The Three Woes of Bitcoin's Fee Market (and How BlockDAGs Can Fix Them) – Part II: Blockchain Fee Markets

Parallel Thoughts (Shai Deshe) · 11 Mar 2025

> This post was written as a cursory version of a future part in my open book, understanding blocckDAGs and GHOSTDAG

<< Part I, Part III, Part IV.

In the previous post we qualitatively surveyed the Bitcoin fee market and introduced some concepts and machinery that would help us quantify our observations. In this post, we will apply these tools to better understand the fee markets of block*chains*.

The Pure Fee Market Game

We model the fee market as a game played between many *users* and a single *miner*.

In this game, there are two types of turns, many *offering turns* followed by a single *payoff turn*. In an offering turn, users can add transactions or increase the fees of the transactions they added in previous offering turns. In the payoff round, the miner chooses up to ℓ transactions to include in the block and gains their fees. A *round* is comprised of a payoff turn following the procession of offering runs since the previous payoff turn.

Let m be the number of transactions at the start of the payoff turn. We call a round *congested* if $m > \ell$, and *non-congested* otherwise. We assume that this property does not change during the round. This captures the reality that most of the time, a blockchain throughput is not at the exact capacity, and typically has few and far between transitions between prolonged periods of under or overutilization. In other words, we assume that m is "rather far" from ℓ.

The utility of the miner is to maximize their profit, the utility of the users will be discussed shortly.

This description already seems a bit odd. First of all, why is there only one miner, shouldn't there be many? Isn't this the point? Other than that, it is clear that a rational miner will just pick the ℓ highest paying transactions available (or all transactions if there are fewer than ℓ). So if their behavior is so simple, why even bother including it in the model, and not just say the users compete for space in the top ℓ paying transactions?

For the first question, yes, when we talk about *mining* there are indeed many miners. But currently, we are *not* talking about the mining game but about the *transaction selection game*. In the latter game there is *a single miner* who gets to choose the transaction. The former game is how that single miner is *chosen*.

So why even bother with modeling the miner as a participant at all? Because this paves the way for further down the post when we compare blockchains to blockDAGs. It is instructive to model the two as closely as possible, so while considering the miner a rational entity (and not just a deterministic rule of the game) might ostensibly seem superfluous, it will serve us well.

The utility of the *user* requires a more nuanced discussion. The user utility has two parameters: *speed* (measured in how many rounds they have to wait before their transaction is chosen) and *cost*. In general, a user wants to transact fast and cheaply. Their willingness to trade off one to the other can be very complex and may depend on many external factors such as the time of day or the weather. Worse yet, each and every user has a different utility function. So how can we tame this complexity? We agree on *some* properties that arguably apply to almost all users, two of them to be exact:

- Users will always avoid paying more to get the same (or worse) service. In other words, if a user can *pay less* to transact *(at least) as fast*, they would.
- Users will always agree to pay *very little more* (say, two sats) to transact *considerably faster* (say, ten minutes faster).

Before moving forward, we should note that realistic fee markets are *not* pure, and are actually subject to many externalities. Miners can be coerced by many means, from threats to bribery. They have ways to obtain coins beyond mining them and might be interested in other valuables. Grievance attacks are not unheard of. Even the assumption that the number of miners and their proportion is fixed is unrealistic. However, this simplified model captures the most dominant dynamics of a healthy fee market and exhibits enough intricacy to derive very meaningful conclusions. Arguably, more realistic models of fee markets should be obtained by *refining* this model rather than *discarding* it.

Equilibria and Aberration of non-Congested Networks

Consider the situation where we know the number of transactions will not exceed m by the payoff turn. What are the consequences? Well, in this case, users know that they will be chosen by the miner in the next payoff round *regardless* of how much they pay. Recall that we said a user will always pay less for the same speed if they can. This implies that in this situation, users will pay *zero* fees. This decline in fees is our first woe, which we call **race-to-the-bottom**.

The consequence of race-to-the-bottom is that to have **any** security budget, congestion is **required**.

There are reasonable objections to this analysis, claiming it is just a bit *too* oversimplified, and I agree. However, we only need to slightly complicate the model to dispel the central objections.

First objection: transaction inclusion isn't *actually* free.

In reality, the effort of including a transaction has a slight cost. If the fee is actually zero, then the miner has nothing to gain, but a bit to lose, by including the transaction. In principle, even the electricity cost of verifying the transaction validity, or the weight it adds to the block, makes it more profitable to ignore it. Not to mention the costs of the mining operation itself. When taken into account, the actual equilibrium does not plummet to zero, but to some positive number. If this number is not enough to cover the miner's operational costs, mining becomes unprofitable and the coin security collapses to a level the network can afford. But even if the fees are sufficient, they are stagnated, and the nice feedback between an increase in security and in value is broken.

