Luxor

Forecasting Network Difficulty

For Bitcoin Miners, Hosters, Lenders and Hashrate Traders

By: Ben Harper, Jon Conley, Ethan Vera, Colin Harper and Matt Williams

Table of Contents

1. Introduction: Satoshi’s Control Valve

2. What is Network Difficulty?

3. How Has Network Difficulty Behaved Historically?

4. What Determines Network Difficulty?
    a. Block Times
    b. Hashrate
    c. Hashrate Supply
    d. Hashrate Demand
    e. Mining Luck

5. How Can We Forecast Network Difficulty?
    a. Qualitative Techniques
    b. Time Series Analysis and Projections
    c. Causal Models

6. Hashrate Index Premium Hashrate, Difficulty, and Hashprice Projection Updates

Bitcoin’s network difficulty is at an all time high. This has Bitcoin miners, hosting providers, lenders, financiers and hashrate forward traders all wondering what comes next. Will Bitcoin’s price keep up with hashrate and difficulty growth? Can investors with Bitcoin mining exposure use difficulty trends to plan for the future?

This report will answer these questions by covering:

- **What is network difficulty?**
- **How has network difficulty behaved historically?**
- **What determines network difficulty?**
- **How can we forecast network difficulty?**

We also introduce:

- A new time series forecasting method for upcoming difficulty adjustments which improves overall performance relative to the constant block time method, particularly at the beginning of the epoch.
- Updated hashrate supply and demand model projections and sensitivity tables, which will be made available to Hashrate Index Premium subscribers and refined on a quarterly basis. Please reach out to hashrateindex@luxor.tech if you are interested in subscribing to quarterly updates of Luxor’s hashprice, hashrate and difficulty models and forecasts.

What is Network Difficulty?

Bitcoin's network difficulty measures how much work a miner or mining pool must produce to submit a valid hash with the Bitcoin network's cryptographic hashing function (SHA-256). A higher network difficulty indicates that a miner must produce more work to submit a valid hash, meaning it takes more hashrate to produce a valid hash. If a miner or mining pool submits a valid hash, they mine the next block in the chain and have the right to collect the block reward, currently 6.25 bitcoin plus any fees from the transactions the miner includes in the block.

The more hashrate that Bitcoin miners add to the network, the faster miners will produce blocks, and this prompts the Bitcoin network to raise the difficulty of mining by adding extra zeros to the difficulty level; miners then need to produce a number under the target to produce a valid hash that entitles them to the next block in the chain's sequence.

To make the concept of network difficulty more concrete, we recommend checking out this SHA-256 Hash Generator. In the *Input Your Text Here* box, you can submit a hash (any string of characters) to the same cryptographic hashing function that the Bitcoin network uses. How long does it take you to find a valid hash with a theoretically difficulty of 1, defined as the *SHA256 hash box* showing a string that starts with a single zero? Now how about 20 zeros?

Spoiler: you'll need a specialized Bitcoin mining computer to find it. And the more zeros a number needs to start with, the more difficult it is to find a valid input. This is how Bitcoin network difficulty works. Miners can calculate their probability of submitting a valid hash and mining pools use this probability to calculate expected earnings and payouts.

Network difficulty is important for hashrate market participants because it significantly impacts hashprice – the revenue miners can expect to earn from a quantity of hashrate. As more hashrate comes onto the network, network difficulty rises and miners receive a smaller slice of block rewards – and vice versa.

The Bitcoin network programmatically adjusts this difficulty every 2,016 blocks (approximately every two weeks) to a level that targets an average block time of ten minutes. If network hashrate grows and miners find blocks more quickly than 10 minutes on average, the network difficulty will increase. If network hashrate falls and miners find blocks more slowly than 10 minutes on average, the network difficulty will decrease. To avoid extreme volatility, Bitcoin's code limits difficulty adjustments to a maximum of +300% and minimum of -75% (i.e., a factor of four).

Interestingly, due to an early bug in Bitcoin's code, the first block in each epoch is excluded from the difficulty adjustment calculation. This means that the difficulty adjustment is calculated from 2,015 blocks rather than the expected 2,016 blocks. Since the impact is rather minimal and would require a hard fork to fix, developers have decided to leave the bug be for now.

Anyone can view Bitcoin's network difficulty with Bitcoin mining data sites like Hashrate Index. Bitcoin users who operate their own nodes can also pull network difficulty directly from their node's custom dashboard or by using the node's RPC commands.

How Has Network Difficulty Behaved Historically?

Since Bitcoin's inception, network difficulty has grown from 1 to 48.71 trillion. This means it is 48.71 trillion times harder to mine a Bitcoin block today than when mining first began in 2009 – a compound increase of 20.64% per month.

