Reviewing “Modelling Bitcoin’s Value with Scarcity” - Part I to IV

An extensive review of PlanB’s work: does the popular model hold up according to everything we know about statistical modelling?

Foreword

This document is based on all earlier research provided by BurgerCrypto.com which was published between June 2019 and May 2020 by means of different Medium articles. The reader may notice that the research flips from falsifying the model to supporting the model and ends up concluding that we are not able to show that there is enough evidence to be convinced about absence of a spurious relationship. This material should not be considered investment, legal or financial advice. BurgerCrypto.com can under no circumstances be held liable for any financial loss incurred as a result of using this research.

Part I: The first review of the S2F model

Introduction

Recently Plan B, an early Dutch bitcoin adopter who anonymously shares his insights via Twitter, came out with an article [1] that dives into the relation between bitcoin’s price (and market cap) and the stock-to-flow ratio. It has attracted a lot of attention and the crypto community really praises his analysis. People fell in love with the conclusions that followed from his analysis, which is easy to understand as it predicts future price levels of six figures and more. Even though the idea of investigating market capitalisation or price as a function of stock-to-flow ratio’s is very interesting, one should respect the underlying model assumptions that hold for ordinary least squares regression. In this article I’ll pinpoint the flaws of the analysis, and I’ll have a look at alternative methods to come up with an improved model in a follow up article. I can’t guarantee I can find a model that is statistically significant and respects all the imposed assumptions though.

Respecting the assumptions

Plan B attempted to fit a model that describes the relation between the natural logarithm of bitcoins market cap and the natural logarithm of the stock-to-flow ratio by means of ordinary least squares regression. The standard set of 4 (Gaus-Markov) assumptions that needs to be respected (see Verbeek, Modern Econometrics [2]) are as follows:

(1) The expected value of the error term is zero (which means that on average the regression should be correct)

(2) The error and the independent variables should be independent.

(3) The error term should be homoskedastic (i.e. error terms have the same variance)

(4) There should be zero correlation between different error terms (i.e. autocorrelation is excluded)

Together, (1), (3) and (4) imply that the error terms are uncorrelated drawings from a distribution with expectation zero and constant variance. If the above assumptions are not met, the standard errors of the coefficients might be biased and therefore the results of the significance tests don't mean that much anymore. In order to draw conclusions based on the ANOVA results, we also have to check if the underlying model assumptions hold.

A review of plan B's model

To make sure the review is appropriate, I have used Plan B's dataset. In my own dataset I'll introduce some improvements, but for now his dataset is used. After running the exact same regression analysis in which the natural logarithms of the stock-to-flow of bitcoin acts as the independent variable to explain the natural logarithm of the market cap of bitcoin as the dependent variable, I found the same test results.

Checking the assumptions: Autocorrelation

A quick glance at the residual plots (one vs time, the other vs the predicted values), shows that there is some kind of periodic effect observable in the second scatter plot.

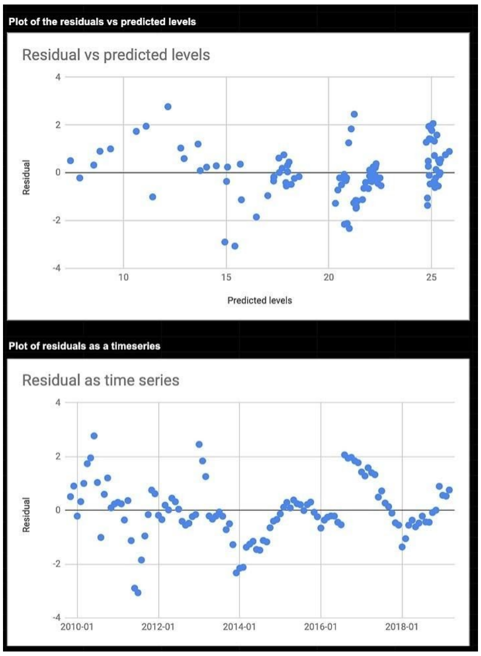

Residual analysis vs time and vs predicted levels

As the second chart clearly shows a certain pattern, we have to test for autocorrelation. Both the Durbin-Watson test and the test for

autoregression (by regressing the residual terms on their lagged values in case of first order autocorrelation) show we have to reject the hypothesis that there is no first order autocorrelation in the error terms.

The Durbin Watson statistic turned out to be ~0.4 which was much smaller than the lower bound value of 1.654~1.72 (see the table [4] in the appendix) which we would have to use for our dataset as the 5% lower bound of the critical value. The anova results for the regression test are also shown below.

Regression Statistics						
Multiple R	0.7967272841					
R^2	0.6347743652					
Standard Error	0.6416202174					
Adjusted R^2	0.6314840441					
Observations	113					
ANOVA						
	df	*SS*	*MS*	*F*	*Significance of F*	
Regression	1	79.4213356	79.4213356	192.9217114	0	
Residual	111	45.69609188	0.4116765034			
Total	112	125.1174275				
	Coefficients	*Standard Error*	*t-Statistics*	*p-Value*	*Lower 95%*	*Upper 95%*
Intercept	-0.00152920752(	0.06035893106	-0.02533523208	0.979833	-0.1211344588	0.1180760438
error_t-1	0.7964674401	0.05734261225	13.88962604	0.000000	0.6828392259	0.9100956543

Regression results of regressing error_(t) on error_(t-1)

Or in layman terms; we have to find a better model as the significance tests are biased as a result of serial correlation in the residual. The model indicates that the error at time t is a function of 0.796 times the error at time t-1 plus some noise.

This should not come as a complete surprise as both time series are non-stationary. There is a clear upward trend for both time series as can be seen in the graphs below. One way to deal with this can be found in differencing the time series or in case of asset price evolution looking at the periodic returns instead. This is however not a guarantee that the problem of autocorrelation is solved immediately or that the model returns significant estimators.

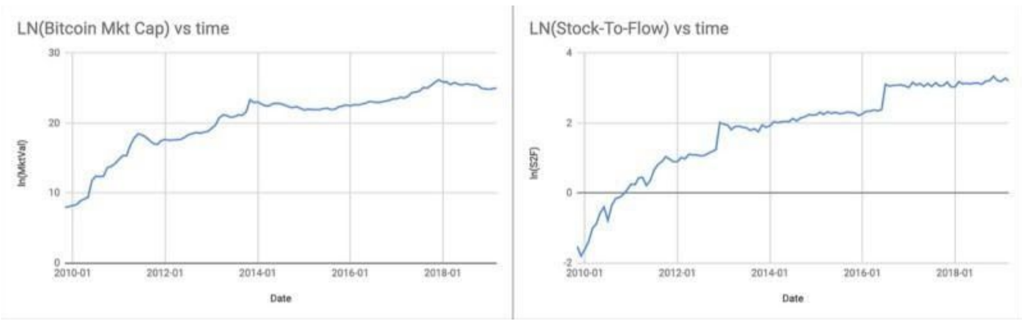

The log S2F and Bitcoin market cap timeseries

Econometricians learn to always use returns, log returns or differencing when studying timeseries of asset prices to get rid of non stationarity in the time series. Non-stationary time series in financial models are likely to produce unreliable and spurious results and as a consequence lead to poor understanding and forecasting.

Checking the assumptions: Heteroskedasticity

Aside from checking for autocorrelation, we also have to check for heteroskedasticity. A plot of the squared residuals could show us if we might expect heteroskedasticity.

