null hypothesis that each variable has a unit root. Thus, we conclude that dgoogle and dtransactions are stationary. Hence, Google and transactions are I(1) variables. Next we check to see if they are cointegrated. First, we regressed transactions on Google and saved the residuals in a new variable called ehat. Next, we test to see if ehat is stationary. Table 7 shows the steps in our ADF test.

Table 7:

Lags =1

Breusch–Godfrey LM test for autocorrelation

lags(*p*)	chi2	df	Prob > chi2
1	**0.706**	**1**	**0.4009**
2	**0.707**	**2**	**0.7021**
3	**1.148**	**3**	**0.7656**
4	**2.163**	**4**	**0.7059**
5	**2.217**	**5**	**0.8183**
6	**4.689**	**6**	**0.5843**
7	**4.956**	**7**	**0.6654**
8	**5.489**	**8**	**0.7043**
9	**5.501**	**9**	**0.7886**
10	**5.520**	**10**	**0.8539**

H0: no serial correlation

. dfuller ehat, noconstant lags(1)

Augmented Dickey–Fuller test for unit root   Number of obs = **80**

		Interpolated Dickey–Fuller		
	Test Statistic	1% Critical Value	5% Critical Value	10% Critical Value
Z(t)	**-1.879**	**-2.608**	**-1.950**	**-1.610**

We determined that one lag was appropriate. Then we ran an ADF test. We suppressed the constant because the mean of ehat should be zero. The ADF test has different critical values, called Engle-Granger critical values, when used for residuals of a prospective cointegrating regression than with a standard time series. The appropriate 5% critical value for a cointegration test is -3.337. Since our test statistic of -1.879 is less than the critical value, we cannot reject the null hypothesis. Thus, we conclude that Google and transactions are not cointegrated. Since our variables are I(1) and not cointegrated, the appropriate time series model is the VAR model.

Estimating a VAR Model

In searching for the best model, we want to use enough a lags such that we can minimize AIC and BIC and eliminate serial correlation to a reasonable degree. We decided that eliminating serial correlation for the first 15 lags (1 year and 1 quarter) would be enough. We use the varsoc command to compare VAR models with different lags. We started by comparing models with 4 lags as shown in the following table.

Table 8:

```
. varsoc dtransactions dgoogle, maxlag(4)

   Selection-order criteria
   Sample:  2010w41 - 2012w13                    Number of obs      =        77
```

lag	LL	LR	df	p	FPE	AIC	HQIC	SBIC
0	-900.295				5.2e+07	23.4362	23.4606	23.4971
1	-891.096	18.397*	4	0.001	4.5e+07	23.3012	23.3743*	23.4838*
2	-886.728	8.7361	4	0.068	4.5e+07*	23.2916*	23.4134	23.596
3	-884.518	4.42	4	0.352	4.7e+07	23.3381	23.5086	23.7643
4	-880.219	8.5989	4	0.072	4.7e+07	23.3304	23.5495	23.8783

```
   Endogenous:  dtransactions dgoogle
    Exogenous:  _cons
```

We see that 2 lags minimizes AIC and 1 lags minimizes BIC. So we estimated the model trying both 1 and 2 lags. Then we used the varlmar command to check for serial correlation in the first 15 lags. The output for varlmar is shown in Table 9.

Table 9:

```
. qui var dtransactions dgoogle if t>5, lags(1)

. varlmar, mlag(15)

   Lagrange-multiplier test
```

lag	chi2	df	Prob > chi2
1	8.1554	4	0.08605
2	9.4834	4	0.05009
3	1.9295	4	0.74872
4	4.9182	4	0.29579
5	0.6547	4	0.95680
6	5.3112	4	0.25683
7	0.2843	4	0.99081
8	0.5052	4	0.97299
9	1.2465	4	0.87038
10	0.8305	4	0.93431
11	1.2425	4	0.87105
12	1.2768	4	0.86530
13	1.1016	4	0.89402
14	1.2612	4	0.86792
15	0.7593	4	0.94382

H0: no autocorrelation at lag order

```
. qui var dtransactions dgoogle if t>5, lags(1/2)

. varlmar, mlag(15)

   Lagrange-multiplier test
```

lag	chi2	df	Prob > chi2
1	3.7212	4	0.44506
2	5.9716	4	0.20128
3	4.8691	4	0.30100
4	5.7938	4	0.21508
5	0.6057	4	0.96243
6	5.0818	4	0.27900
7	0.3477	4	0.98653
8	0.3422	4	0.98693
9	1.2116	4	0.87618
10	0.7675	4	0.94276
11	2.5649	4	0.63306
12	0.8468	4	0.93207
13	1.5642	4	0.81521
14	1.4227	4	0.84025
15	0.6416	4	0.95833

H0: no autocorrelation at lag order

We see that 2 lags eliminates serial correlation, but there is still serial correlation of the second order after using 1 lag. Thus, two lags is appropriate for the model. Table 10 shows our estimated model using two lags.

Table 10:

```
. var dtransactions dgoogle if t>5, lags(1/2)

Vector autoregression

Sample:  2010w41 - 2012w13                      No. of obs      =         77
Log likelihood = -886.7284                      AIC             =   23.29165
FPE            =  4.47e+07                      HQIC            =    23.4134
Det(Sigma_ml)  =  3.45e+07                      SBIC            =   23.59604

Equation           Parms      RMSE     R-sq      chi2     P>chi2
----------------------------------------------------------------
dtransactions        5     909.879   0.2524   25.99927   0.0000
dgoogle              5     7.24967   0.1135   9.859479   0.0429
----------------------------------------------------------------
```

	Coef.	Std. Err.	z	P>|z|	[95% Conf.	Interval]
dtransactions						
dtransactions						
L1.	-.2596049	.1149102	-2.26	0.024	-.4848247	-.034385
L2.	-.3159358	.1099629	-2.87	0.004	-.5314592	-.1004125
dgoogle						
L1.	67.16184	14.97193	4.49	0.000	37.8174	96.50627
L2.	27.66185	15.99648	1.73	0.084	-3.690674	59.01438
_cons	115.3852	101.0133	1.14	0.253	-82.59734	313.3677
dgoogle						
dtransactions						
L1.	.0005772	.0009156	0.63	0.528	-.0012173	.0023717
L2.	-.0013886	.0008762	-1.58	0.113	-.0031059	.0003286
dgoogle						
L1.	.2517644	.1192923	2.11	0.035	.0179557	.485573
L2.	.0911541	.1274557	0.72	0.474	-.1586544	.3409627
_cons	.1506612	.8048473	0.19	0.852	-1.426811	1.728133

After estimating our model, we used the vargranger command to perform the appropriate Granger causality test, as shown in the following table.

Table 11:

```
. vargranger

  Granger causality Wald tests
```

Equation	Excluded	chi2	df	Prob > chi2
dtransactions	dgoogle	**24.061**	**2**	**0.000**
dtransactions	ALL	**24.061**	**2**	**0.000**
dgoogle	dtransactions	**3.5014**	**2**	**0.174**
dgoogle	ALL	**3.5014**	**2**	**0.174**

The results of the Granger causality test suggest that dgoogle has a causal effect on dtransactions, but not vice versa. Graphs of our estimated impulse response function are shown in Figure 5. In Figure 5, we can see that a shock to transactions has no effect on Google hits, but a shock to dgoogle, showing a small increase of publicity, causes an increase of about 67 transactions

Figure 5:

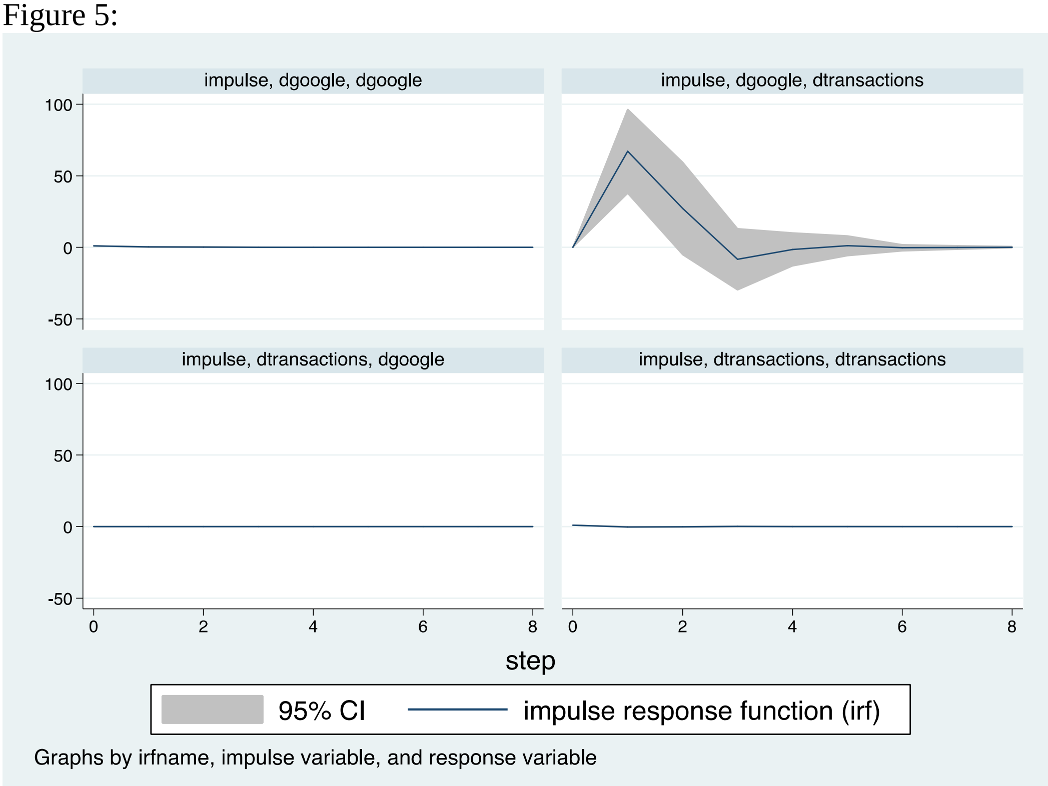

B. Relationship between Price Shocks and Total Transaction Value

In this section we explore the dynamic relationship between the price of bitcoins in dollars and the total value of Bitcoin transactions, also measured in dollars. We will examine the impulse response functions to assess how price shocks affect Bitcoin use.