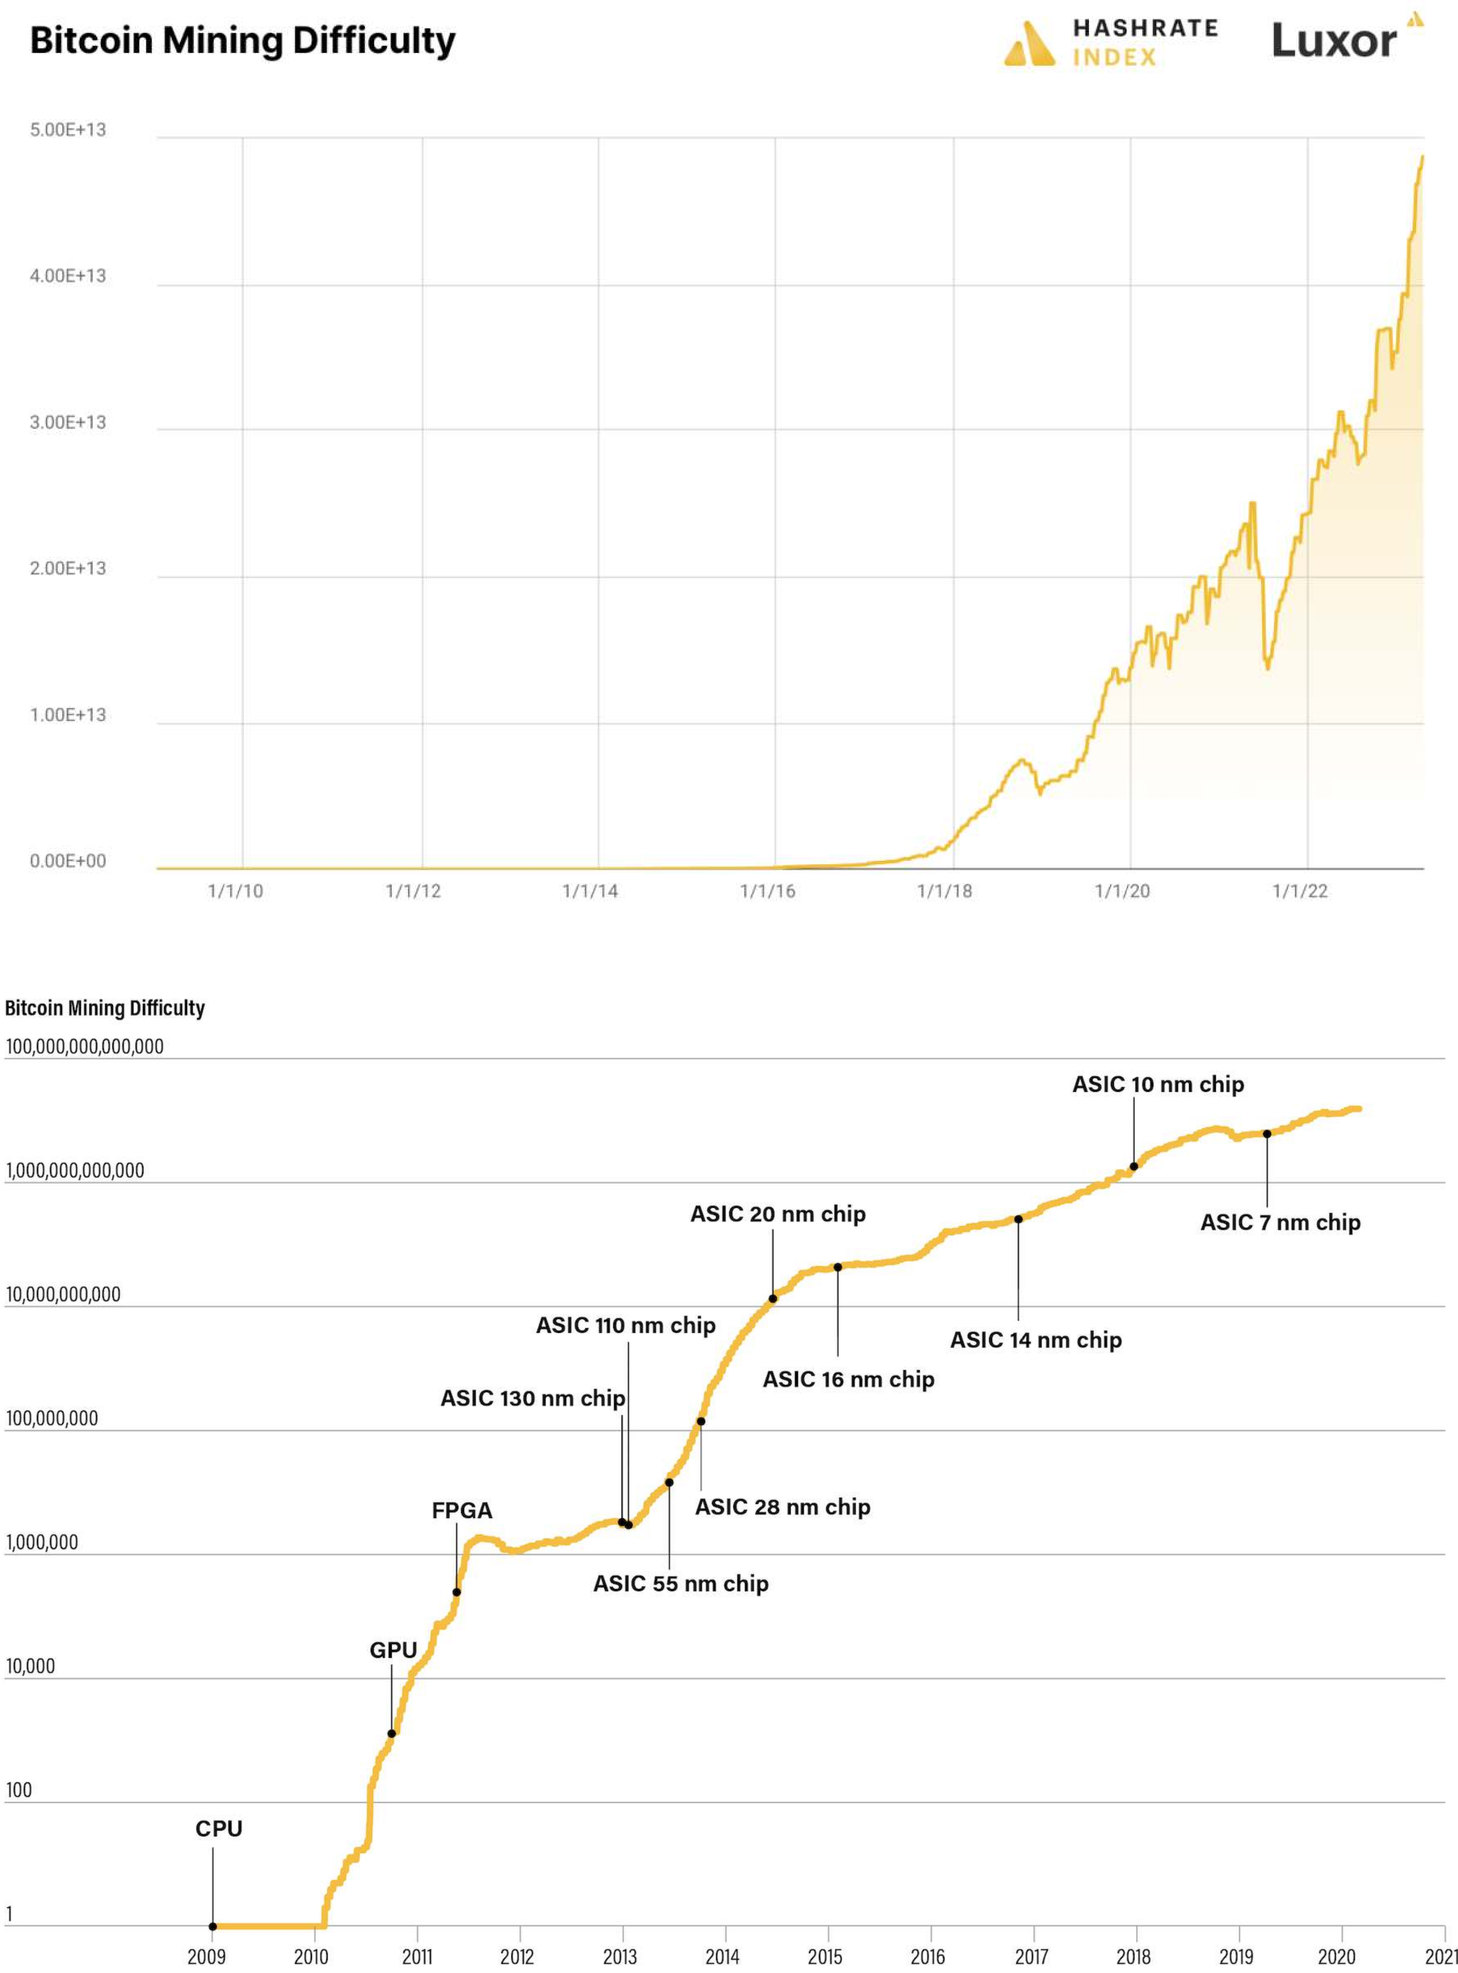

However, when we look at difficulty adjustments over time in the previous chart from CoinDesk, we see that the growth rate is positive, but this growth is declining over time. Large decreases in Bitcoin's mining difficulty sometimes follow political and policy decisions (e.g., China's mining ban), which in turn can lead to larger increases in difficulty once hashrate relocates and comes back online. The most dramatic increases in network hashrate have coincided with the adoption of new, more energy efficient mining hardware (e.g., when GPUs replaced CPUs and when ASICs replaced GPUs/FPGAs). Bitcoin price changes can also significantly impact difficulty changes as hashrate supply comes online to capture excess margins.

Year	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
Avg. Difficulty Adjustment	45.28%	18.03%	3.74%	21.84%	12.81%	3.65%	4.36%	7.02%	4.27%	3.37%	1.62%	1.50%	1.56%

What Determines Network Difficulty?

Block Times

Bitcoin network difficulty is a function of block times and previous network difficulty. Block times, in turn, are determined by network hashrate and luck.