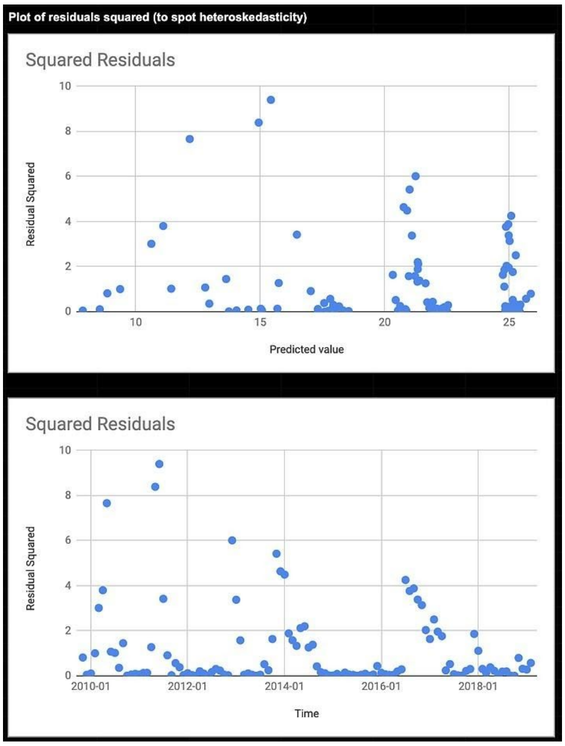

Squared residuals plotted vs predicted value and vs time

At a first glance it seems like the error doesn't have the same variance everywhere, which would imply heteroskedasticity. In order to confirm the suspicions, I ran both the Whites test and the Breusch-Pagan test. In both cases we have to reject the null-hypothesis of heteroskedasticity being present in the model. In laymen terms; we are not able to state that the model suffers from heteroskedasticity, which is good! Test results for Whites and Breusch-Pagan are below.

Variables			Intercept	TRUE		
dependent	Error^2					
independent	ln(S2F)					
Regression Statistics						
Multiple R	0.1073649039					
R^2	0.0115272226					
Standard Error	1.794077192					
Adjusted R^2	0.002701572799					
Observations	114					
ANOVA						
	df	*SS*	*MS*	*F*	*Significance of F*	
Regression	1	4.20397611	4.20397611	1.30610469	0.2555382438	
Residual	112	360.4958529	3.218712973			
Total	113	364.699829				
	Coefficients	*Standard Error*	*t-Statistics*	*p-Value*	*Lower 95%*	*Upper 95%*
Intercept	1.35673565	0.2806702915	4.833912568	0.000004	0.800623445	1.912847855
ln(S2F)	-0.145756065	0.1275374239	-1.142849373	0.255538	-0.3984551221	0.1069429921

Breusch-Pagan showing no sign of Heteroskedasticity

Variables			Intercept	TRUE		
dependent	Error^2					
independent	Pred value	Pred Value^2				
Regression Statistics						
Multiple R	0.1080997987					
R^2	0.01168556647					
Standard Error	1.801996153					
Adjusted R^2	-0.00612190080					
Observations	114					
ANOVA						
	df	*SS*	*MS*	*F*	*Significance of F*	
Regression	2	4.261724094	2.130862047	0.6562172089	0.5208110081	
Residual	111	360.4381049	3.247190135			
Total	113	364.699829				
	Coefficients	*Standard Error*	*t-Statistics*	*p-Value*	*Lower 95%*	*Upper 95%*
Intercept	1.664259288	2.242968804	0.7419894941	0.459661	-2.780333131	6.108851707
Pred value	-0.00785588782	0.2553985685	-0.030759326	0.975517	-0.5139452009	0.4982334252
Pred Value^2	-0.00092927785	0.006968367305	-0.1333566115	0.894153	-0.01473756285	0.01287900715

Whites Test showing no sign of Heteroskedasticity

Minor data issue

I have downloaded the data used by Plan B from his github repository to replicate and review his work. I noticed there was a little error in the number of bitcoins created over time. In his data set he did not change mining rewards every 210.000 blocks (as follows from the protocol) but instead he shifted to a new mining reward regime once the monthly measured block height passed another 210.000 blocks. I expect the resulting error to be not really significant, but decided to use a closed form formula that exactly returns the number of bitcoins into existence as a function of block height. The function follows from solving the initial value problem for the differential equation that describes the growth of the bitcoin supply. This is perfectly explained by Onur Solmaz [3] in his article on the bitcoin inflation curvature. In this function the following variables are defined:

- S: the total supply
- R0: the initial mining reward = 50
- alpha: the reward decay factor = 0.5
- h: blockheight
- beta: the milestone number of blocks before a decrease of the reward kicks in is 210.000

Please note that the floor function is used where h is divided by beta.

$$\begin{aligned} S(h) &= \sum_{i=0}^{\lfloor h/\beta \rfloor -1} \alpha^i \beta R_0 + \alpha^{\lfloor h/\beta \rfloor} (h \bmod \beta) R_0 \\ &= R_0 \left(\beta \frac{1-\alpha^{\lfloor h/\beta \rfloor}}{1-\alpha} + \alpha^{\lfloor h/\beta \rfloor} (h \bmod \beta) \right) \end{aligned}$$

Closed formula to determine the bitcoin supply as a function of block height

The graph below shows how the number of bitcoins over time in PlanB's dataset differ from the true supply.

Difference in bitcoin supply over time in Plan B's dataset vs actual supply

Conclusion

With the model suffering from autocorrelation in the residuals, we have to reject the current model and come up with a better model. For most quants it is nothing new that explanatory models w.r.t. asset performance are always using returns or log returns to avoid autocorrelation ruining the model. In a follow up article I will check different approaches to help us find a model that meets the assumptions. Unfortunately I can't guarantee there will be a model that will meet the requirements and will be significant at the same time, but the quant in me is thirsty for improvements to the work that got so much attention from the community.

June 30th 2020

Part II: The hunt for cointegration

Introduction

In my first review of the work of PlanB, I concluded that the relation between stock-to-flow and bitcoin price as pointed out by the author was invalid because the general assumptions of ordinary least squares regression were not met. When two variables are non-stationary and we estimate a regression model, there is a good chance we find highly autocorrelated residuals and a significant value for the coefficient. This phenomenon is well known as spurious regression. But, spurious regression isn't always the case. Sometimes the variables might be cointegrated, which would imply that the estimated relation is superconsistent. Another review by Nick pointed out that in this specific case we could be dealing with the exceptional case of cointegration. For a better understanding of cointegration, I would recommend to have a look at a very good visual introduction of the concept here.

In this article, I will investigate if the log of bitcoin price and the log of its stock-to-flow ratio are indeed cointegrated. If cointegration applies, it turns out that the OLS estimates of the coefficients are consistent. If this is the case, I would have to reject my earlier conclusion where I said that the relation between the two variables as indicated by PlanB is nonsense since the OLS assumptions are not met.

As the concept of cointegration wasn't really on top of my mind, I had to take a deep dive in some of my college books and academical literature during my holiday to refresh my mind on the concept and how to test for it.

TL;DR in layman terms

In an earlier analysis I showed that assumptions that should be met, were not met and that the resulting model therefor was flawed. In this article I looked into an exceptional case. If I would be able to confirm we are dealing with that specific exception, the resulting model would be validated and could be used to quantify the relation between stock-to-flow and bitcoin price. It turned out that the exception indeed applies and that we CAN use the model.

Difference between correlation and cointegration

Before we continue, it's good to understand the difference between cointegration and correlation. Correlation is describing the in tandem movement of two (or more) variables. Cointegration is about the constant difference (with a stationary distribution) between the means of the same variables. Or a bit shorter: cointegration means that two time series both share a stochastic drift.