The standard money demand function relies on consumption because individuals hold money in order to decrease transaction costs (necessary for consumption). Our variable "total transaction value" measures aggregate consumption in bitcoins over time. We convert this variable into dollars so that its value can be properly understood. We transformed both price and total transaction value into their log forms so that the first differences are interpretable as growth rates.

Table 1. Summary statistics logged variables

Variable	Obs	Mean	Std. Dev.	Min	Max
lprice	593	.4842691	1.712232	-2.821456	3.367455
lval_dol	593	31.44301	2.731991	25.41629	37.43035
dlprice	592	.0071683	.0717751	-.3575828	.3665798
dlval_dol	592	.0115607	.6189698	-3.663157	2.826921

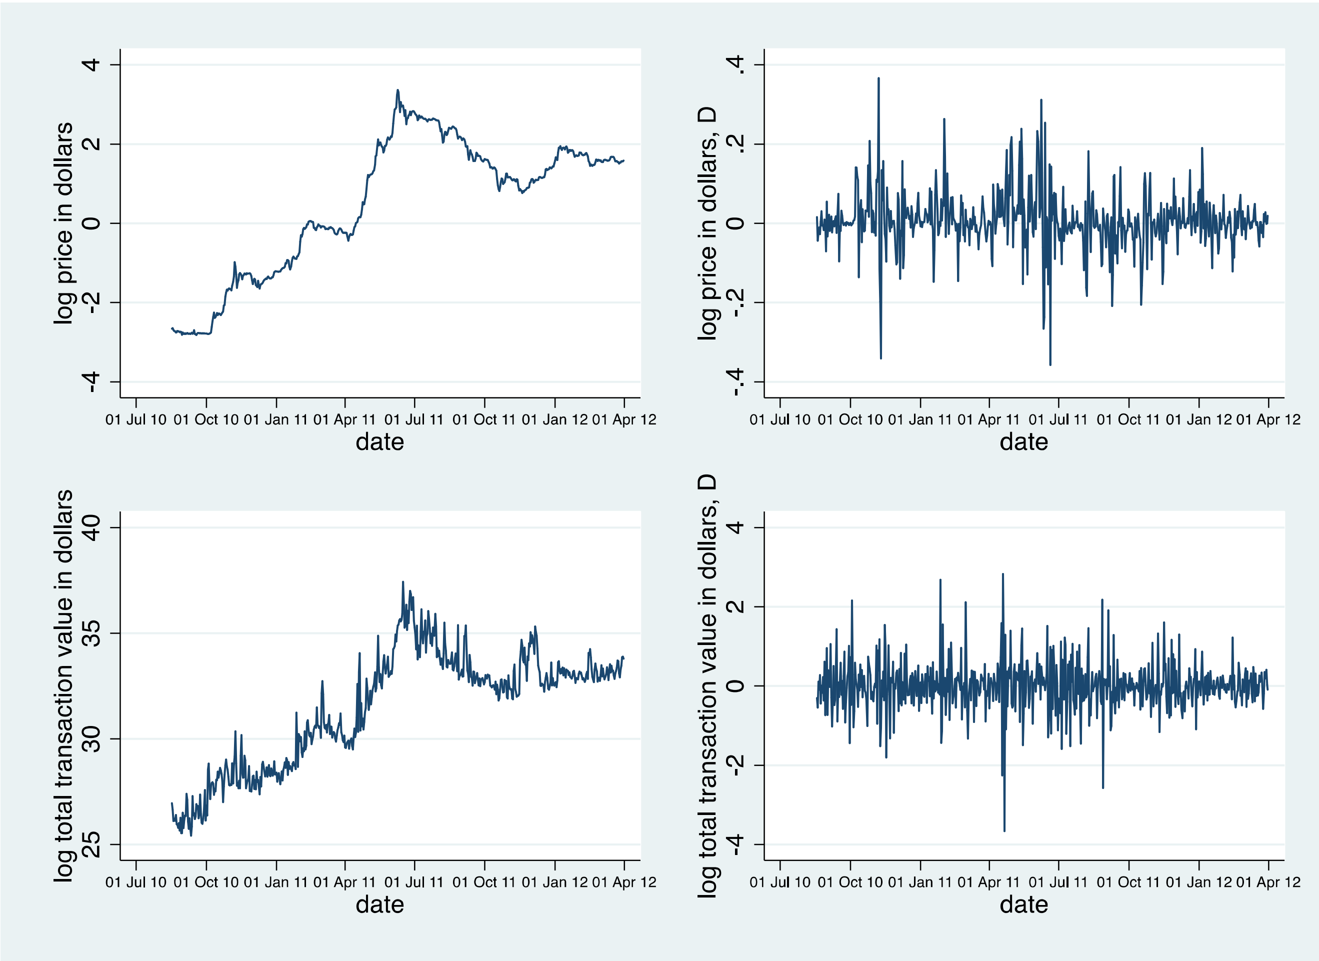

Figure 7. The logs of price and total transaction value (in dollars) over time, along with their first differences

We must first determine whether our variables are stationary or non-stationary. Qualitatively, neither series appears to be stationary, but it is possible they fluctuate around a trend. We can formally test for stationarity using a unit root test. We utilize an augmented Dickey-Fuller (ADF) test, which adds lagged first difference terms to eliminate autocorrelation in the errors.

In order to use the ADF test for the stationarity of lprice, we must specify the number of lags to include. We test numerous possible model specifications using the following do-file in Stata. The results are shown in Table 2.

```
. forvalues p=1/10 {
  2. qui reg dlprice L.lprice L(1/`p').dlprice
  3. display "p=`p'
  4. modelsel
  5. }
```

Table 2. Information criteria for possible lag specification for lprice

Lags	AIC	SC	Obs.
1	**-5.3266040**	**-5.3043613**	591
2	-5.3262737	-5.2965779	590
3	-5.3221085	-5.2849402	589
4	-5.3186563	-5.2739958	588
5	-5.3158029	-5.2636306	587
6	-5.3211457	-5.2614417	586
7	-5.3218337	-5.2545782	585
8	-5.3216832	-5.2468561	584
9	-5.3201039	-5.2376852	583
10	-5.3155356	-5.2255052	582

The information criteria, Akaike information criterion (AIC) and Schwarz criterion (SC), both indicate we should include no more than one lag term. We use a Breusch-Godfrey LM test, which checks for autocorrelation in the errors, to confirm that including one lag of the first difference of log price eliminates the serial correlation which would otherwise have biased our ADF test.

```
. reg dlprice L.lprice L.dlprice
. estat bgodfrey, lags(1)

Breusch-Godfrey LM test for autocorrelation
---------------------------------------------------------------------------
  lags(p)  |        chi2            df              Prob > chi2
-------------+-------------------------------------------------------------
      1    |        1.978            1                0.1596
---------------------------------------------------------------------------
                    H0: no serial correlation
```

The Breusch-Godfrey test agrees with the results of the information criteria, and thus we include one lag.

```
. dfuller lprice, regress lags(1)

Augmented Dickey-Fuller test for unit root         Number of obs   =       591

                               ---------- Interpolated Dickey-Fuller ---------
                  Test         1% Critical       5% Critical      10% Critical
               Statistic           Value             Value             Value
------------------------------------------------------------------------------
 Z(t)             -1.785            -3.430            -2.860            -2.570
------------------------------------------------------------------------------
MacKinnon approximate p-value for Z(t) = 0.3880
```

Because the approximate *p*-value for this test is greater than 0.05, we fail to reject the null hypothesis of nonstationarity. We now repeat this process on lval_dol to determine whether lval_dol is stationary or nonstationary. We use the same do-file (with appropriate adjustments) to test various lag specifications. The results are shown in Table 3.

```
. forvalues p=1/10 {
  2. qui reg dlval_dol L.lval_dol L(1/`p').dlval_dol
  3. display "p=`p'
  4. modelsel
  5. }
```

Table 3.