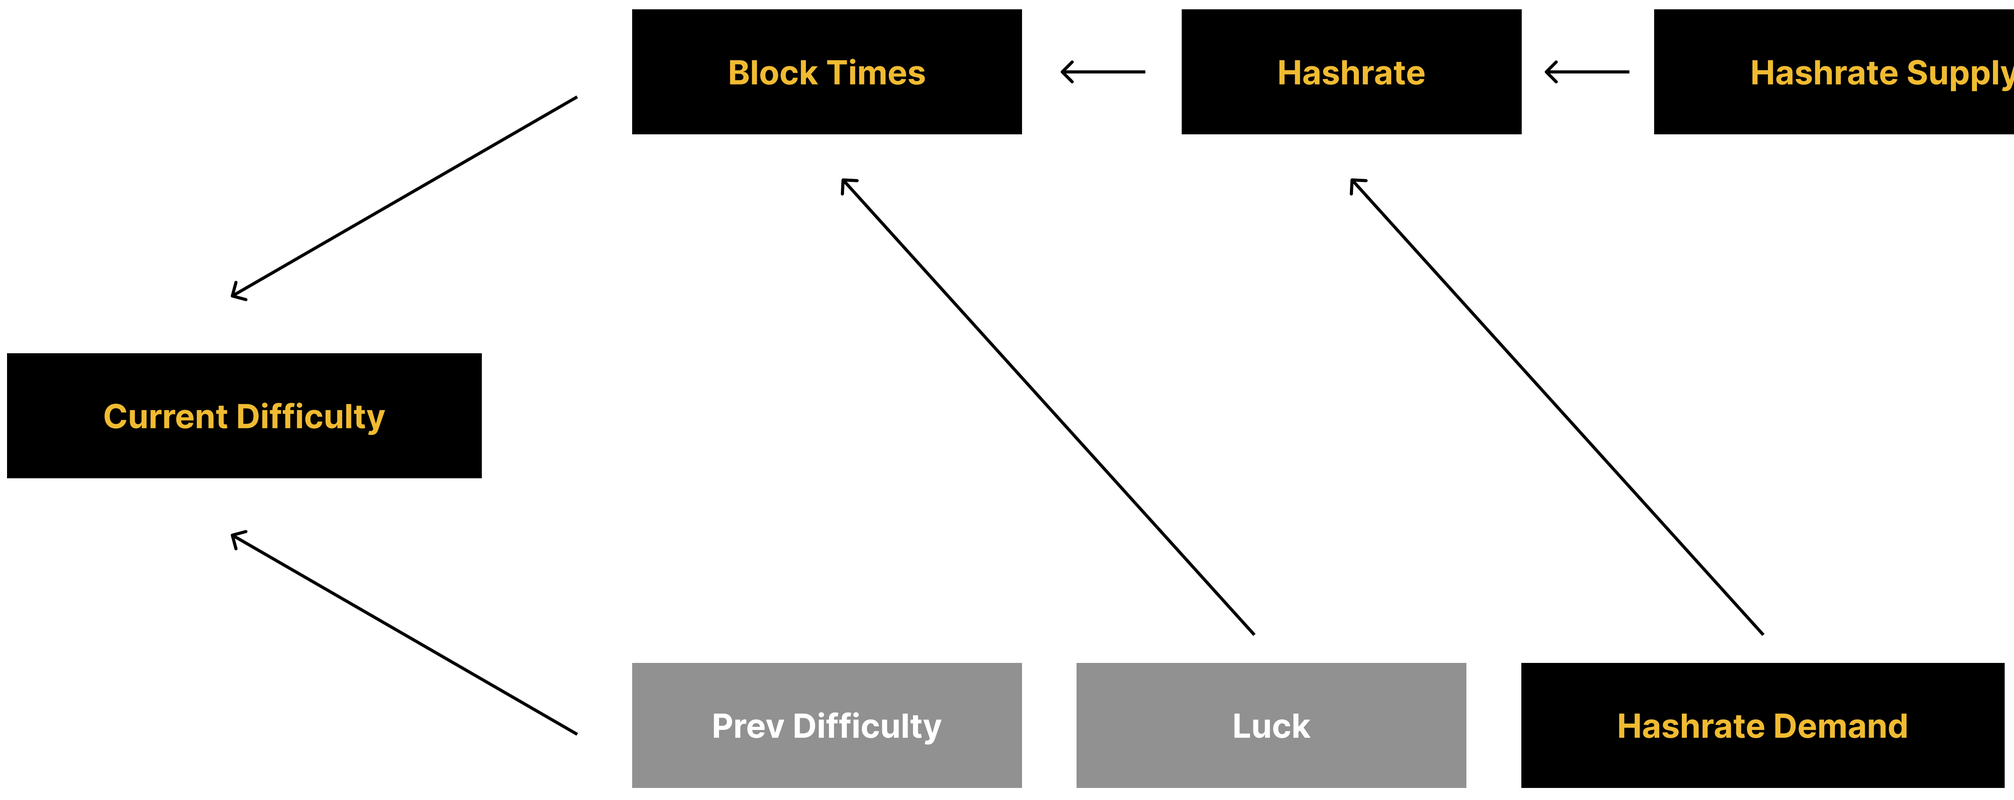

$$Difficulty_t = f\ (\ Block\ Times_{t-1}, Difficulty_{t-1}\ )$$

$$Difficulty_t = f\ (\ Hashrate_{t-1}, Luck_{t-1}, Difficulty_{t-1}\ )$$

$$Difficulty_t = f\ (\ Hashrate\ Supply_{t-1}, Hashrate\ Demand_{t-1}, Luck_{t-1}, Difficulty_{t-1}\ )$$

Network Hashrate: Supply, Demand, and Mining Luck

In contrast with block times and difficulty, data points which we can observe on the Bitcoin blockchain in real-time, we can only estimate network hashrate and mining luck. For this reason, popular Bitcoin mining data platforms produce network hashrate estimates based on block times.

Hashrate

Hashrate Index and other websites typically estimate Bitcoin's hashrate across three simple-moving-averages (SMA): 3 days (432 blocks), 7 days (1,008 blocks), and 30 days (4,320 blocks). The 3 day SMA is useful because it is very current. You can more easily spot massive disruptions to Bitcoin mining hashrate from events like seasonal curtailment on ERCOT or China's Mining Ban. The downside of the 3 day view is that faster or shorter blocks influenced by luck can distort the hashrate estimate, making Bitcoin's total hashrate appear larger or smaller than it really is.

While less current than the 3 day, the 7 day SMA hashrate metric is less influenced by Bitcoin mining luck, and so miners see it as a more accurate estimate. The 7 day SMA is the industry standard for hashrate reporting, but suffers more lag than the 3 day. Lastly, the 30 day SMA smooths out most of the noise caused by luck but greatly lags short-term trends.

Hashrate Demand

In a Q4 2022 Hashrate Index post, we developed a supply and demand model of hashrate and hashprice. Like markets for other commodities, we argued that price and quantity – hashprice and hashrate – are determined by the interaction of supply and demand.

On the demand side, in contrast to a typical commodity market with standard buyers, demand for hashrate comes from the blockchain in the form of the block reward. Bitcoin miners have the option tlo deliver hashrate directly to the blockchain if they are self mining, or indirectly through a mining pool. In return, they receive a share of the Bitcoin block reward, which has a U.S. dollar equivalent market price. As such, our model treats Bitcoin price, block subsidy, and transaction fees as exogenous determinants of hashrate, difficulty, and hashprice on the demand side.

Hashrate Supply

On the supply side, with several different cost profiles, there is an overall amount of hashrate that miners are willing to supply to the network at various hashprices. Theoretically, all else being equal, miners should be willing to provide more hashrate the higher hashprice goes to capture excess profits. Conversely, the lower hashprice goes, a greater share of miners become unprofitable, so they shut off operations and reduce Bitcoin's total hashrate in the process. Our model uses estimates of global ASIC capacity and the distribution of operating costs as exogenous determinants of hashrate, difficulty, and hashprice on the supply side.

For long term planning, we believe it is important to recognize **that difficulty and hashprice are determined endogenously by the interaction of supply and demand**. This is in contrast with some other models and techniques we have seen, which treat hashrate and difficulty as exogenous determinants of hashprice, or vice versa. In reality, this can't be the case. While difficulty determines hashprice, difficulty is also simultaneously determined by hashprice and miner profitability. An accurate model should capture and account for this dynamic.

Mining Luck

We can estimate Bitcoin mining luck (which is a measure of how many blocks a miner produces compared to their expected production) by looking at self-reported hashrate of mining pools and comparing it to actual block times. Of course, this assumes pools publish accurate data, which is not always the case. Mining pools have no legal obligation to report an accurate amount of hashrate on their API. However, if the self-reported hashrate from mining pools is accurate, then we can calculate luck and block time probability.

The following table presents the statistical distribution of mining luck outcomes over a 2,016 block epoch, where mining luck is represented by the ratio of estimated hashrate based on block times over the self-reported hashrate of mining pools.

The higher this number, the more lucky the network would be mining Bitcoin, and vice versa if the number is low. The cumulative distribution function is the probability that a given luck ratio will take a value less than or equal to the value in the corresponding column. For example, over a 2,016 block period, we should only expect a luck ratio of 97% or less 9% of the time.

Luck Ratio	94%	95%	96%	97%	98%	99%	100%	101%	102%	103%	104%	105%	106%
Cumulative Distribution Function	0%	1%	3%	9%	18%	33%	50%	68%	82%	91%	96%	99%	100%

Suppose, based on reported hashrate and actual block times, a hashrate market participant believes a recent difficulty adjustment was impacted by highly lucky or unlucky mining. For a 2,016 block epoch, this could have a +/- 6-8% impact in extreme scenarios.