June 30th 2020

Method

All analysis is performed in Python where I used the following packages:

- numpy
- pandas
- statsmodels
- matplotlib

The dataset originates from my earlier analysis and a download can be found here. I figured out how to use Jupyter Notebook to visualise the analysis, because learning how to work with Jupyter was still on my wish list. The best way to learn these things is by just having a go at it.

Testing

I use three different approaches to test for cointegration of the natural logarithms of bitcoins price and stock-to-flow ratio. To easily refer to those series we refer to them as lnBTCprice and lnS2F. I use the following tests:

- Cointegrating Regression Durbin-Watson test (CRDW test);
- the two step Engle Granger test;
- the Johansen test.

All approaches are briefly summarised below.

CRDW Test

Test whether the Durbin-Watson statistic is significantly larger than 0. If a unit root exists the value should be close to zero. If we can't reject the presence of a unit root in the residuals, this implies we can't reject that the variables are not cointegrated.

Engle Granger Test

1. Determine the integration order of the two time series; lnS2F and lnBTCprice. (i.e. how often do we need to difference the series in order to find a stationary time series).
2. If both lnS2F_ and lnBTCprice_t are integrated of order one (abbreviated to I(1)), we know that if these two series cointegrate then there will exist coefficients, and such that: lnBTCprice_ = + lnS2F_ +u_ . The residuals that follow from running a regression will be tested for unit root, as residuals should be stationary in case variables are cointegrated.
3. If for the residuals we can reject the null hypothesis of the presence of a unit root, we can say with at least 99% certainty that the residuals are not integrated of the first order.

Johansen Test

We know the natural logarithms of bitcoin price and S2F are both non stationary, which means they are integrated of an order larger than 0. That implies we can model both series by means of an autoregressive model. As we model both series at once, we can use a vector auto regressive (VAR) model in which y is the *nx1* vector of variables integrated of order one (lnBTCprice and lnS2F).

$$\mathbf{y}_t = \boldsymbol{\mu} + \mathbf{A}_1 \mathbf{y}_{t-1} + \cdots + \mathbf{A}_p \mathbf{y}_{t-p} + \boldsymbol{\varepsilon}_t$$

This can be rewritten as:

$$\Delta \mathbf{y}_t = \boldsymbol{\mu} + \boldsymbol{\Pi} \mathbf{y}_{t-1} + \sum_{i=1}^{p-1} \boldsymbol{\Gamma}_i \Delta \mathbf{y}_{t-i} + \boldsymbol{\varepsilon}_t$$

where:

$$\boldsymbol{\Pi} = \sum_{i=1}^{p} \mathbf{A}_i - \mathbf{I} \text{ and } \boldsymbol{\Gamma}_i = -\sum_{j=i+1}^{p} \mathbf{A}_j$$

In the second equation above we have multidimensional variables and multiplication would happen via matrix multiplication. The Johansen tests consists of two tests: the maximum eigenvalue test, and the trace test. For both test statistics we test the null hypothesis of no cointegration against the alternative of cointegration, by means of comparing the test statistics to the critical values for the test.

Running the tests

In this section we carry out the mentioned tests and have a closer look at the results of these tests.

Engle Granger and CRDW

Both series (natural logarithms of S2F and bitcoins price) are clearly not stationary, but trending over time. After differencing the series, we might find stationarity for both though.

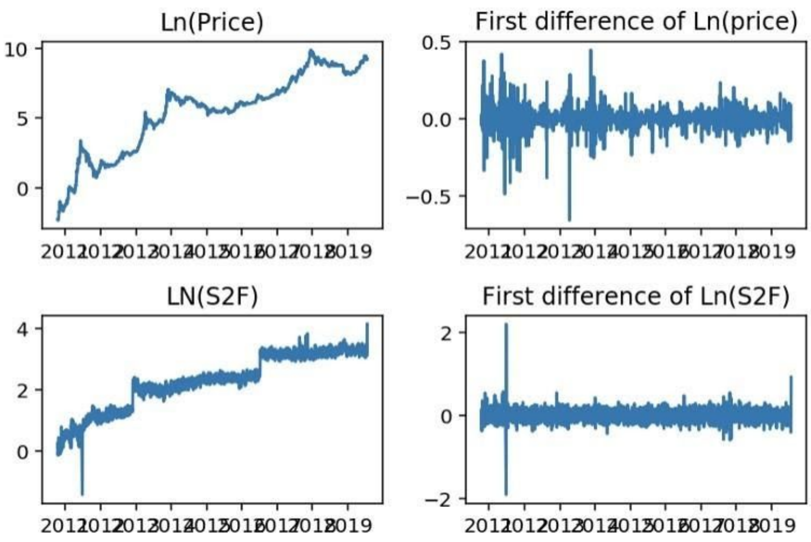

By the naked eye I would say there is a very good chance that the differenced series are both stationary, but we need to check that as well. To verify whether the differenced series are stationary I ran the augmented Dickey-Fuller (ADF) test for both differenced series. Code for the test is in the appendix.

ADF Statistic: -12.843153 p-value: 0.000000

Critical Values:
1%: -3.432
5%: -2.862
10%: -2.567

ADF test result for first order difference of ln(price)

ADF Statistic: -15.426991 p-value: 0.000000

Critical Values:
1%: -3.432
5%: -2.862
10%: -2.567

ADF test result for first order difference of ln(S2F)

For both series we can reject the null hypothesis of the presence of a unit root, which tells us we can say with at least 99% certainty that both variables are not integrated of the first order. Time to run an OLS regression to estimate the coefficients in:

$$lnBTCprice\_t = \alpha + lnS2F\_t + e\_t$$

Here's the regression summary which we use for both CRDW and the Engle and Granger procedure.

OLS Regression Results

Dep. Variable:	lnprice	**R-squared:**	0.875
Model:	OLS	**Adj. R-squared:**	0.874
Method:	Least Squares	**F-statistic:**	2.229e+04
Date:	Thu, 05 Sep 2019	**Prob (F-statistic):**	0.00
Time:	13:39:03	**Log-Likelihood:**	-4543.8
No. Observations:	3200	**AIC:**	9092.
Df Residuals:	3198	**BIC:**	9104.
Df Model:	1		
Covariance Type:	nonrobust		

	coef	std err	t	P>|t|	[0.025	0.975]
const	-1.2696	0.048	-26.463	0.000	-1.364	-1.176
LNS2F	2.9650	0.020	149.285	0.000	2.926	3.004

Omnibus:	116.008	**Durbin-Watson:**	0.161
Prob(Omnibus):	0.000	**Jarque-Bera (JB):**	309.481
Skew:	0.120	**Prob(JB):**	6.27e-68
Kurtosis:	4.505	**Cond. No.**	7.54

OLS Regression Results from statsmodels package in Python

We'll have a closer look at the residuals from that regression. The residuals as shown below don't look like a stationary series, but the Durbin Watson statistic is just significantly larger than zero for~3200 observations, so even though the residual plot indicates no cointegration, the CRDW test statistic (value=0.161) doesn't support this!