Lags	AIC	SC	Obs.
1	-1.0291495	-1.0069068	591
2	-1.0683272	-1.0386315	590
3	-1.0842398	-1.0470715	589
4	-1.096005	**-1.0513445**	588
5	-1.0919877	-1.0398154	587
6	-1.0957468	-1.0360428	586
7	-1.1144142	-1.0471587	585
8	-1.110934	-1.0361069	584
9	**-1.1188671**	-1.0364485	583
10	-1.117041	-1.0270107	582

SC is known to be more stringent than AIC, and it is expected that AIC would indicate a more liberal lag specification. We test both a 4-lag model and a 9-lag model using Breusch-Godfrey. In the case of 4-lags, we can reject the null hypothesis of no serial correlation in some cases. This is not the case when we test the 9-lag model. Hence, we determine 9 lags to be the appropriate specification. The results of the Breusch-Godfrey of 9 lags are shown below.

```
. reg dlval_dol L.lval_dol L(1/9).dlval_dol
```

. estat bgodfrey, lags(1 2 3 4 5 6 7 8 9)

Breusch-Godfrey LM test for autocorrelation

lags(p)	chi2	df	Prob > chi2
1	0.736	1	0.3908
2	0.741	2	0.6904
3	0.752	3	0.8609
4	0.803	4	0.9381
5	0.848	5	0.9739
6	1.423	6	0.9645
7	1.588	7	0.9791
8	1.672	8	0.9895
9	1.810	9	0.9941

H0: no serial correlation

We then run an ADF test using the 9-lag model.

```
. dfuller lval_dol, regress lags(9)

Augmented Dickey-Fuller test for unit root         Number of obs   =       583

                               ---------- Interpolated Dickey-Fuller ---------
                  Test         1% Critical       5% Critical      10% Critical
               Statistic           Value             Value             Value
------------------------------------------------------------------------------
 Z(t)             -1.874            -3.430            -2.860            -2.570
------------------------------------------------------------------------------
MacKinnon approximate p-value for Z(t) = 0.3444
```

Because the approximate *p*-value for this test is greater than 0.05, we do not reject the null hypothesis of nonstationarity. Therefore, we conclude that both series are nonstationary in their log-levels. Our next step is to use ADF to test the stationarity of the differenced series.

```
. dfuller dlprice, noconstant lags(1)

Augmented Dickey-Fuller test for unit root         Number of obs   =       590

                               ---------- Interpolated Dickey-Fuller ---------
                  Test         1% Critical       5% Critical      10% Critical
               Statistic           Value             Value             Value
------------------------------------------------------------------------------
 Z(t)            -15.568            -2.580            -1.950            -1.620

. dfuller dlval_dol, noconstant lags(9)

Augmented Dickey-Fuller test for unit root         Number of obs   =       582

                               ---------- Interpolated Dickey-Fuller ---------
                  Test         1% Critical       5% Critical      10% Critical
               Statistic           Value             Value             Value
```

```
------------------------------------------------------------------------------
 Z(t)            -10.353            -2.580            -1.950            -1.620
```

Because the test statistics are less than their critical values, we reject the null hypothesis of nonstationarity for the first differences. Therefore, we conclude that, lprice and lval_dol are integrated of order 1, or I(1).

Because both times series are I(1), we next test whether they are cointegrated – do the series tend to move together over time? To accomplish this, we utilize an Engle-Granger test. The Engle-Granger test regresses one I(1) variable on the other by OLS, then uses ADF to test the null hypothesis that the residuals are nonstationary.

```
. regress lprice lval_dol

      Source |       SS       df       MS              Number of obs =     593
-------------+------------------------------           F(  1,   591) = 6531.82
       Model |  1591.58334     1  1591.58334           Prob > F      =  0.0000
    Residual |  144.006677   591   .24366612           R-squared     =  0.9170
-------------+------------------------------           Adj R-squared =  0.9169
       Total |  1735.59001   592  2.93173989           Root MSE      =  .49363

------------------------------------------------------------------------------
      lprice |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    lval_dol |   .6001703    .007426    80.82   0.000     .5855857     .614755
       _cons |  -18.38689   .2343754   -78.45   0.000     -18.8472   -17.92658
------------------------------------------------------------------------------

. predict ehat, res
(31 missing values generated)

. dfuller ehat, noconstant lags(9)

Augmented Dickey-Fuller test for unit root         Number of obs   =       583

                               ---------- Interpolated Dickey-Fuller ---------
                  Test         1% Critical       5% Critical      10% Critical
               Statistic           Value             Value             Value
------------------------------------------------------------------------------
 Z(t)             -3.537            -2.580            -1.950            -1.620
```

Table 4. Critical values for the cointegration test (from HGL table 12.4 on pp. 489)

Model	1%	5%	10%
$y_t = \beta x_t + e_t$	-3.39	-2.76	-2.45

The Dickey-Fuller test statistic is less than the critical value, even at the 1% level. Hence we reject the null hypothesis that the residuals are nonstationary. Therefore, we conclude that lprice and lval_dol are cointegrated.

In order to analyze the cointegrated relationship, we estimate a vector error-correction (VEC) model. First, we search for the proper lag specification.

```
. varsoc dlprice dlval_dol, maxlag(30)

   Selection-order criteria
   Sample:  17 Sep 10 - 31 Mar 12               Number of obs      =       562
  +---------------------------------------------------------------------------+
  |lag |    LL      LR      df    p      FPE       AIC      HQIC      SBIC    |
  |----+----------------------------------------------------------------------|
  |  0 |  146.997                      .002046  -.516005  -.509987   -.50059  |
  |  1 |  188.086  82.178    4  0.000  .001793  -.647994   -.62994   -.60175  |
  |  2 |  203.881   31.59    4  0.000  .001719  -.689969  -.659879  -.612896* |
  |  3 |  212.073  16.383    4  0.003  .001694  -.704885  -.662758* -.596982  |
  |  4 |  219.015  13.885    4  0.008  .001676  -.715355  -.661193  -.576624  |
  |  5 |  221.253  4.4766    4  0.345  .001687  -.709086  -.642887  -.539525  |
  |  6 |  227.103    11.7    4  0.020  .001676  -.715669  -.637435   -.51528  |
  |  7 |   238.22  22.234    4  0.000  .001634  -.740997  -.650726  -.509778  |
  |  8 |   240.01    3.58    4  0.466  .001647  -.733132  -.630825  -.471084  |
  |  9 |  246.953  13.887    4  0.008   .00163* -.743606* -.629263  -.450729  |
  | 10 |  247.345  .78417    4  0.941  .001651  -.730767  -.604388   -.40706  |
  | 11 |  248.291  1.8904    4  0.756  .001669  -.719896   -.58148   -.36536  |
  | 12 |  249.726  2.8697    4  0.580  .001684  -.710767  -.560315  -.325402  |
  | 13 |  252.764  6.0779    4  0.193   .00169  -.707347  -.544859  -.291153  |
  | 14 |  253.981  2.4332    4  0.657  .001707  -.697442  -.522918  -.250418  |
  | 15 |  255.265  2.5671    4  0.633  .001724  -.687775  -.501215  -.209922  |
  | 16 |  256.377  2.2248    4  0.694  .001742  -.677498  -.478902  -.168817  |
  | 17 |  256.814  .87291    4  0.928  .001764  -.664817  -.454185  -.125306  |
  | 18 |  257.603  1.5792    4  0.813  .001784  -.653392  -.430724  -.083052  |
  | 19 |  259.561  3.9167    4  0.417  .001797  -.646126  -.411422  -.044957  |
  | 20 |  262.645  6.1667    4  0.187  .001803  -.642864  -.396124  -.010866  |
  | 21 |  264.998  4.7066    4  0.319  .001814  -.637004  -.378227   .025824  |
  | 22 |  267.658   5.319    4  0.256  .001823  -.632234  -.361421   .061423  |
  | 23 |  271.252  7.1878    4  0.126  .001826  -.630788   -.34794   .093698  |
  | 24 |  273.381  4.2593    4  0.372  .001838  -.624132  -.329247   .131183  |
  | 25 |  274.289  1.8148    4  0.770  .001859  -.613127  -.306206   .173018  |
  | 26 |  281.507  14.437*   4  0.006  .001838   -.62458  -.305623   .192394  |
  | 27 |  285.468  7.9227    4  0.094  .001838  -.624442  -.293449   .223361  |
  | 28 |  287.045  3.1525    4  0.533  .001854  -.615817  -.272787   .262815  |
  | 29 |  288.067  2.0451    4  0.727  .001874  -.605221  -.250155    .30424  |
  | 30 |  290.827  5.5189    4  0.238  .001883  -.600806  -.233705   .339484  |
  +---------------------------------------------------------------------------+
   Endogenous:  dlprice dlval_dol
    Exogenous:  _cons
```

Nine lags minimizes two of our criteria, FPE and AIC, while two other criteria result in fewer lags. Thus, we take nine lags to be the most likely appropriate specification for our VEC model. While we include 9 lags in the vec command in stata, the underlying VAR model, which requires one fewer lag, will only use 8.

```
. vec lprice lval_dol, lags(9) alpha
```