Pool Hashrate and Luck Factor Scenarios

Pool	Self Reported Hashrate (EH)	Market Share (%)	Expected Block (24h)	Actual Blocks (24h)	Luck Factor
Foundry	10.0	16.7%	24.00	23	95.8%
Antpool	9.0	15.0%	21.60	24	111.1%
F2Pool	8.0	13.3%	19.20	17	88.5%
Binance Pool	7.0	11.7%	16.80	20	119.0%
ViaBYC	6.0	10.0%	14.40	10	69.4%
Braiins Pool	5.0	8.3%	12.00	12	100.0%
Luxor	4.0	6.7%	9.60	10	104.2%
BTC.com	3.0	5.0%	7.20	7	97.2%
Poolin	2.0	3.3%	4.80	4	83.3%
Other	1.0	10.0%	14.40	10	69.4%
Total	**60.0**		**144.0**	**137.0**	**95.1%**

In the example above, we show self-reported hashrate by mining pool and actual hashrate estimates, and we calculate a corresponding luck factor for the first 24 hour period of an epoch. During the first 24 hours, mining pools are reporting ~95% luck factor, meaning they have been relatively unlucky. In this case, we could expect block times to speed up as miners' fortunes revert to mean luck. In a situation like this, hashrate estimates based on block times would underestimate total network hashrate.

Under such scenarios, it may be profitable to trade network difficulty products like Luxor's Hashprice Non-Deliverable Forward on the assumption that luck should revert to the mean. However, in practice, trading hashrate around network mining luck is very difficult because it relies on assumptions, cannot be directly observed, and only has a short-term impact. Over time, network mining luck will even out.

How Can We Forecast Network Difficulty?

There are three basic approaches to forecasting, which we can employed in a number of ways to estimate future network difficulty:

1. Qualitative techniques
2. Time series analysis and projections
3. Causal models

We will describe each of these approaches and provide examples for how they can be employed to forecast Bitcoin network difficulty. For more information on approaches to forecasting, we encourage readers to check out How to Choose the Right Forecasting Technique from the Harvard Business Review.

Qualitative Techniques

Qualitative techniques are the most basic approach to forecasting. They are usually employed when data is scarce or resources are unavailable for quantitative techniques. These techniques take expert opinion, human judgment, and information about special events to apply rating schemes to turn qualitative information into quantitative estimates. Examples of qualitative techniques include the Delphi method, market research, panel consensus, and historical analogy.

We can think of two areas in particular where qualitative techniques could be useful for forecasting network difficulty. The first is for aggregation of forecasts and research from third party sources. Perhaps, instead of creating and running their own quantitative model, a Bitcoin mining company may source forecasts and relevant information for forward planning from third-party sources, and with a predetermined methodology, use the acquired information to produce a network difficulty forecast.

A second area where qualitative methods are useful are public policy changes. Take for example China's mining ban in May 2021. This event was massive in terms of its impact on network difficulty and hashprice. Bitcoin had its largest downward adjustment ever, falling 27.9%, as well as five negative adjustments out of six epochs, for a total difficulty drawdown of 42.4%. However, this event would not be captured in historical time series data and may be missed with more quantitative economic forecasting approaches. Therefore, while qualitative methods may seem rudimentary, they do serve a purpose for forecasting network difficulty and forward planning in hashrate markets.

[1] *https://hbr.org/1971/07/how-to-choose-the-right-forecasting-technique*

Time Series Analysis and Projections

Time series forecasting relies on historical data and the trends and patterns this data reveal. These techniques are based on the assumption that previous trends will continue into the future. As you might expect, these techniques are best suited for situations where historical data is available and abundant, and trends and patterns in the data are clear and stable. In general, these analyses are better for short term forecasts, as they are likely to miss turning points in the data where historical patterns change. Examples of time series forecasting techniques include moving averages (e.g., ARIMA), exponential smoothing, and trend projection.

Before preparing a quantitative forecast, it's helpful to perform a time series analysis. It involves identifying and explaining trends in the data, growth rates in the trend data, and cyclical or seasonal patterns. For example, in the section above on historical network difficulty data, we identified a growth trend in Bitcoin network difficulty which has remained positive but declined over time. To evaluate the seasonal impact on difficulty, the table below shows the average difficulty adjustment by month for different historical periods relative to the average difficulty adjustment for all months. In line with reports from jurisdictions like Texas about the seasonal behavior of miners due to weather, energy prices, and demand response, seasonal patterns in network difficulty adjustments appear to be strengthening as the bitcoin mining industry matures.

Avg. Difficulty Adjustment By Month, Relative to Avg Difficulty Adjustment All Months			
Month	2018 to Present	2015 to Present	All Adjustments
January	2.34	1.78	0.81
February	1.24	1.54	1.44
March	0.67	0.53	0.67
April	1.46	0.89	0.96
May	0.39	0.57	0.99
June	0.90	0.92	1.12
July	0.15	0.49	1.76
August	2.18	1.17	1.02
September	1.73	1.54	1.03
October	1.48	1.24	0.88
November	-0.40	0.15	0.66
December	-0.47	1.11	0.63

As we described above, time series forecasts are best suited for shorter time periods. In the context of forecasting network difficulty, this approach is well suited for short term average difficulty forecasts (i.e., up to a few months) or for forecasting block times for upcoming difficulty adjustments. With respect to block times, Hashrate Index publishes estimates of the upcoming difficulty adjustment using the constant block time method, which we evaluated in this post. The constant block time method provides an unbiased forecast which is highly accurate near the end of each epoch, but extremely inaccurate at the beginning of the epoch.

For this paper, we have developed a new time series forecasting method for upcoming difficulty adjustments, which improves accuracy at the beginning of the epoch compared to the constant block time method. We call this the succinctly named rolling 2,015 block, square root weighted, epoch adjusted block time method (or just "rolling block method," "adjusted block time method," or "dual epoch method").

This new method improves upon the constant block time method early in the epoch by including block times from the previous 2,015 blocks, instead of just the blocks from the current epoch, which can skew forecasts early in the epoch for lack of data points. To account for the change in network difficulty between epochs, block times in the previous epoch are adjusted by the previous adjustment. And finally, we weight the average block times of the current epoch with the square of the proportion through the epoch. This final step is to diminish the impact of block times from the previous epoch as the current epoch progresses since these values do not actually determine the upcoming adjustment. The method is represented by the following formula:

$$Estimated\ Difficulty\ Adjustment\ =\ (\frac{600}{Adjusted\ Block\ Time}-1)\ *\ 100\%$$

Where,

$$Adjusted\ \ Block\ Time\ =\ \sqrt{\frac{Block\ \#\ in\ Current\ Epoch}{2{,}016}}\ *\ (Avg.\ Block\ Time_{Current\ Epoch})$$

$$+\ (1-\sqrt{\frac{Block\ \#\ in\ Current\ Epoch}{2{,}016}})\ *\ (Avg.\ Block\ Time_{Previous\ Epoch})\ *\ Latest\ Difficulty\ Adjustment$$

The table on the following page compares the overall (i.e., through an entire epoch) forecasting accuracy of the constant block time method and the dual epoch method. While the new method is slightly more biased, it is far more accurate overall.