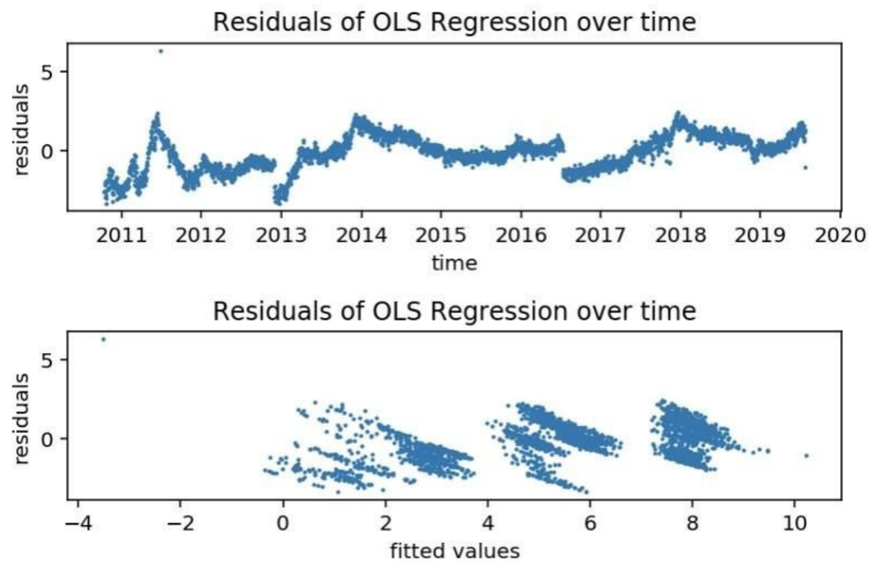

Residuals vs fitted values and time

I ran the ADF test to check for unit root in the residuals. According to the ADF test we have to reject the null hypothesis and conclude that the residuals are stationary. The concept of cointegration is again not rejected!

ADF Statistic: -3.714701 p-value: 0.003911
Critical Values:
1%: -3.432
5%: -2.862
10%: -2.567

ADF test result for regression residuals

Johansen test

As mentioned the Johansen test consists of two separate tests; themaximum eigenvalue test and the trace test. The statsmodels package in Python was used to conduct the tests. Code can be found in the Appendix.

Trace Statistic:
[77.61330689 8.83704667]
Critical Values Trace Statistic [90% 95% 99%]:
[[13.4294 15.4943 19.9349]
[2.7055 3.8415 6.6349]]
Maximum Eigenvalue Statistic
[68.77626022 8.83704667]
Critical Values Maximum Eigenvalue Statistic [90% 95% 99%]
[[12.2971 14.2639 18.52]
[2.7055 3.8415 6.6349]]

For both Johansen tests we fail to reject the null hypothesis (as the test statistics are higher than critical values for all confidence intervals).

Conclusion

The estimated relation between lnBTCprice and lnS2F is consistent (even though the OLS assumptions are not met) as we have shown that the time series are cointegrated. My former conclusion is thereby falsified. As cointegration applies we are able to use the coefficients coming from the OLS to quantify a model that describes the relation between the two series. We could set up a Vector Error Correction Model to model both the short term and the long term dynamics of the relation, which I leave for a follow up article.

Part III: The Fall Of Cointegration

Introduction

We are a little more than a year further down the road since PlanB wrote his initial piece on the S2F model. Both phraudsta and I have put in quite some work to run model validations and quite recently our findings were criticised by Sebastian Kripfganz (Assistant Professor in Econometrics at the University of Exeter), who pointed out that the deterministic elements in the S2F timeseries should be accounted for.

Now that the halving is behind us, the majority is probably less interested in these model validations. Most likely I'm writing this piece as a rectification of my earlier work and to bring some clarity to myself. Clarity w.r.t. the question whether or not it would be possible that the two main variables within PlanB's model are cointegrated or not. Why did cointegration matter? If the two variables are cointegrated, the key take away is that there is a long term relation between the two non stationary timeseries.

This piece is actually the missing piece in my series on the stock to flow model. (If you paid close attention, there never was a part III, while I did publish a part IV). In part IV I already wrote what the theoretical framework should look like and how it should lead us in terms of model selection. This piece will mainly readdress the concepts of cointegration and integration order and what they mean for the model so many people fell in love with. I will evaluate the definitions of these important statistical concepts applied in earlier research and check how they hold up. Something that I should have done before I started to run calculations.

Definitions of cointegration and integration order

As per the paper on cointegration by Engle and Granger [1], the initial definition of cointegration is as stated below:

> DEFINITION: The components of the vector x_t are said to be *co-integrated of order d, b*, denoted $x_t \sim CI(d, b)$, if (i) all components of x_t are $I(d)$; (ii) there exists a vector $\alpha(\neq 0)$ so that $z_t = \alpha' x_t \sim I(d-b)$, $b > 0$. The vector α is called the *co-integrating vector.*

Definition of cointegration by Engle and Granger

I'll try to describe their definition in my own words and tweak it a bit so that it better fits with our quest. A number of timeseries are said to be cointegrated if:

1. all these timeseries are integrated of the same order *d*;
2. there exists a linear combination *z* of these timeseries that results in a timeseries which is integrated of at least one order smaller.

This definition leads us to the following question. How to define integration order? As per the same paper a timeseries integrated of order *d* is defined as:

DEFINITION: **A series with no deterministic component which has a stationary, invertible, ARMA representation after differencing d times, is said to be integrated of order d, denoted $x_t \sim I(d)$.**

Definition of integration order as per Engle and Granger

This definition tells us that:

- the original timeseries can't have any deterministic component, so before we check order of integration we should get rid of any such component;
- the timeseries should have an ARMA representation after differencing *d* times;
- the resulting ARMA representation should be stationary and invertible.

The definition of integration order was introduced much earlier in 1938 by Wold and is better known as Wold's decomposition theorem.

Wold's Theorem: *A stationary time series process, after removal of any deterministic components, has an infinite moving average (MA) representation which, in turn, can be represented by a finite autoregressive moving average (ARMA) process.*

Back to the model

We have two timeseries of interest; the natural logarithm of stock-to-flow and bitcoin market cap and the question is how these timeseries should be modeled. Before we start running any tests to check for integration order and stationarity of the timeseries, we first look again how the timeseries is constructed.

"*It is sometimes very difficult to decide whether trend is best modeled as deterministic or stochastic, and the decision is an important part of the science — and art — of building forecasting models*."
— Diebold, Elements of Forecasting, 1998

Sometimes it's hard to see whether a timeseries is stochastic or deterministic. Let's check stock to flow once again. Let's start here with a visual inspection of the series. The series is shown below.

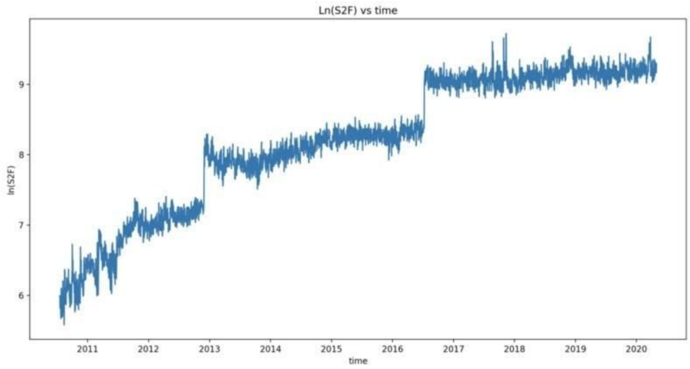

Chart 1: Stock to Flow timeseries

Taken as a whole, the timeseries seems to have both deterministic and stochastic elements. We can see there are three epochs and that each epoch seems to have its own trend. Next to that we also see that there is a clear jump (aka the halving or halvening) between the different epochs. These jumps are referred to as structural breaks in literature. So, we seem to have both a change in trend and intercept at the breaks.

The way I like to think about the series is as follows. Given that the blockreward is known for a certain blockheight, the S2F metric is fully deterministic vs blockheight. To show this please note how the flow evolves as a function of blockheight.