Vector error-correction model

```
Sample:  26 Aug 10 - 31 Mar 12                    No. of obs      =       584
                                                  AIC             = -.7846038
Log likelihood =  266.1043                        HQIC            = -.6766977
Det(Sigma_ml)  =  .0013781                        SBIC            = -.5077437

Equation           Parms      RMSE     R-sq      chi2     P>chi2
----------------------------------------------------------------
D_lprice             18     .069527   0.1089   69.16266   0.0000
D_lval_dol           18     .557891   0.2198   159.4543   0.0000
----------------------------------------------------------------

------------------------------------------------------------------------------
             |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
D_lprice     |
        _ce1 |
         L1. |  -.0073697   .0074789    -0.99   0.324    -.0220281    .0072887
             |
      lprice |
         LD. |   .2781615   .0426516     6.52   0.000      .194566     .361757
        L2D. |  -.0646237   .0439127    -1.47   0.141     -.150691    .0214436
        L3D. |  -.0130674   .0438501    -0.30   0.766     -.099012    .0728773
        L4D. |   .0249231   .0437735     0.57   0.569    -.0608713    .1107176
        L5D. |   .0257942   .0437632     0.59   0.556    -.0599802    .1115686
        L6D. |    .122515   .0435422     2.81   0.005     .0371739    .2078561
        L7D. |  -.0896402   .0437946    -2.05   0.041    -.1754761   -.0038044
        L8D. |   .0703818   .0424113     1.66   0.097    -.0127428    .1535064
             |
    lval_dol |
         LD. |   -.000171   .0065585    -0.03   0.979    -.0130255    .0126836
        L2D. |  -.0051439   .0065253    -0.79   0.431    -.0179333    .0076455
        L3D. |      .0019   .0064706     0.29   0.769    -.0107822    .0145822
        L4D. |  -.0064413   .0064386    -1.00   0.317    -.0190608    .0061781
        L5D. |   .0012646   .0062882     0.20   0.841    -.0110601    .0135892
        L6D. |  -.0053587   .0060574    -0.88   0.376    -.0172309    .0065135
        L7D. |   .0047898   .0057223     0.84   0.403    -.0064257    .0160053
        L8D. |  -.0032424   .0052905    -0.61   0.540    -.0136115    .0071268
             |
       _cons |   .0056523    .003047     1.86   0.064    -.0003198    .0116243
-------------+----------------------------------------------------------------
D_lval_dol   |
        _ce1 |
         L1. |   .2368706   .0600115     3.95   0.000     .1192503    .3544909
             |
      lprice |
         LD. |   .9785287   .3422404     2.86   0.004     .3077498    1.649308
        L2D. |  -.6344659   .3523598    -1.80   0.072    -1.325079    .0561467
        L3D. |    .398152   .3518577     1.13   0.258    -.2914764     1.08778
        L4D. |   .0397664   .3512428     0.11   0.910    -.6486569    .7281896
        L5D. |  -.3634944   .3511608    -1.04   0.301    -1.051757    .3247681
        L6D. |  -.0086925   .3493869    -0.02   0.980    -.6934782    .6760932
        L7D. |   .5014279   .3514123     1.43   0.154    -.1873275    1.190183
```

```
        L8D. |  -.1601195   .3403126    -0.47   0.638      -.82712     .506881
             |
    lval_dol |
         LD. |  -.2807132   .0526264    -5.33   0.000    -.3838591   -.1775673
        L2D. |  -.2256627   .0523599    -4.31   0.000    -.3282864   -.1230391
        L3D. |  -.1717799    .051921    -3.31   0.001    -.2735433   -.0700166
        L4D. |  -.1484872   .0516641    -2.87   0.004    -.2497469   -.0472275
        L5D. |  -.0523144   .0504572    -1.04   0.300    -.1512087    .0465799
        L6D. |  -.1174619   .0486048    -2.42   0.016    -.2127256   -.0221982
        L7D. |  -.1316266   .0459163    -2.87   0.004    -.2216209   -.0416324
        L8D. |  -.0117934   .0424514    -0.28   0.781    -.0949966    .0714097
             |
       _cons |   .0001759   .0244496     0.01   0.994    -.0477444    .0480961
------------------------------------------------------------------------------

Cointegrating equations

Equation           Parms    chi2     P>chi2
-------------------------------------------
_ce1                  1   364.3426   0.0000
-------------------------------------------

Identification:  beta is exactly identified

                 Johansen normalization restriction imposed
------------------------------------------------------------------------------
        beta |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
_ce1         |
      lprice |          1          .        .       .            .           .
    lval_dol |  -.6546289   .0342957   -19.09   0.000    -.7218473   -.5874105
       _cons |   20.19613          .        .       .            .           .
------------------------------------------------------------------------------

Adjustment parameters

Equation           Parms    chi2     P>chi2
-------------------------------------------
D_lprice              1   .9710093   0.3244
D_lval_dol            1   15.57952   0.0001
-------------------------------------------
------------------------------------------------------------------------------
       alpha |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
D_lprice     |
        _ce1 |
         L1. |  -.0073697   .0074789    -0.99   0.324    -.0220281    .0072887
-------------+----------------------------------------------------------------
D_lval_dol   |
        _ce1 |
         L1. |   .2368706   .0600115     3.95   0.000     .1192503    .3544909
------------------------------------------------------------------------------
```

The estimate of the coefficient [D_lval_dol] L._ce1 (shown in the adjustment parameters table) is .24 and statistically significant. This indicates that when price is out of equilibrium with total transaction value, total transaction value adjusts in the same direction as the price shock. This can be seen further in the following graphs of the impulse response functions.

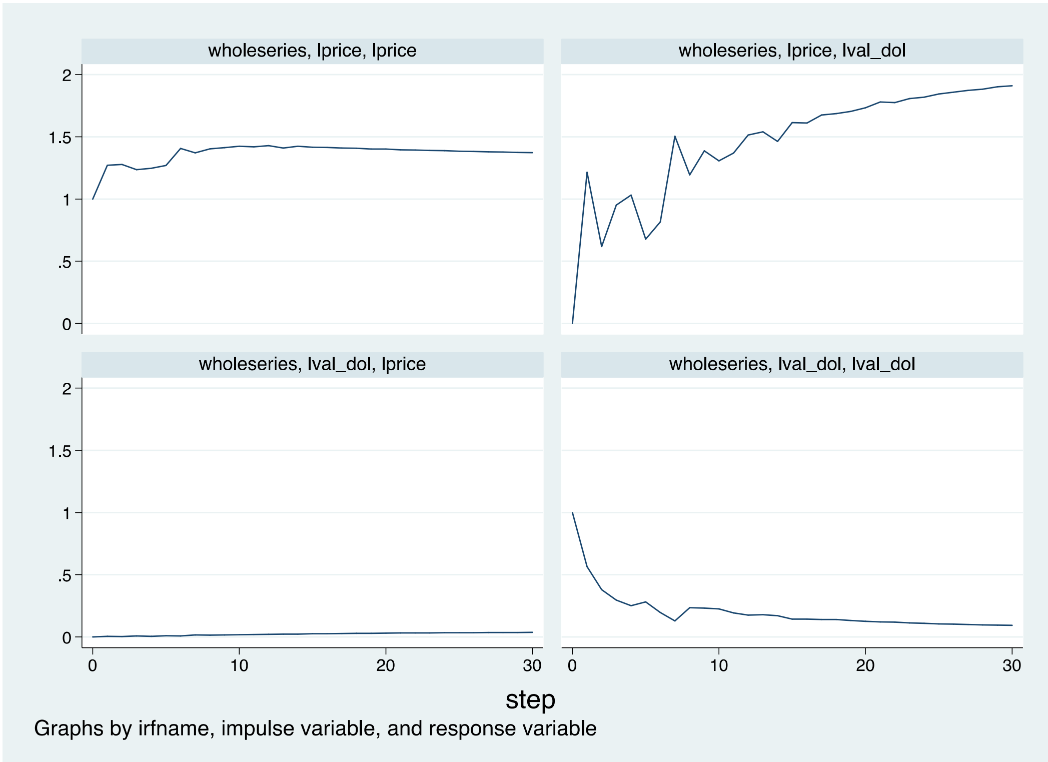

In the initial impulse response functions, we have constrained the contemporaneous effect of price on total transaction value to be zero. We assume that individuals using Bitcoin to engage in transactions are not continually aware of movements in the exchange rate with dollars.

The impulse responses reflect our earlier assessments of the VEC model results. The effect of shocks to transaction value on price is insignificant or zero. Shocks to total transaction value are transitory while shocks to price appear to be permanent. Total transaction value moves in the same direction as price shocks in order to return the system to equilibrium. Price shocks on total transaction value have permanent effects.

We are concerned that the responses of the variables to shocks may change before and after the price peak. The reason this might be the case is that while the price of bitcoins was rising, individuals viewed price shocks as opportunities to increase their stock of bitcoins, or as indications of increasing demand and value. In the words of one bitcoin enthusiast, "'I knew it wasn't a stock and wouldn't go up and down,' he explains. 'This was something that was going to go up, up, up.'" (Wallace 2011).

Sentiment such as this was proven misplaced after the peak and subsequent crash. Thus, we might expect individuals to be more cautious. To investigate this possibility, we re-ran the VEC model twice, once for those observations occurring before the peak, and once again for after. The following figures show the results.