[2] *We would like to thank Haley Thomson of Imperium Digital for inspiring this idea.*

Measure of Forecast Accuracy	Constant Block Time Method	Rolling 2,015 Block, Square Root Weighted, Epoch Adjusted Block Time Method
MFE	-0.07	0.97
MAFE	3.65	2.68
MSE	852.65	16.44
RMSE	29.2	4.06

In the chart below, we can see through confidence intervals that the new method performed better than the old model at the beginning of the epoch up to block 650, but it performed slightly poorer thereafter.

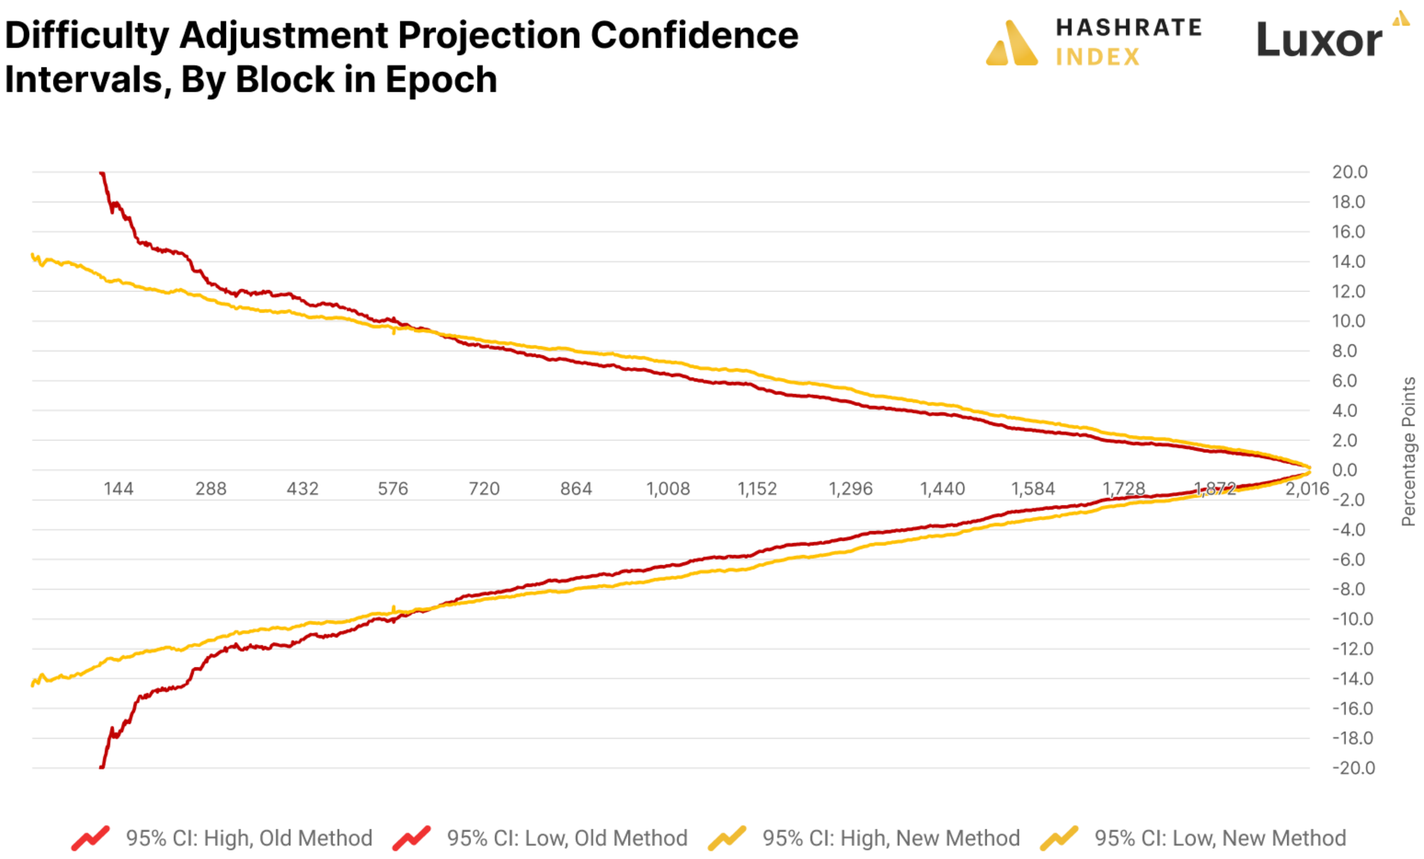

Causal Models

Causal models are the most sophisticated approach to forecasting. Forecasting with causal models involves analyzing historical data to identify relationships between related variables or elements and constructing mathematical, multi-variable models to create projections. They are best suited for medium-to-long-term forecasting and for anticipating trend reversals or turning points, but they can over-complicate simple forecasting exercises. Examples of these types of forecasting techniques include regression and econometric models, as well as leading indicator models. The most sophisticated techniques leverage artificial intelligence and machine learning.

As described earlier, Luxor has developed its own supply and demand model of hashrate and hashprice. It takes Bitcoin price, transaction fees, and block subsidy as inputs on the demand side, and internal data on ASIC production estimates and operating cost distributions across the industry on the supply side. Using these input, the model produces an equilibrium hashrate, difficulty, and hashprice for 18 month periods. The model structure reflects reality; hashrate, difficulty and hashprice are endogenous to the system, not exogenous determinants of one another. We can conduct sensitivity analysis with the model across all inputs as well. For example, we can forecast an equilibrium hashrate, difficulty, and hashprice across a range of Bitcoin prices.

The charts below present projections from our updated hashrate supply and demand model. It provides estimates for flat, bull, and bear Bitcoin price scenarios.