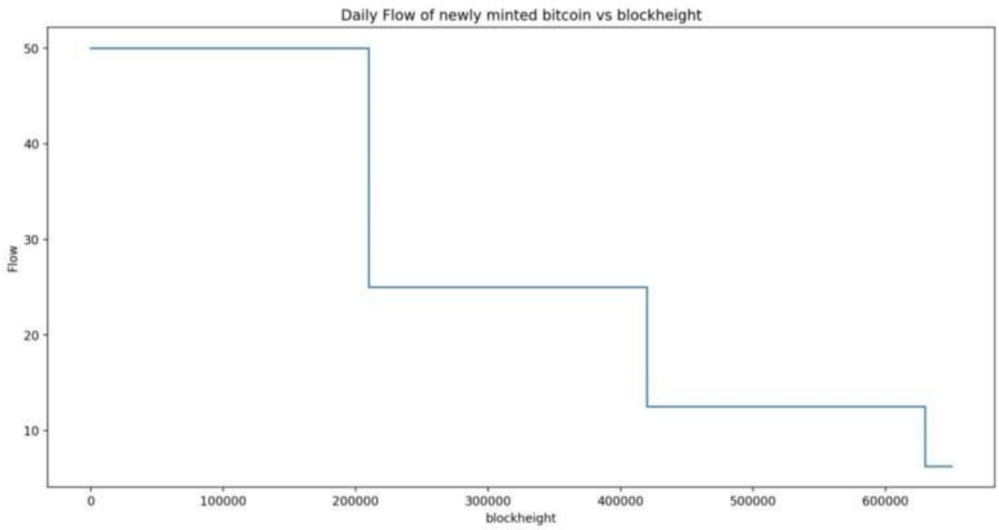

Chart 2: Flow of bitcoin vs bitcoin blockheight

Now look at both supply and stock to flow vs blockheight.

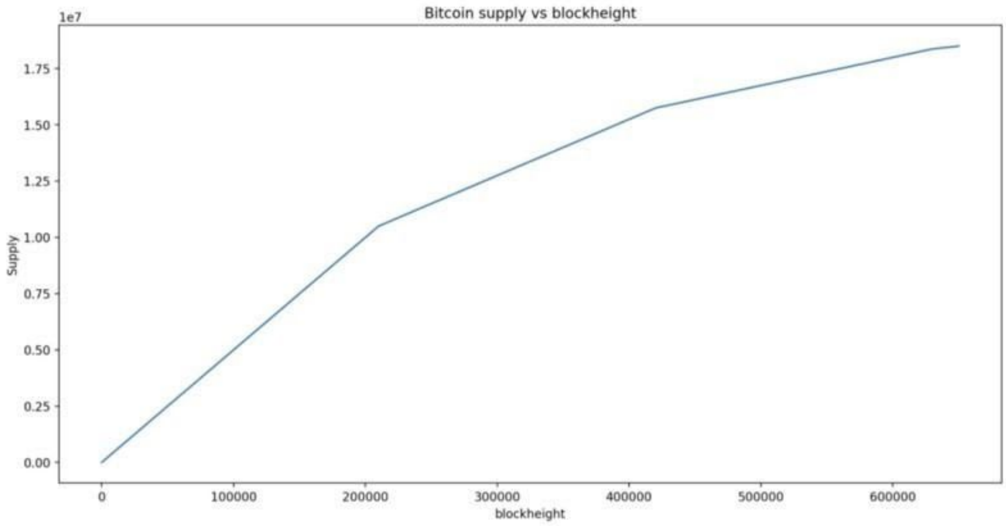

Chart 3: Bitcoin supply vs bitcoin blockheight

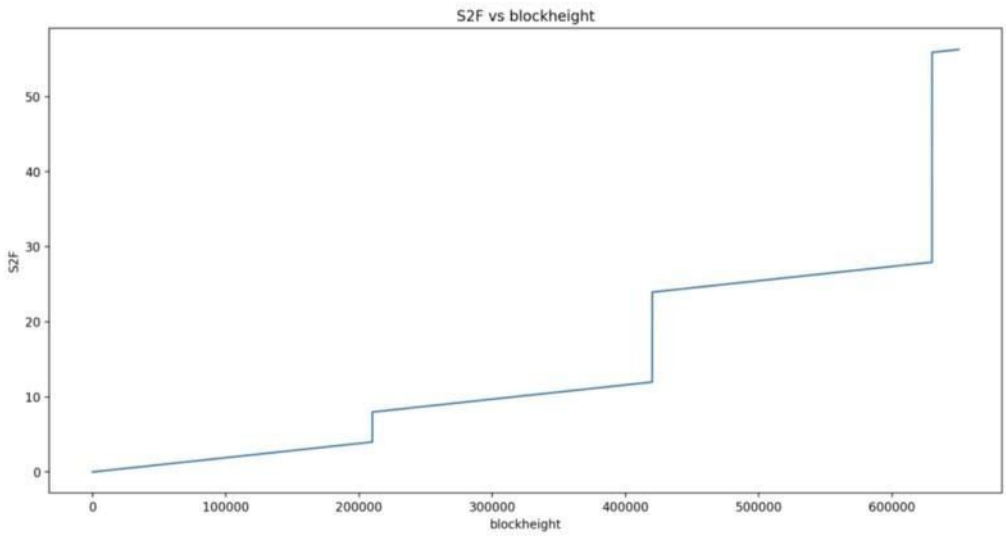

Chart 4: Bitcoin S2F vs bitcoin blockheight

All these evolutions over blockheight are fully deterministic. The stochasticity we observe in the timeseries of S2F vs time comes from the uncertainty in the interarrival time for each block. I took a sample of 10.000 inter arrival times and had a look into the distribution.

	sumsquare_error	aic	bic	kl_div
exponweib	0.000072	1414.812807	-187442.460173	inf
gengamma	0.000072	1415.657417	-187427.840699	inf
betaprime	0.000078	1411.306255	-186691.059259	inf
gamma	0.000108	1442.313012	-183428.021006	inf
chi2	0.000108	1442.306581	-183427.571284	inf
pearson3	0.000108	1442.313220	-183427.514062	inf
beta	0.000126	1458.868053	-181877.326776	inf
weibull_min	0.000127	1440.741984	-181773.159405	inf
frechet_r	0.000127	1440.741984	-181773.159405	inf
halfgennorm	0.000150	1433.334318	-180093.390879	inf

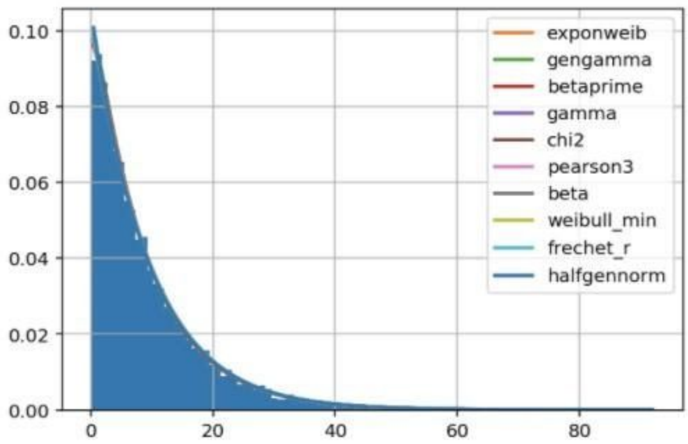

Distribution Fitting of blocktimes for bitcoin

The deviation in the interarrival times causes the S2F metric over time to show some stochasticity. If miners would mine new blocks exactly every 10 minutes, then S2F over time would actually be fully deterministic. But this is how S2F looks vs time.

