The Three Woes of Bitcoin's Fee Market (and How BlockDAGs Can Fix Them) – Part IV: Discussion and Conclusion

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 III

This short post concludes our series. We started by isolating the three woes of blockchain fee markets. We proceeded to mathematically analyze the market (to an extent) and conclude that they are a consequence of a blockchain's *single leader* dynamic, we then progressed to understand the multi leader dynamic of block*DAGs*, and are now finally ready to leverage the observations we accumulated to understand how if at all, these woes manifest in a multi-leader environment.

Race-to-the-bottom

Obviously, if $m \leq \ell$ then all miners include all transactions, and the fee value has no bearing on quality of service. As we noted, this would cause users to decrease fees (because why pay more to get the same). Miners might try to collude for a minimal fee, but that would create a very unstable equilibrium as a single miner could cut off the rest by reducing their threshold just a little, causing the fee threshold to race-to-the-bottom as well.

However, once there are more than ℓ transactions, fees *immediately* affect the quality of service, *regardless of how many leaders there are*. For N leaders per round, then the maximal capacity of the network is $N \cdot \ell$ transactions per round, which means that the network need only be at $1/N$ capacity for the fee market to activate! If we have 10 leaders per round, it becomes rational to increase fees when as little as 10% of the throughput is utilized!

A blockDAG's ability to arbitrarily increase block rates becomes extremely valuable here, as higher block rates mean more leaders per round, which in turn implies that lower utilization is required to avoid a race-to-the-bottom.

In that sense, we can say that race-to-the-bottom is *not* a blockDAG's woe.

Starvation

Our computations show that for a completely flat market, a transaction will not be starved even if it is a considerably small fraction of the average fee. These computations could be extended to a general fee distribution by replacing F with the correct form of average (if you must know, if there are N leaders, then you should use the power mean $M_{-1/(N-1)}$).

This means that transactions have *some chance* of being included in the next block even if their fee is relatively low. Now, "some chance" might not sound like a lot, but consider this: if a transaction has a probability of p to be included in a round, then it would have to wait an expected number of $1/p$ rounds. In other words, even if you "only" have a 1% chance of inclusion, and a round takes three seconds, then you would have to wait an average of 5 minutes, and the probability you would have to wait more than 10 is practically zero. This is definitely much shorter than the time you'd have to wait in Bitcoin which is **checking notes** forever.

You actually have two things going for you: that it is rather cheap to get a positive, if small, inclusion probability, and that rounds are as fast as network latency (well, at least they are in Kaspa) where the real-world time you wait is inversely proportional to the length of a round.

Given the above, it is safe to say that starvation is not a blockDAG woe.

Price Aberration

Aberration is harder to discuss than the other two woes. When considering race-to-the-bottom or starvation we know exactly what we want: neither. Our job is to make these dynamics as scarce as possible. But what do we want here? The term *aberration* describes a situation where the market dynamics distort the fee-to-service curve, but who is to say what is the "correct", undistorted curve? This is where you come in.

I find the fact that slight changes to fees only incur slight changes to the inclusion probabilities (in contrast to a congested Bitcoin, where two sats can change the included sets of transactions with certainty) a very good sign. It

means that significantly increasing the fee is required to significantly improve your inclusion times. I didn't get into this here, but we can actually see qualitative phenomena where the actual mapping seems to be nicely bound from above by a low-degree polynomial. That is, you won't have to pay 100 times the fees to be included half as fast. And if you decrease your fees by half, the time you'll wait will increase by a factor of maybe 10, but never a 100 or a 1000. Not only is the resulting curve continuous, but it seems rather tame.

I conjecture that a good economy researcher could come up with a reasonably justified model such that the fee-to-service curve quickly converges to the model's supply and demand curve in the "many tiny blocks" limit. If you have an idea for such a model, hit me up, we might become co-authors.

Given the above, aberration *might* be a woe of blockDAGs, but currently the evidence draws a different picture, and I conjecture that it is not.

Epilogue

- The three woes eternally hover above the Bitcoin market, undermining its security. The way to tame them, it seems, is to build around the market a wide wall, made of many parallel blocks.

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