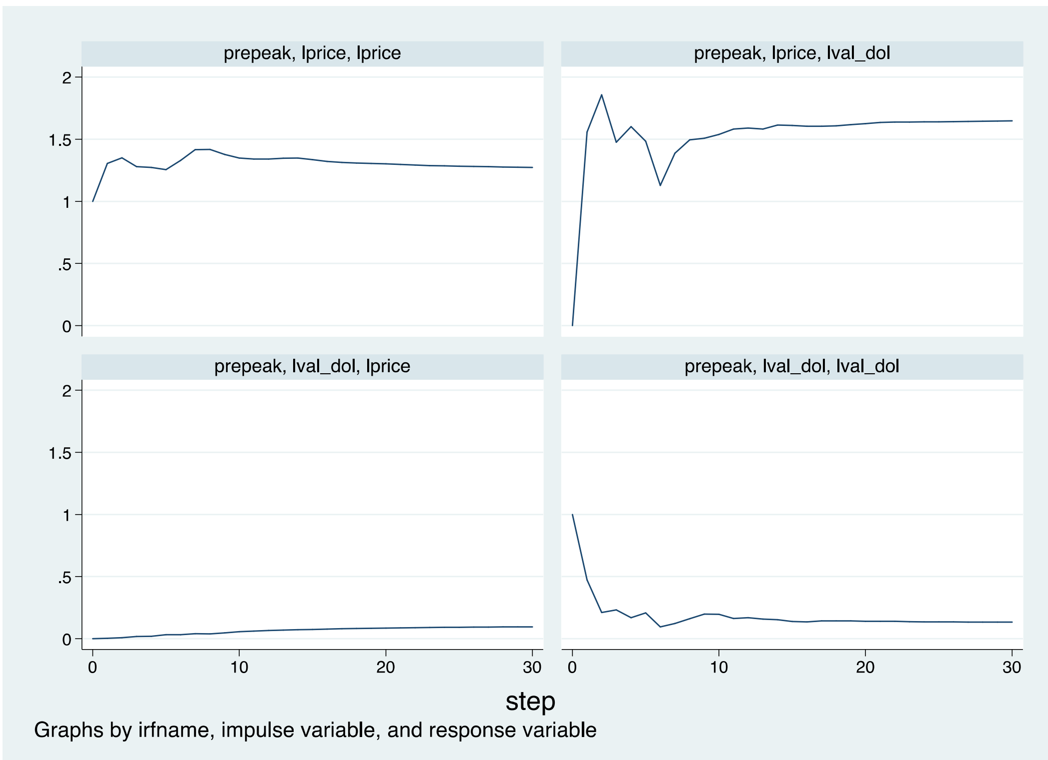

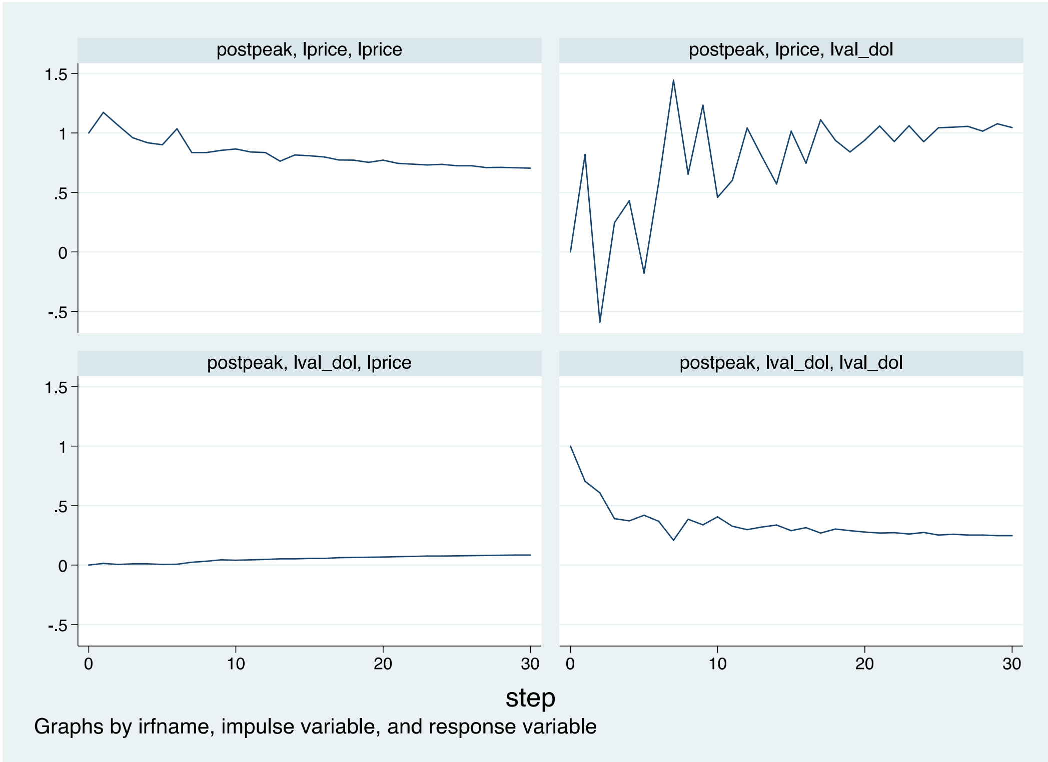

These impulse response functions support our earlier suspicions. Before the peak, transaction behavior responded to price shocks immediately and without much hesitation. After the peak, total transaction value spikes back and forth before the system moves in the direction of the price shock. In addition, the magnitude of the effect of price shocks on total transaction value is much less post-peak.

We test these results further using orthagonalized IRF. By doing this we remove the assumption of no contemporaneous correlation between the impulse and response variables. We do this because the immediate reaction of transaction behavior to price shocks may be of interest.

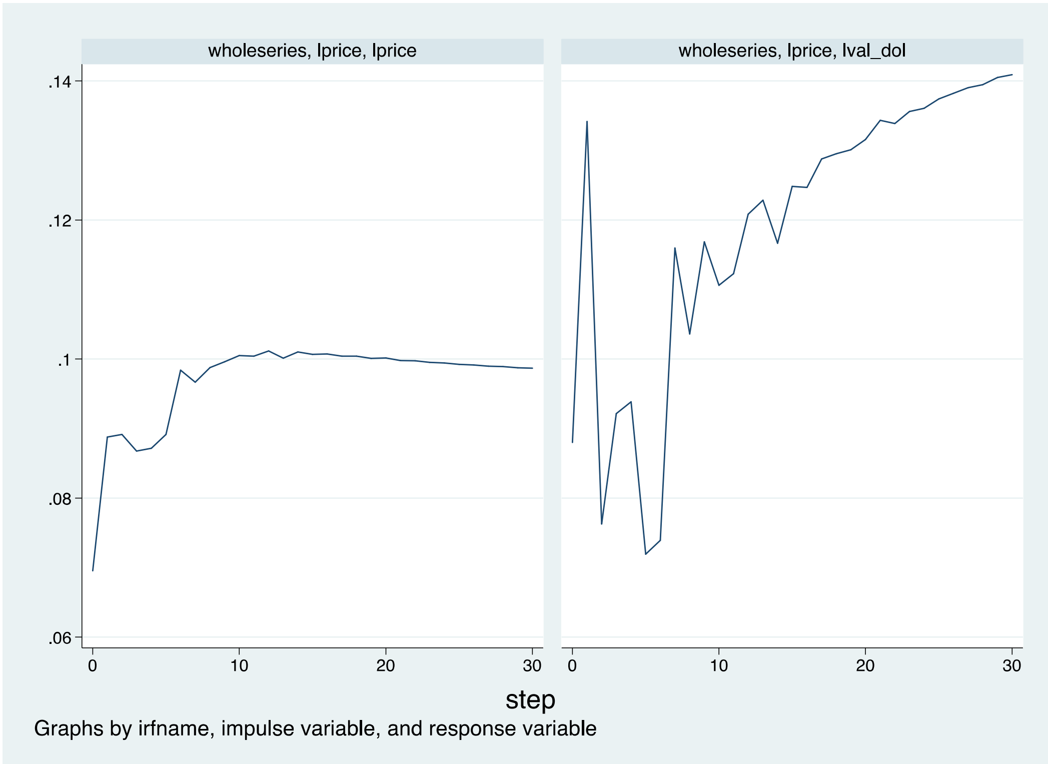
wholeseries, lprice, lprice
wholeseries, lprice, lval_dol
.14
.12
.1
.08
.06
0
10
20
30
step
Graphs by irfname, impulse variable, and response variable