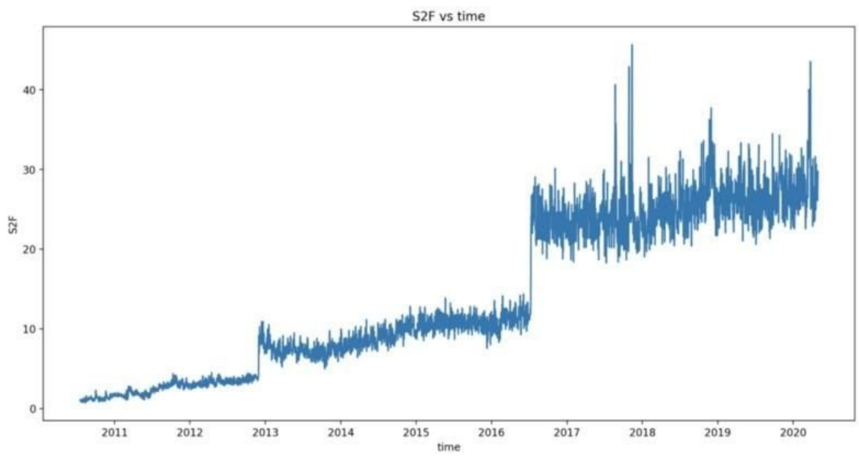

Chart 5: Stock-to-flow over time

Let's have a closer look at the distribution of daily flow per reward era.

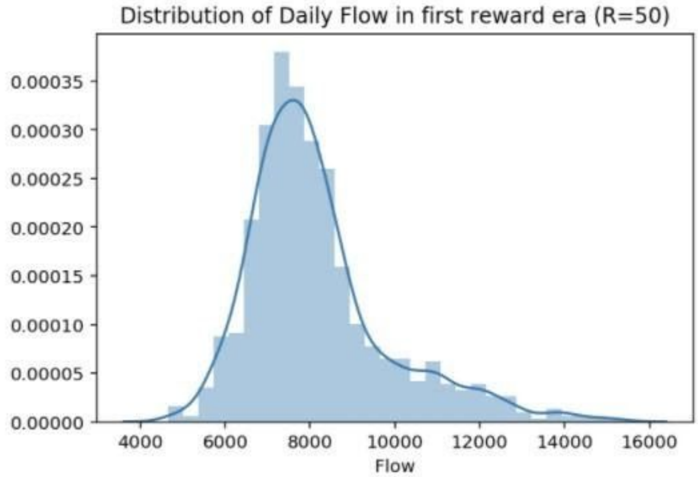

Chart 6: Distribution of daily flow in first reward era

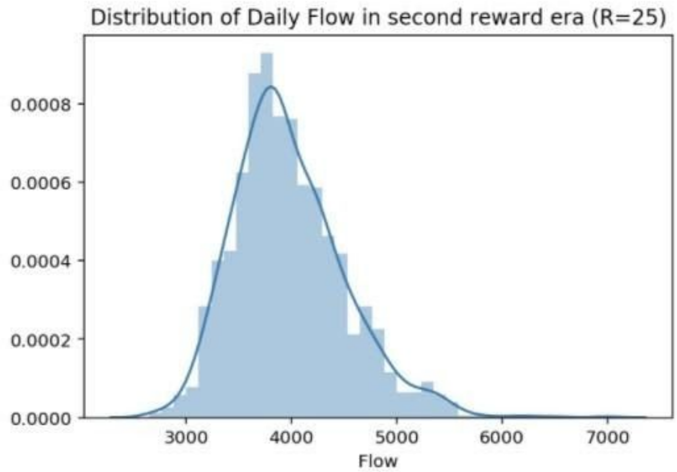

Chart 7: Distribution of daily flow in second reward era

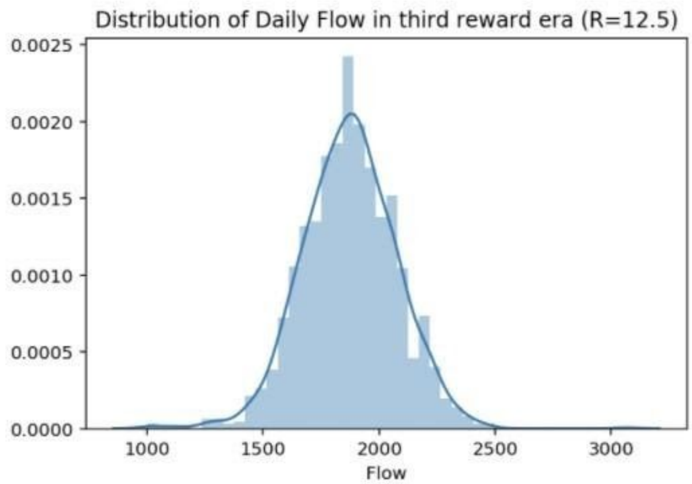

Chart 8: Distribution of daily flow in third reward era

It's exactly this variance in Flow that you'll see back in the day to day change of the S2F metric as well.

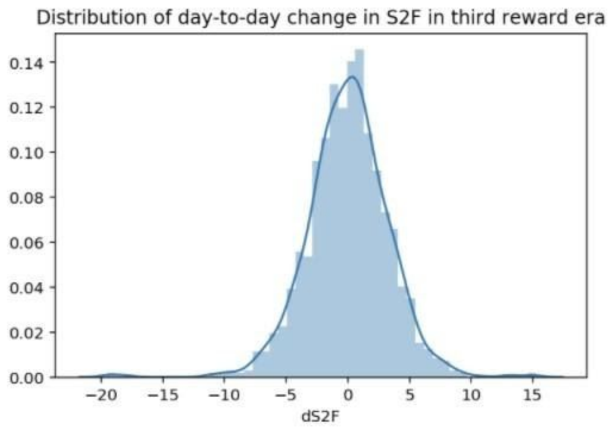

Chart 9: Distribution of day-to-day change in S2F in third reward era

So now we know exactly where the variance is coming from and how it looks like. Given all the analysis above, I think there is enough reason to assume that S2F has 2 deterministic elements;

- a deterministic intercept for every reward era and
- a deterministic trend for every reward era

This would leave us with a model for S2F that looks like:

$$S2F_{r,t} = \alpha_r + \beta_r t + \epsilon_{r,t}$$

Eqn 1: stochastic model with deterministic elements to describe S2F over time

In the above model alpha is the reward era dependent intercept and beta represents the reward era dependent trend. Epsilon represents the noise resulting from variance in blocktimes.
After correcting for the deterministic components, we are only left with a most likely stationary process; epsilon.

What does this mean for cointegration?

In order for two variables to be cointegrated, we require two time series with an equal order of integration. In the previous analysis, I showed bitcoins price is first order integrated, and that still stands. Bitcoin doesn't show any deterministic elements in the evolution of price over time.
For the S2F timeseries I did not correct for the deterministic elements in the series in my earlier work, which according to Engle and Granger should have been corrected before determining the order of integration.
If we strictly follow Engle and Grangers theory and the Wold decomposition theorem, we see that after correction for deterministic elements in the timeseries, we are left with a most likely stationary process.
After correction for the deterministic elements, S2F turns into a I(0) timeseries, while price still is a I(1) timeseries. In this case cointegration as defined by Engle and Granger is impossible, as per definition both processes should have the same order of integration.

Conclusion

The concept of cointegration once saved the by now famous stock-to-flow model from a potential spurious relation between S2F and bitcoin price and is now shown to be improperly applied in among others my earlier work. This means that my earlier work "Reviewing "Modelling Bitcoin's Value with Scarcity" — Part II: The hunt for cointegration" is invalidated; there is no such thing as cointegration between ln(S2F) and ln(price).