Flat case hashprice and hashrate

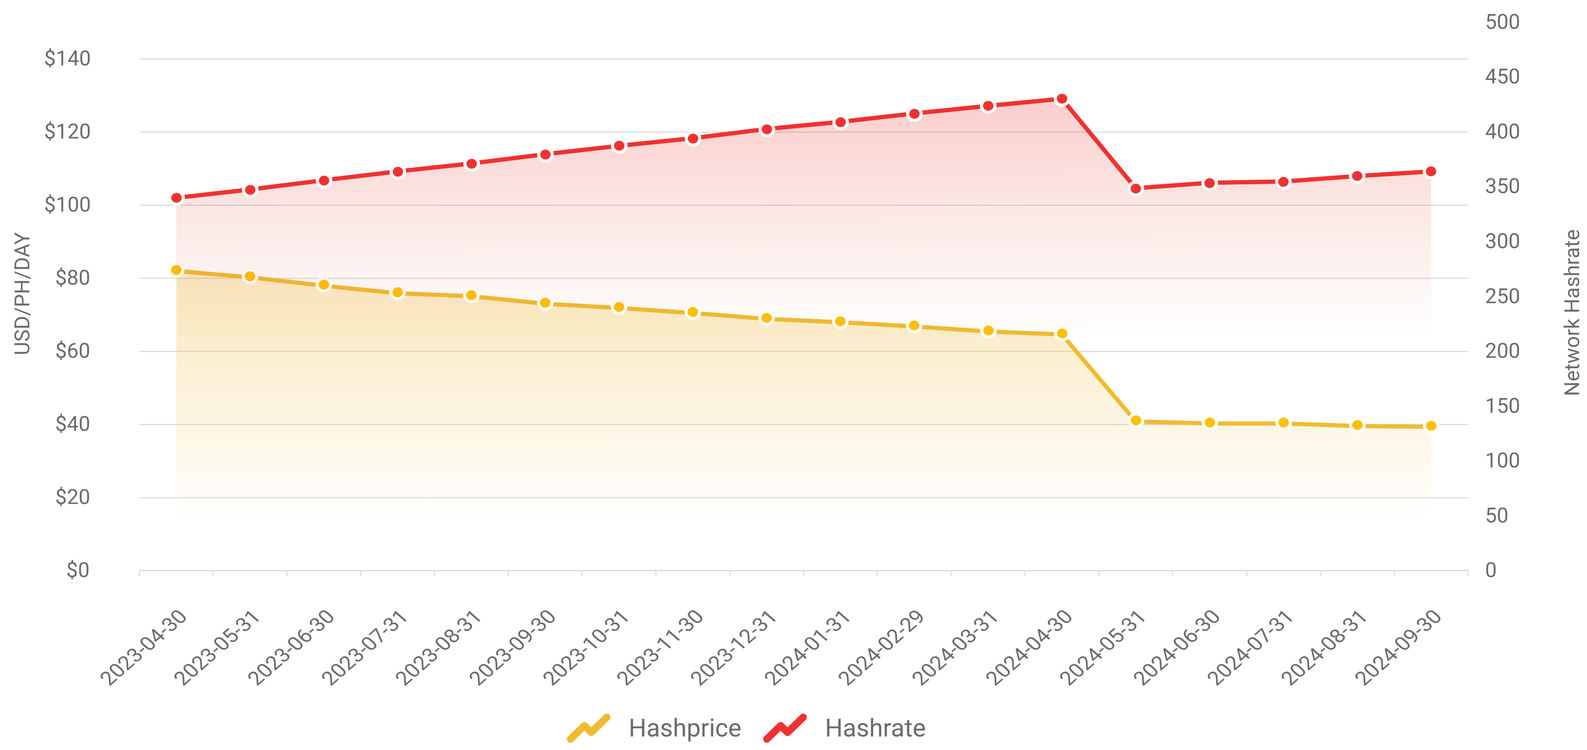

Bull case hashprice and hashrate

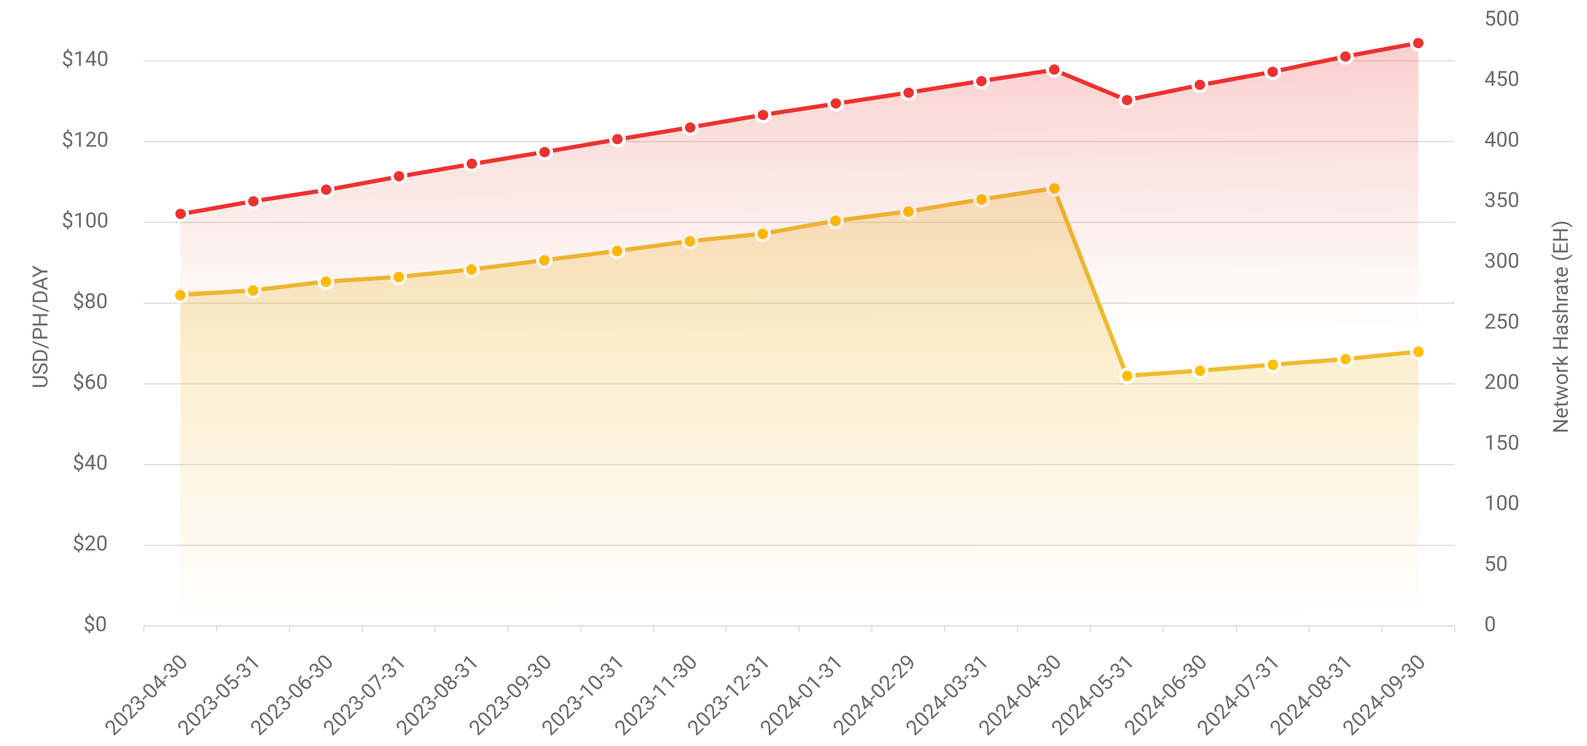

Bear case hashprice and hashrate

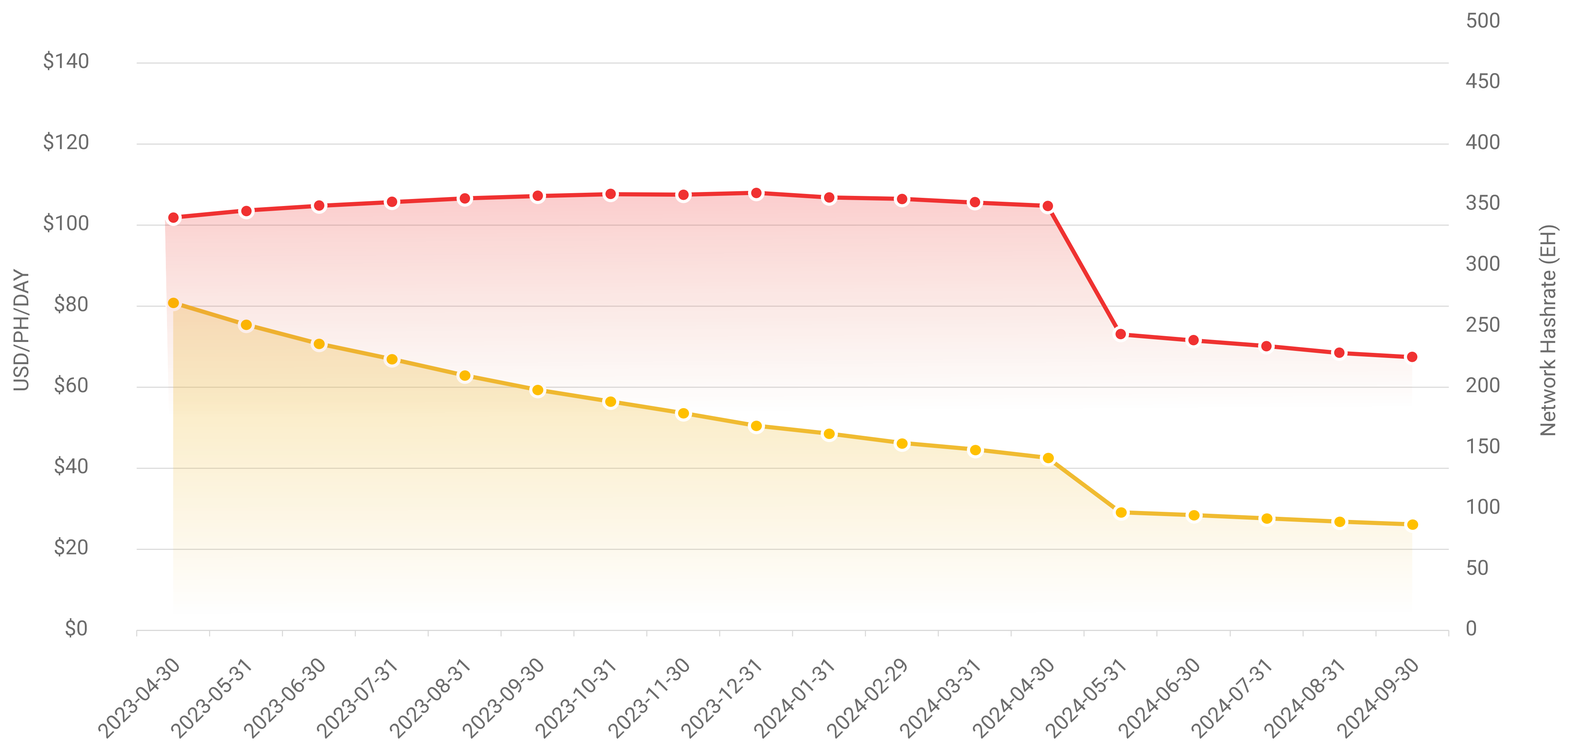

[3] Flat scenario assumes a $30,000 Bitcoin price and 0.18 BTC transaction fees. Bull scenario assumes 5% monthly Bitcoin price growth from $30,000 to $72,000 in September 2024 and 0.18 BTC transaction fees. Bear scenario assumes 5% monthly Bitcoin price declines from $30,000 to $11,000 in September 2024 and 0.10 BTC transaction fees.

Hashrate Index Premium Hashrate, Difficulty, and Hashprice Projection Updates

Hashrate is an emerging asset class and digital commodity market. Hashrate market participants like Bitcoin miners, hosters, lenders, investors, and traders need access to the rigorous economic analysis and data available in other commodity markets. As such, and due to positive feedback we received on version 1 of our hashrate supply and demand model, Luxor has decided to dedicate resources to continually refining and improving our hashprice, hashrate, and difficulty models and forecasts, which will be available to Hashrate Index Premium customers on a quarterly basis.

Hashrate Index Premium customers will receive the following in PDF and Excel format:

- 18-month, flat, bull, and bear BTC price scenario projections;
- Full Bitcoin price sensitivity tables (i.e., an equilibrium hashrate, difficulty and hashprice will be provided across all probable Bitcoin prices).
- A detailed description of assumptions, which will be continually refined and improved.

Please reach out to hashrateindex@luxor.tech if you are interested in subscribing to quarterly updates of Luxor's hashprice, hashrate, and difficulty models and forecasts.