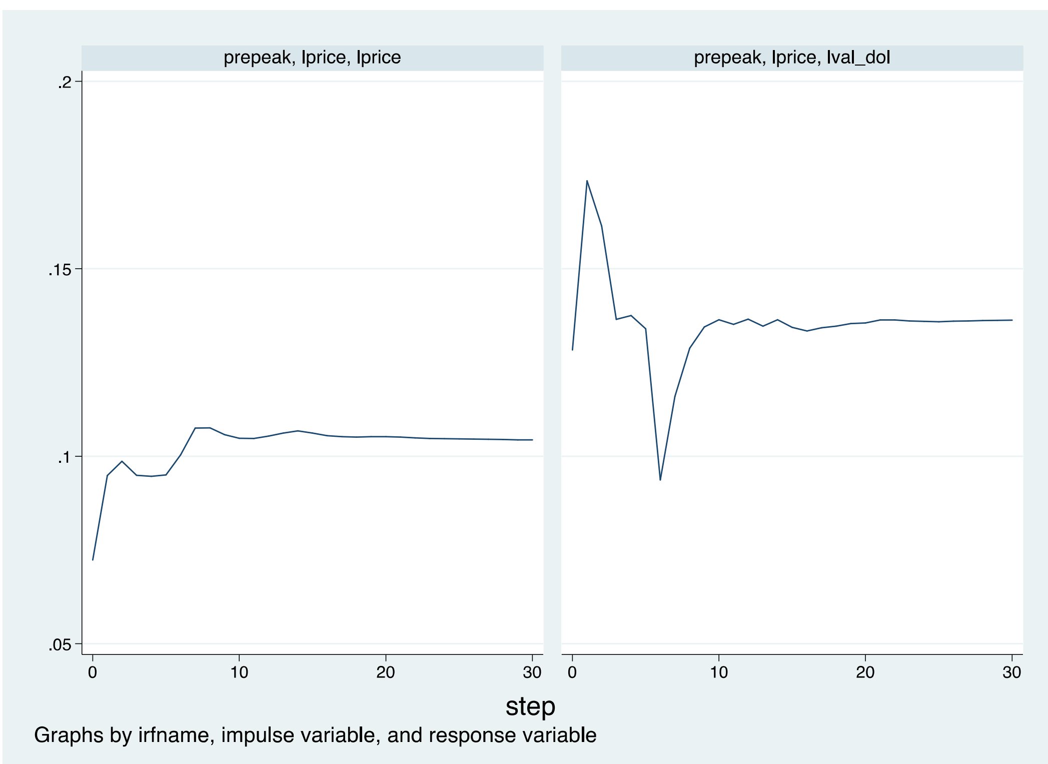
prepeak, lprice, lprice
prepeak, lprice, lval_dol
.2
.15
.1
.05
0
10
20
30
step
Graphs by irfname, impulse variable, and response variable

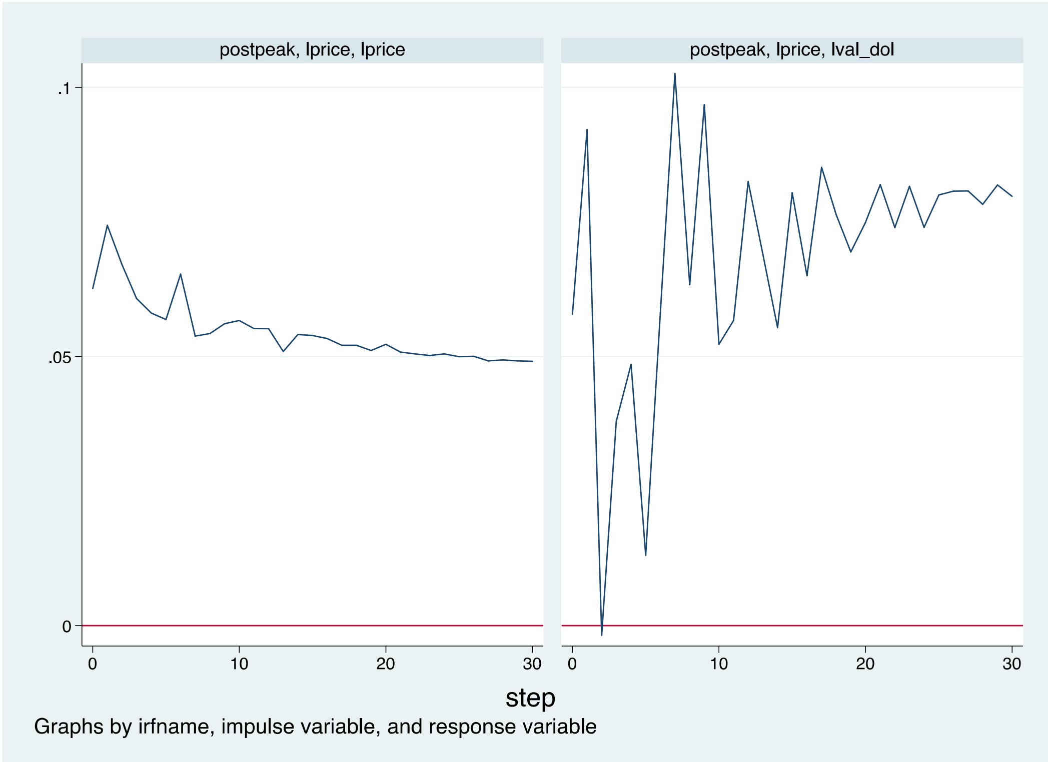

As can be seen, our qualitative results are similar. The orthagonalized results show total transaction value responding even slower to price shocks than the previous IRF. A limitation of the approach in this section is that it does not take into account possible changes in the variance of the responses. It is also possible that users of Bitcoin respond differently to positive price shocks than to negative price shocks. These are potentially informative areas to investigate, which we will begin to examine in the next section.

C. GARCH-In-Mean Model

Here we investigate the volatility of price and its effect on price, which we use as a proxy for demand for bitcoins. We use ARCH and GARCH models to model the effects of volatility on price. Particularly, in line with the rest of our paper, we are interested in the effects of volatility on demand.

We first set the data up for the ARCH model. Previously we showed that log(price) is nonstationary. Standard ARCH and GARCH models assume stationarity. We correct for nonstationarity by producing first differences. First differences are consistently used in GARCH models, and Vale (2004) performs a GARCH-In-Mean model using first difference data.

We first test to see if there are potential ARCH effects. Although we have addressed this earlier in the paper, we reproduce the graph of the first difference of price below make an initial gauge:

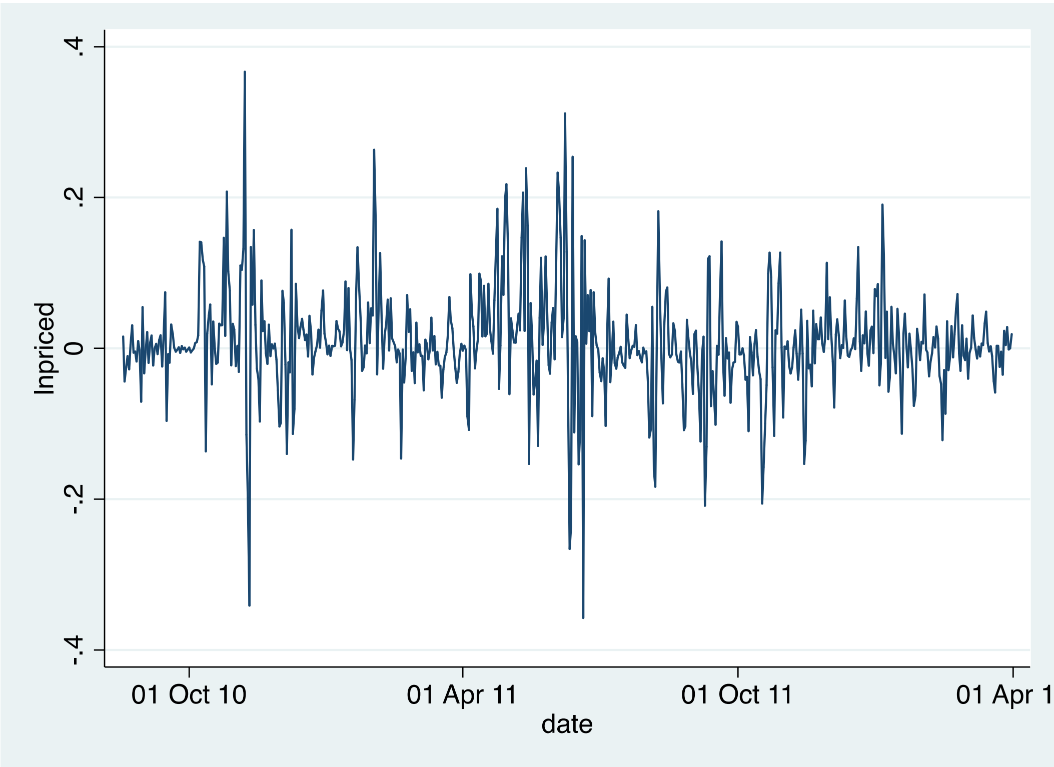

The graph shows that there may be ARCH effects. Especially volatile times are clustered together.

```
. estat archlm, lags(1)
LM test for autoregressive conditional heteroskedasticity (ARCH)
---------------------------------------------------------------------------
    lags(p)  |          chi2               df                 Prob > chi2
-------------+-------------------------------------------------------------
       1     |         62.215               1                    0.0000
```

```
---------------------------------------------------------------------------
         H0: no ARCH effects      vs.  H1: ARCH(p) disturbance
```

From here we note that there are ARCH effects with the first difference price data. We reject the null hypothesis that there are no ARCH affects.