So, now that cointegration is off the table for the model as is, we are no longer able to prove the relation between S2F and bitcoin price isn't spurious. This doesn't mean that the current model and all the metrics around it are useless, but I'd consider them as additional equipment in your Technical Analysis toolkit. It's up to you whether you deem that useful or not.

A final note

I'm not concluding there is no relation between bitcoins price and scarcity. I still think there is a connection between the scarcity of bitcoin and price. I am concluding that we can no longer prove that the S2F model as laid out by PlanB is the model that is able to describe this relation while respecting the definitions and assumptions of the applied statistical concepts.

June 30th 2020

Part IV[1]: The Theoretical Framework leading to the Error Correction Model

Introduction

In my first review of the work of PlanB, I concluded that the relation between stock-to-flow and bitcoin price as pointed out by the author was invalid because the general assumptions of ordinary least squares regression were not met. When two variables are non-stationary and we estimate a regression model, there is a good chance we find highly autocorrelated residuals and a significant value for the coefficient. This phenomenon is well known as spurious regression. But, spurious regression isn't always the case. Sometimes the variables might be cointegrated, which would imply that the estimated relation is super consistent. Nick pointed out that we could very well be dealing with the exceptional case of cointegration and showed that he wasn't able to falsify the cointegrating relationship between stock-to-flow and bitcoin's market cap. After Nick showed that the variables were cointegrated, I verified his findings. Since I still was skeptic, I chose to run the analysis on my own dataset and ran three different cointegration tests to make sure there was no doubt. Even though I expected I would be able to show that at least one of those tests would lead me to reject cointegration, I could not. Initially, I was a bit too fast with drawing my conclusions and as a result I warned people that the model was flawed. So, I offered my apologies in public for drawing a wrong conclusion and engaged in many discussions to explain why the model was eventually right. Because I noticed in those discussions that the basis underneath the material we discuss is poorly understood by most people, I decided to start writing a book that will help people to better understand the econometric concepts we're dealing with. To stay in the loop about that development, I recommend subscribing at www.bitcoinometrics.io to get notified once the book is available. But that's not what this piece is about. This piece is about further development of the model and presenting a framework which helps to understand the developments. In the process of writing my book I also conduct some academical research. Not only to refresh my mind on time series analysis, but also to check with academical researchers if there were any important developments in that field of research. Nick in his write up already mentioned the Vector Error Correction Model ('VECM') and estimated the coefficients for the model as part of his attempt to falsify cointegration. In my 'hunt for cointegration' I also touched on it without estimating any of the model coefficients, but I ended the article by stating that setting up a VECM would be a nice subject for a follow up article. This is still work in progress, but I like to share a bit more on how we actually get to that point and how to go about.

Model Selection Framework

Usually when one is looking to quantify the relation between non-stationary time series, the first step is to difference the series until a stationary series is found. This is basically the first thing you learn as an Econometrics student when you follow classes on Regression Analysis or Time Series

[1] *Please note that this publication was issued before part III was issued.

Analysis. But differencing the time series to make them stationary is only one possible direction to come to a solution. And it comes at the cost of throwing away data that might identify long run relationships between the time series.

Another possibly better solution is to test whether the time series are cointegrated. If the cointegration test tells us cointegration exists, we can set up a model that is able to describe both the long run relationship and the short term corrections.

One of the things I noticed in the literature is that it's hard to find a basic method selection decision tree. I like to attack these kind of problems as structured as possible, so that the chance of actually finding meaningful relations increases while you also prevent yourself from misspecifying a model. In my search for a helpful framework, I found this useful article.

The framework as shown below is my slightly adjusted version that is based on the one I found. It will serve as a guide in the steps we will take. All the shapes that contain bold text, show the path we follow in case we like to construct a model to quantify the relation between stock-to-flow and price (or market cap). In the earlier articles (here and here) Nick and I both independently showed that both variables are first order integrated (after applying differencing we end up with a stationary series over time) and that cointegration couldn't be rejected by running different tests.

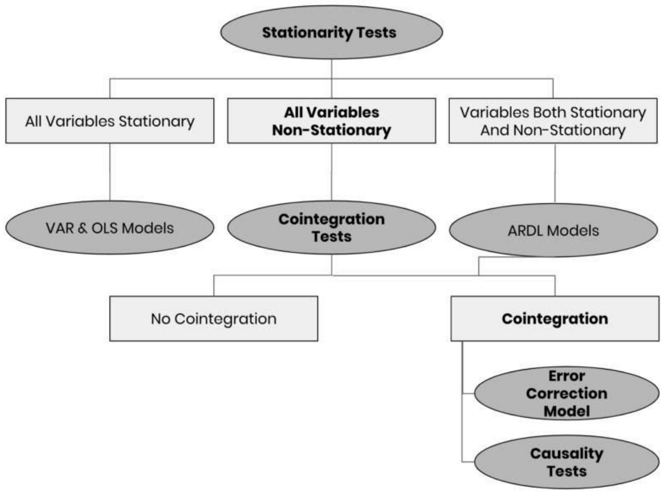

Fig 1: Simplified Model Selection Framework for time series analysis

Note that even though the overview above is far from complete, it offers a useful overview for those models that are often used. Most practitioners who use linear regression models, just go straight to the OLS models without even testing for stationarity. Using frameworks like these would be beneficial to many of them. It also helps people with less of a statistical background to check whether the beautiful model they have been introduced to was indeed the appropriate model to use.

Please also note that this overview is only about model selection and not about estimation of the model. I like to emphasise again that this overview is very simplified.

We found cointegration. Now what?

As cointegration couldn't be falsified, this means that the two variables are linked to form an equilibrium relationship spanning the long run. One of the issues though with the different cointegration tests is that some of them have weaknesses. Johansen (1988) addressed those and came up with an improved cointegration test model, which is widely applied nowadays and incorporated in many different econometric software packages. Both Nick and I used that test in our model validation and we concluded that the model as introduced by PlanB couldn't be falsified.

If both the variables are first order integrated and there exists a cointegration relationship then we can derive an Error Correction Model. When both variables are put into a vector this can also more generally be referred to as a Vector Error Correction Model ('VECM').

Following the framework

First, I set up all the different equations, and briefly touch upon them. Keep in mind that the long term relation is what matters most in terms of showing us the road, and realise that the short term corrections are modelled to return to the middle of the road. In order to run a full blown analysis we need to define:

- the linear model
- the cointegration representation
- the error correction models for both S2F and BMC
- the models to run causality analysis

Setting up the model equations

If we consider the logarithm of bitcoin market cap ('Log(BMC)') as *Y* and we consider the logarithm of stock-to-flow ('LogS2F') as *X,* then the relationship between the two is written as:

$$Y_t = \alpha + \beta X_t + \epsilon_t$$

Equation 1: The Linear Model

Based on the representation theorem of Engle and Granger (1987), we rewrite the equation to display the cointegrating relationship:

$$\epsilon_t = Y_t - \alpha - \beta X_t$$

Equation 2: Cointegration Representation

Since both variables have their own Error Correction Models, I introduce them here:

$$\Delta Y_t = \alpha_Y + \rho_Y \epsilon_{t-1} + \sum_{h=1}^{l} a_{1h} \Delta Y_{t-h} + \sum_{h=1}^{l} b_{1h} \Delta X_{t-h} + u_{Y_t}$$

Equation 3: Error Correction Model for bitcoin market cap

$$\Delta X_t = \alpha_X + \rho_X \epsilon_{t-1} + \sum_{h=1}^{l} a_{2h} \Delta Y_{t-h} + \sum_{h=1}^{l} b_{2h} \Delta X_{t-h} + u_{X_t}$$

Equation 4: Error Correction Model for Stock-To-Flow

The first expressions at the right hand side of the Error Correction Model equations indicated by α are both stationary white noise processes for some number of lags l. In case we look at only at one period lag, the equations become:

$$\Delta Y_t = \alpha_Y + \rho_Y \epsilon_{t-1} + a_1 \Delta Y_{t-1} + b_1 \Delta X_{t-1} + u_{Y_t}$$

Equation 5: Error Correction Model for BMC with one lag.

$$\Delta X_t = \alpha_X + \rho_X \epsilon_{t-1} + a_2 \Delta Y_{t-1} + b_2 \Delta X_{t-1} + u_{X_t}$$

Equation 6: Error Correction Model for S2F with one lag

The coefficients in the cointegration equation are used to show the estimated long term relation among S2F and BMC. The coefficients in the Error Correction Models will provide more information on how deviations from that long term relation affect the changes on them in the next period. In this case measures the speed of adjustment to the long term equilibrium. So, when the model runs away from the long term theoretical price, these coefficients tell us how fast we will return to it.