Second objection: what if miners agree on a minimal fee?

Say the miners agree to ignore transactions with fees below a certain threshold T. Any transaction paying a fee of less than T will have to wait forever. If T is reasonable, most people will agree to pay this cost to transact.

The thing is that to model this strategy, we can no longer get away with the assumption that there is only one miner. We need to take all miners into account. And when we do, we notice that this strategy is *unstable*. At any point, any miner can suddenly agree to take transactions that pay a fee slightly lower than T, because that would increase their utility. This creates a race-to-the-bottom between *the miners*. In fact, when Bitcoin theorists talk about a race-to-the-bottom, this is the one they usually mean. Miners racing with each other to accept lower and lower fees, all the way to bankruptcy.

The equilibrium of this process is very hard to estimate, as it relies on many real-world considerations that are not amenable to modeling. But the mere possibility that, as long as the network is typically non-congested, a small change of tide could suddenly make the entire mining industry highly unprofitable is, well, a problem.

Third objection: what about users who try to "cut in" by only posting transactions right before the payoff turn? Wouldn't they incentivize users to pay higher fees to avoid being boxed out in the last second?

This is another example of something that can happen in the game as we described it, and not in reality. We assumed that the number of payoff rounds is known in advance. In reality, it is **not**. Block creation follows something called a Poisson process. When we say that a block is created once every ten minutes, this does not mean that once a block is created you can measure ten minutes to know when the next block arrives. It means that if we divide time into very small intervals, say, microseconds, and assume that the block time is, say, 600 seconds, then in any such interval there is a chance of one in 600 million that a block will be produced. If we consider each such interval as an offering turn, then we can rewrite the rules of the game such that in any offering turn, there is a one in 600 million chance that the next turn will be a payoff round. The crux is that this process is independent and it is memoryless. The probability that the next round is an offering round remains the same, regardless of how many offering turns there have been since the last payoff turn. There is no point in time where you have any better chance to guess when the next pay-out round is going to happen, making cutting in impossible.

Now regarding abberation. It is a bit funny to call the consequences of race-to-the-bottom an aberration of the price and not just a complete obliteration of the fee market creating zero prices that represent nothing.

Equilibria and Aberration of Congested Networks

Now say that m is larger than ℓ, sufficiently large to assume the network isn't going to become available any time soon. How do you transact in this reality? You *fight*.

In order to be included in the next block, you need to find the least paying transaction that is about to enter the block (that is, the ℓth highest fee) and pay *one sat more* to guarantee your position. At least, until the next offering round, when the guy you boxed out retaliates by increasing *their* transactions by *two* sats, kicking you back out.

Note what's going on here. First, we see that as little as two sats can make a *significant* change to what transactions are included: to change the inclusion probability of one transaction all the way from zero to one, while changing the other all the way from one to zero (a mathematician would say that the probability a transaction is included as a function of how much fee it pays is *highly discontinuous*). This means that the contents of the next block are very erratic and brittle. But that's not the problem. The problem is that the users competing for a seat on the block never make any *significant* increase to the fee. They just keep tossing their spare change into the pile, culminating in a total sum of next to nothing.

This is the price aberration of congested blockchain fee markets: instead of increasing the price to reflect the service they desire, rational users compete over who can make the largest pile of fractions of a penny.

On the other side of congestion, we have *low-paying transactions*. Imagine that the current ℓth highest fee hovers around $100, but you cannot afford to pay more than $10. Will you ever be able to send your transaction? Not unless the fees drop below $10 at some point.

In the previous situation, we had that increasing the fees by just a little the service (in terms of speed) increased considerably. In the last one, we see that outside the high-fee regime, it doesn't really matter how much you are willing to pay (within that regime), the time to inclusion will remain *infinite*. If the ℓth fee is $100, then it doesn't matter if you pay $1, $10, or $50 bucks, you are going to get the same speed: none.

This is the third woe, **starvation**. It raises concerns that as demand increases, using the network will become the exclusive prerogative of a limited oligarchy. While this last scenario might be a bit exaggerated, it demonstrates a tangible tension between Bitcoin's egalitarian ethos and the natural dynamics of its fee markets.

Part III >>

> If you find this content educational, interesting, or useful, please consider supporting my work.