We wish to see if price volatility has an effect on price of currency. Common literature, as described by H.G.L., states that in the case of stock returns, which are a common subject matter for ARCH models, volatility has a positive effect on stock returns because higher volatility will lead to a higher risk premium. We believe that price volatility will have a negative effect on the price of currency. There is no potential risk premium effect that we know of with bitcoins, and so we believe that potential holders of currency will start selling bitcoins in response to higher volatility. We perform a GARCH-In-Mean model as described by H.G.L. From here onward we use GARCH models at first to see if the garchL1. coefficient is statistically significant. This is because the garchL1. coefficient captures lags much farther back and thus explains the momentum come from previous lags. We look at our results and compare it to hypothesis by running a GARCH-In-Mean model.

```
Sample: 18 Aug 10 - 31 Mar 12                   Number of obs   =       592
Distribution: Gaussian                          Wald chi2(1)    =      4.69
Log likelihood =  840.9978                      Prob > chi2     =    0.0303

------------------------------------------------------------------------------
             |                 OPG
    lnpriced |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
lnpriced     |
       _cons |  -.0010876   .0028902    -0.38   0.707    -.0067523    .0045771
-------------+----------------------------------------------------------------
ARCHM        |
      sigma2 |   1.331854    .614861     2.17   0.030     .1267486     2.53696
-------------+----------------------------------------------------------------
ARCH         |
        arch |
         L1. |    .379535   .0689776     5.50   0.000     .2443413    .5147286
             |
       tarch |
         L1. |   .2503921   .0959487     2.61   0.009      .062336    .4384481
             |
       garch |
         L1. |   .4568105   .0374615    12.19   0.000     .3833874    .5302337
             |
       _cons |   .0006515   .0000727     8.96   0.000      .000509    .0007941
------------------------------------------------------------------------------
```

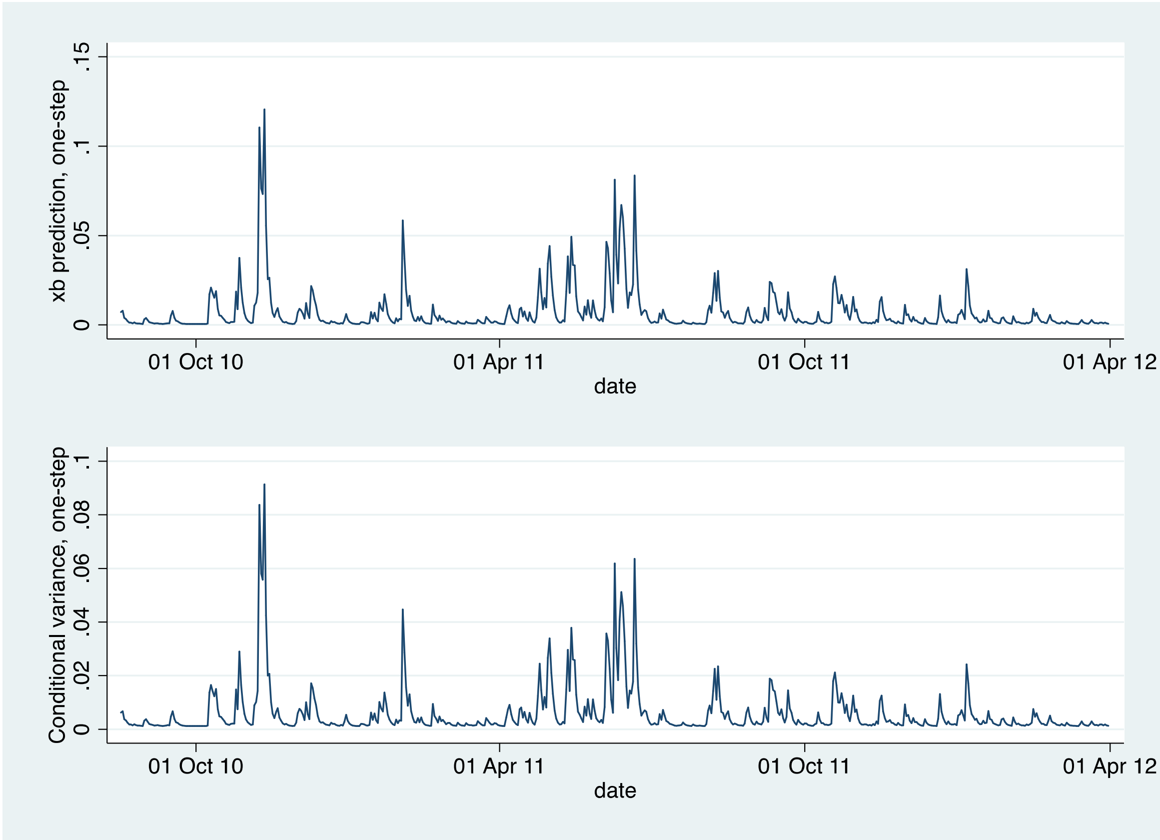

The graph above shows that the predictions of mean and variance take very similar patterns.

sigma2 is the coefficient for the effect of volatility on the first difference of logged price. Note that the sigma2 variable is a positive 1.331854 with a statistically significant p-value of 0.03. It is a very curious result because the results state that an increased volatility leads to increased price.

We perform more GARCH tests to see the nature of the volatility, in hopes to find clues as to why volatility might affect price positively. We perform a T-GARCH test. The importance of the T-GARCH test is that it accommodates asymmetry in the types of shocks involved. The effects of a negative shock on volatility are separated from the effects of a positive shock on volatility.

```
ARCH family regression

Sample: 18 Aug 10 - 31 Mar 12                     Number of obs   =       592
Distribution: Gaussian                            Wald chi2(.)    =         .
Log likelihood =  838.7289                        Prob > chi2     =         .

------------------------------------------------------------------------------
             |                 OPG
    lnpriced |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
lnpriced     |
       _cons |   .0025982    .002394     1.09   0.278    -.0020939    .0072903
-------------+----------------------------------------------------------------
ARCH         |
        arch |
         L1. |   .4456171   .0795054     5.60   0.000     .2897893    .6014449
             |
       tarch |
         L1. |   .1604125   .0941781     1.70   0.089    -.0241733    .3449982
             |
       garch |
         L1. |   .4324251   .0363378    11.90   0.000     .3612043     .503646
             |
       _cons |   .0006786   .0000695     9.77   0.000     .0005424    .0008148
```

This leads to a more illuminating result. We notice that from the tarch coefficient   that the extra effect of a positive shock has a positive coefficient, which implies that if the shock is positive, volatility would be affected by a factor of (0.1604125 + 0.4456171) = 0.6060296, while the effect of a negative shock is simply a factor of -0.4456171. If we can accept this with a 10% level of significance, this means that that negative shocks have less of an effect on volatility than positive shocks. This could explain why volatility leads to an increase in price. If we try to stretch out argument, most of the long lasting and heavy volatility comes from positive shocks, and these positive shocks positively affect price. This is the complete opposite of the typical financial market analyzed by ARCH and GARCH models, where negative shocks lead to much more and much longer volatility than positive shocks, which reach equilibrium quickly.

However, the fact that it is only significant under a 10% level of significance, meaning that the volatility effects of a price shock are symmetric under a 5% level of significance, indicates that we should be looking closer at the data. The fact that volatility is symmetric to both positive and negative shocks as well as a positive effect of volatility on price is a problematic result, and we feel we may not know a crucial part of the story. As we explained before, the history of bitcoin in 2010-2011 states that there was a bubble that brewed up until the 9th of June, 2011, where the price peaked at 29.004612 dollars per bitcoin. All the while, speculators, arbitragers, and other market participants were shorting and making money of the growing bubble. After the point the bubble burst, marked by the peak on the 9th of June, 2011, the market participants realized they could lose money afterwards, and those that feared losing money left the novelty of bitcoins and moved out of the currency market. We hypothesize that our GARCH-In-Mean results will be different if we split our time series before and after this peak. We first run a M-GARCH test with the time series before 9th of June:

Before 9th of June 2011

```
ARCH family regression

Sample: 18 Aug 10 - 09 Jun 11                     Number of obs   =        296
Distribution: Gaussian                            Wald chi2(1)    =       6.83
Log likelihood =  409.5917                        Prob > chi2     =     0.0090

------------------------------------------------------------------------------
             |                 OPG
    lnpriced |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
lnpriced     |
       _cons |   .0001576   .0046968     0.03   0.973     -.009048    .0093632
-------------+----------------------------------------------------------------
ARCHM        |
      sigma2 |   2.772423   1.060783     2.61   0.009     .6933274    4.851519
-------------+----------------------------------------------------------------
ARCH         |
        arch |
         L1. |   .1337787   .0596704     2.24   0.025     .0168268    .2507306
             |
       tarch |
         L1. |   .5285886    .132724     3.98   0.000     .2684544    .7887228
             |
       garch |
         L1. |   .5464072    .049737    10.99   0.000     .4489243      .64389
             |
       _cons |   .0006224   .0001061     5.87   0.000     .0004144    .0008303
------------------------------------------------------------------------------
```

After 9th of June, 2011

```
ARCH family regression

Sample: 10 Jun 11 - 31 Mar 12                     Number of obs   =        296
Distribution: Gaussian                            Wald chi2(1)    =       0.43
Log likelihood =  439.9263                        Prob > chi2     =     0.5124

------------------------------------------------------------------------------
             |                 OPG
    lnpriced |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
lnpriced     |
       _cons |  -.0000829   .0039126    -0.02   0.983    -.0077514    .0075857
-------------+----------------------------------------------------------------
ARCHM        |
      sigma2 |  -.6714396   1.025027    -0.66   0.512    -2.680456    1.337577
-------------+----------------------------------------------------------------
ARCH         |
        arch |
         L1. |   .9605993   .1816795     5.29   0.000     .6045141    1.316685
             |
       tarch |
         L1. |  -.6677757   .1952373    -3.42   0.001    -1.050434   -.2851176
             |
       garch |
         L1. |    .406956   .0532386     7.64   0.000     .3026103    .5113017
             |
       _cons |   .0005581   .0001023     5.45   0.000     .0003576    .0007586
------------------------------------------------------------------------------
```

WOW! The results confirm our hypothesis. Before the bubble burst, the effect of volatility on price is positive statistically significant with a sigma2 coefficient of

2.772423. After the bubble burst, the effect of volatility is statistically insignificant with a p-value of .512. We believe that this implies that volatility led to a demand for the currency, and after the bubble burst the novelty of bitcoin left, and only market participants who were averse to volatility stayed in the market, leading to no effect of volatility on price.