It's important to note that all equations above on their own, are just equations. I preferred to keep it readable, by not addressing the assumptions over and over again. In order to turn these equations into models we should say something about what assumptions are made regarding the errors. In case we like to use parametric estimation we have to either use 'the weak' or 'the strong' assumptions regarding the error. Under the weak assumptions we only make assumptions regarding the first 2 moments of the distribution, while under the strong assumption we define the entire distribution itself.

Causality Testing and required models

Since S2F and BMC are cointegrated, one of the following statements regarding this relationship will hold:

- S2F drives BMC,
- BMC drives S2F, or
- BMC and S2F drive each other

If S2F and BMC were not cointegrated, there would be no causal relation and the two variables would be independent. Granger (1969) has developed a causality test method that will enable us to determine the direction of the relationship. If current and lagged values of S2F improve the prediction of the future value of BMC, then it is said that S2F Granger causes BMC. And the opposite can be said as well.

The model equations that we build to test for the direction of the relation are shown below:

$$\Delta Y_t = \sum_{i=1}^{n} \alpha_i \Delta Y_{t-i} + \sum_{j=1}^{n} \beta_j \Delta X_{t-j} + u_{1t}$$

Equation 7: Model to test if X Granger causes Y

$$\Delta X_t = \sum_{i=1}^{n} \lambda_i \Delta X_{t-i} + \sum_{j=1}^{n} \delta_j \Delta Y_{t-j} + u_{2t}$$

Equation 8: Model to test if Y Granger causes X

In both cases we are testing the null hypothesis that beta and delta are equal to zero. If beta in equation 7 is equal to zero then X is not Granger causing Y (or LogS2F is not Granger causing LogBMC). And in equation 8 we test for delta being equal to zero.

Estimating the model coefficients

In order to estimate model coefficients we can either choose between parametric or non parametric approaches or some kind of combination of both. As stated by Green:

"Contemporary econometrics offers the practitioner a remarkable variety of estimation methods, ranging from tightly parameterized likelihood based techniques at one end to thinly stated nonparametric methods that assume little more than mere association between variables at the other, and a rich variety in between. ***Even the experienced researcher could be forgiven for wondering how they should choose from this long menu.*** *As a general proposition, the progression from full to semi- to non-parametric estimation relaxes strong assumptions, but at the cost of weakening the conclusions that can be drawn from the data."*

The main reason to use a parametric estimation is that we can infer stronger conclusions (as long as assumptions are respected), because we use strong assumptions w.r.t. the distributions of the variables we analyse. At the other hand, the advantage of non parametric estimation is that we don't need to impose strong assumptions on the distributions of the variables we analyse, but we can just go with their actual

distribution. As long as the actual distribution comes close enough to the imposed distribution, I would prefer to go with parametric estimation (like OLS estimation), but when the imposed distribution is not really a good match with the actual distribution, non-parametric estimation is preferred.

What's next?

I intended to write a piece on the estimation of the VECM model coefficients, but actual estimation calculations are put on the ToDo list for now. This piece has been in draft mode for a while, so I decided to split up the content into a theoretical part and a more hands on estimation of the parameters. So, next in line is running some actual calculations on different estimation techniques and presentation of the results.

June 30th 2020

BurgerCrypto.com
Digital Asset Management & Consultancy

References Part I

1) https://medium.com/@100trillionUSD/modeling-bitcoins-value-with-scarcity-91fa0fc03e25
2) M. Verbeek, A Guide To Modern Econometrics
3) https://solmaz.io/2019/02/09/inflation-curvature-bitcoin-supply/
4) http://www.real-statistics.com/statistics-tables/durbin-watson-table/

References Part II

[1]:https://medium.com/@100trillionUSD/modeling-bitcoins-value-with

[2]: https://medium.com/burgercrypto-com/challenging-plan-b-a-reviewof-modelling-bitcoins-value with-scarcity-3d4e7e141286

[3]:https://medium.com/@btconometrics/falsifying-stock-to-flow-as-a-model-of-bitcoin-value-b2d9e61f68af

[4]: Co-Integration and Error Correction: Representation, Estimation, and Testing; Robert F. Engle and C. W. J. Granger, 1987

[5]: A guide to Modern Econometrics, second edition, 2005; M. Verbeek

[6]:https://medium.com/@vera.chernova.burfield/how-to-embed-jupyter-notebook-into-medium-posts-in-three-steps-1-2-3-54262edea88a

Appendix Part II

Python code

https://gist.github.com/MarcelBurger/ed216b12e436bb4f07497cecff2b6742

References Part III

[1]: Co-Integration and Error Correction: Representation, Estimation, and Testing; Robert F. Engle and C. W. J. Granger, 1987

[2]: A Study in the Analysis of Stationary Time Series; Wold, H., 1938

[3]: https://www.cairn.info/revue-cahiers-d-economie-politique-2015-1-page-91.htm#no2

[4]:https://medium.com/burgercrypto-com/reviewing-modelling-bitcoins-value-with-scarcity-part-ii-the-hunt-for-cointegration-66a8dcedd7ef

June 30th 2020

[5]:https://medium.com/@100trillionUSD/modeling-bitcoins-value-with-scarcity-91fa0fc03e25

[6]:https://medium.com/burgercrypto-com/challenging-plan-b-a-review-of-modelling-bitcoins-value-with-scarcity-3d4e7e141286

References Part IV:

[1]: A guide to Modern Econometrics, second edition, 2005; M. Verbeek

[2]: Time Series and Panel Data Econometrics, 2015; M. Hashem Pesaran

[3]: https://medium.com/@100trillionUSD/modeling-bitcoins-value-with-scarcity-91fa0fc03e25

[4]: https://medium.com/altcoin-magazine/falsifying-stock-to-flow-as-a-model-of-bitcoin-value-b2d9e61f68af

[5]: https://medium.com/burgercrypto-com/challenging-plan-b-a-review-of-modelling-bitcoins-value-with-scarcity-3d4e7e141286

[6]: https://medium.com/burgercrypto-com/reviewing-modelling-bitcoins-value-with-scarcity-part-ii-the-hunt-for-cointegration-66a8dcedd7ef

[7]: Co-Integration and Error Correction: Representation, Estimation, and Testing; Robert F. Engle and C. W. J. Granger, 1987

[8]:https://www.sciencedirect.com/science/article/pii/S2405918817300405

[9]: https://medium.com/@tylerneylon/how-to-write-mathematics-on-medium-f89aa45c42a0

[10]: Investigating causal relations by econometric models and cross-spectral methods; Econometrica, 37 (3) (1969), pp. 424–438; C.W.J. Granger

[11]: Econometric Analysis, 5th Edition, 2003, William H; Greene