Finally, we look at T-GARCH models split before and after the peak. Before the peak, which is the first set of results, we see that a unit negative shock leads to a -.2023101 change in volatility, and a positive shock leads to a (.2023101 + .49099) = 0.6933001. We see here that the market has a higher volatility due to a positive shock but a low volatility due to a negative shock, showing the evidence of market bubble mentality.

```
ARCH family regression

Sample: 18 Aug 10 - 09 Jun 11                     Number of obs   =       296
Distribution: Gaussian                            Wald chi2(.)    =         .
Log likelihood =   404.166                        Prob > chi2     =         .

------------------------------------------------------------------------------
             |                 OPG
    lnpriced |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
lnpriced     |
       _cons |   .0085614   .0035756     2.39   0.017     .0015533    .0155694
-------------+----------------------------------------------------------------
ARCH         |
        arch |
         L1. |   .2023101   .0945331     2.14   0.032     .0170287    .3875915
             |
       tarch |
         L1. |     .49099   .1469145     3.34   0.001     .2030428    .7789372
             |
       garch |
         L1. |   .4528901   .0534543     8.47   0.000     .3481216    .5576586
             |
       _cons |   .0008033   .0001133     7.09   0.000     .0005812    .0010253
------------------------------------------------------------------------------
```

We then see that after the peak, market participants mimic the behavior of a financial market as described by H.G.L.: in other words, the market reaches equilibrium much more quickly after a positive shock but volatility is more strongly affected and more persistent with a negative shock. This shows that after the peak the market has shifted away from market bubble behavior to that of a typical financial market.

```
ARCH family regression

Sample: 10 Jun 11 - 31 Mar 12                     Number of obs   =       296
Distribution: Gaussian                            Wald chi2(.)    =         .
Log likelihood =  439.6181                        Prob > chi2     =         .

------------------------------------------------------------------------------
             |                 OPG
    lnpriced |      Coef.   Std. Err.      z    P>|z|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
lnpriced     |
       _cons |  -.0015772   .0028771    -0.55   0.584    -.0072162    .0040617
-------------+----------------------------------------------------------------
```

ARCH						
arch						
L1.	.9361392	.1752296	5.34	0.000	.5926954	1.279583
tarch						
L1.	-.6174132	.1916837	-3.22	0.001	-.9931063	-.2417201
garch						
L1.	.3952406	.0519775	7.60	0.000	.2933666	.4971146
_cons	.0005737	.0001019	5.63	0.000	.0003739	.0007734

Conclusion

In conclusion, we learn from running our ARCH/GARCH models that before the peak of the bubble, volatility had a statistically significant positive effect on price. This makes sense because the market bubble implies that, coupled with the positive spirits of a market bubble, speculators, arbitragers, miners, and other market participants caught in the hype viewed the volatility in a positive light as a method to make large amounts of quick money. After the bubble burst, we see that market participants feared holding bitcoins because many realized that they could lose their wealth due to fluctuations in bitcoin price. Only the ones that stayed were tolerant of risk, which is why our sigma2 coefficient was statistically insignificant after the market bubble peak. Furthermore, we also notice in our TGARCH models that before the peak of the market bubble, there were asymmetrical effects to positive and negative shocks. Particularly, there was significantly less volatility as a consequence of negative shocks than there were as a consequence of positive shocks. This implies market bubble and speculative behavior. After the bubble peak, we notice that a correction occurs and the market responds quickly into equilibrium after a positive shock, but responds with high volatility after a negative shock. With price (ie. the price of bitcoins in US dollars) as a proxy for demand, we see how volatility significantly effects demand, with price increases implying demand increases and price decreases implying demand decreases. Altogether, we have a strong explanation and validation of the existence of a market bubble in the bitcoin currency market.

Validity Issues

Data

Our first issue comes from the data. While we have a large number of observations, our observations do not go longer than a year and a half. Considering the age of the currency, it would be difficult to get data much longer than a few years. Furthermore, our data is collected from an open source website, and so verification issues do exist. We also only have weekly data of Google hits, and we are missing so many variables with other measures of publicity such as RSS and Lexis Nexus that we are unable to use them effectively. The weekly form of Google hits required us to decrease our sample set to only weekly observations in order to test for the effects of Google hits. This leads us to our second issue.

Stationarity

While we can correct for stationarity, we might have stationarity issues for data that we cannot adequately test for. For example, we concluded that the data for Google hits was stationary, but we acknowledge that Google hits could be nonstationary if we had enough observations.

Heteroskadasticity

While heteroskadasticity is not too much of problem with our ARCH and GARCH models, the fact that we can clearly see heteroskedasticity in our ARCH and GARCH applications show the existence of heteroskedasticity could affect our VEC and VAR models. Being able to correct for them might give us more accurate results.

References

Arthur, Charles. "Bitcoin Value Crashes below Cost of Production as Broader Use Stutters." *The Guardian*. Guardian News and Media, 18 Oct. 2011. Web. 05 May 2012. <http://www.guardian.co.uk/technology/2011/oct/18/bitcoin-value-crash-cryptocurrency?newsfeed=true>.

Barber, Simon, Xavier Boyen, Elaine Shi, and Ersin Uzun. "Bitter to Better — How to Make Bitcoin a Better Currency." *Palo Alto Research Center, University of California, Berkeley* (2012).

"Bitcoin." *Wikipedia*. Wikimedia Foundation, 05 May 2012. Web. 03 May 2012. <http://en.wikipedia.org/wiki/Bitcoin>.

Chapman, Stephen. "Bitcoin: A Guide to the Future of Currency." *Technology News, Analysis, Comments and Product Reviews for IT Professionals*. ZD NET (CBS), 15 June 2011. Web. 05 May 2012. <http://www.zdnet.com/blog/btl/bitcoin-a-guide-to-the-future-of-currency/50601>.

Golder, Peter N., and Gerard J. Tellis. "Will It Ever Fly? Modeling the Takeoff of Really New Consumer Durables." *Marketing Science* 16.3 (1997): 256-70.

Goodhart, Charles A. E. *Money, Information and Uncertainty*. Basingstoke: Macmillan, 1989.

Greenberg, Andy. "Crypto Currency." *Forbes.com*. Forbes Magazine, 09 May 2011. Web. <http://www.forbes.com/forbes/2011/0509/technology-psilocybin-bitcoins-gavin-

andresen-crypto-currency.html>.

Grinberg, Reuben. "Bitcoin: An Innovative Alternative Digital Currency." *Hastings Science & Technology Law Journal, Vol. 4, P.160* 4 (2011): 160-209.

Hall, Bronwyn H. "Innovation and Diffusion." *The Oxford Handbook of Innovation* (2005): 459-567.

Jeffries, Adrianne. "The New York Observer GA_googleFillSlot("Obs_Article_Bookend_Top_right");." *Bit Oâ™Money: Whoâ€™s Behind the Bitcoin Bubble?* The New York Observer, 14 July 2011. Web. 02 May 2012. <http://www.observer.com/2011/06/bit-omoney-whos-behind-the-bitcoin-bubble/>.

McCallum, Bennett T. *Monetary Economics: Theory and Policy*. New York: Macmillan, 1989.

Nakamoto, Satoshi. "Bitcoin: A Peer-to-Peer Electronic Cash System." *www.bitcoin.org*. Web.

Nelson, R. R. "Why and How Innovations Get Adopted: A Tale of Four Models." *Industrial and Corporate Change* 13.5 (2004): 679-99.

Paul, Krugman. "Golden Cyberfetters." *Paul Krugman Blog*. The New York Times, 07 Sept. 2011. Web. 05 May 2012. <http://krugman.blogs.nytimes.com/2011/09/07/golden-cyberfetters/>.

Selgin, George. "Quasi-Commodity Money." Liberty Fund Symposium. Freidburg, Germany. 02 May 2012.

Surowiecki, James. "Cryptocurrency." *Technology Review*. MIT, Sept.-Oct. 2011. Web. 06 May 2012. <http://www.technologyreview.com/computing/38392/>.

Wallace, Benjamin. "The Rise and Fall of Bitcoin." *Wired.com*. Conde Nast Digital, 23 Nov. 2011. Web. 05 May 2012. <http://www.wired.com/magazine/2011/11/mf_bitcoin/all/1>.

White, Lawrence Henry. *The Theory of Monetary Institutions*. Malden [u.a.: Blackwell, 1999.