Miner Collusion and the BitCoin Protocol

Alfred Lehar
Calgary

Christine A. Parlour
Berkeley Haas

Motivation

- The promise of Decentralized Finance is to replace existing financial intermediaries with technology.
- Question: Does the decentralized system provide services more efficiently than traditional intermediaries?
- Examine one successful use case: BitCoin Blockchain which provides settlement services.
- What can BitCoin tell us about DeFi in the future?

What we do

- Look at all blocks between Jan 3, 2009 and November 14, 2018.
- For each block, we observe
  1. The Coinbase
     - This is the "block reward" that the Protocol automatically assigns to the successful miner.
     - This was originally set at ₿50, and is cut in half every 210K blocks
  2. All transactions: These comprise at least two addresses (from/to)
     - 353,306,421 Transactions
     - 903,976,764 inputs and 961,308,921 outputs
     - Fees: Difference between ₿ "From" to ₿ "to"

- Try and understand what determines demand and supply for settlement services on Bitcoin Blockchain.

Our focus is on fees/Recompense to the Miners

- The miner who is the first to find the correct number receives the coinbase.
  - This reward is independent of *what* he mines
- In addition, he can keep any excess of fees or ‘inputs over outputs”.
- Fees are offered by agents as an incentive to be included in the blocks.

What we find

1. There is frequently excess capacity in mined blocks.
2. There is large heterogeneity in fees paid to miners.

3. Excess capacity is consistent with "price discrimination."
4. Mining Pools, by reducing the number of separate participants facilitate this.
    - Estimate that mining pools have extracted in excess of $200 million USD per year.

Literature on Fees

- Two papers (that we know of) have looked at "fees" and mining efficiency
  1. Easley, O'Hara and Basu "From Mining to Markets: The evolution of BitCoin transaction fees
     - Implicitly assumes that all blocks are full.
     - Mining efficiency/cost determines entry and hence equilibrium efficiency
  2. Huberman, Leshno and Maollemi "An Economic Analysis of the Bitcoin Payment system"
     - Fixed Capacity forces fees.
     - Model as a queueing problem and pay to jump the queue
     - Efficiency
- Both papers view miners as competitive, and agents offering fees to overcome congestion in the system.

Block Capacity

- Blocks have a fixed capacity. $\Longrightarrow$ potentially a scarce resource.
- Important to our analysis that capacity is finite
- There have been three capacity changes.
  1. Implicit limitation due to data base design – 500-750 Bytes
  2. 1MB limit in original design re-established on May 15, 2013
  3. 4MB "Segwit" on August 24, 2017
     - Not all are Segwit compliant $\Longrightarrow$ weight is the correct size metric.

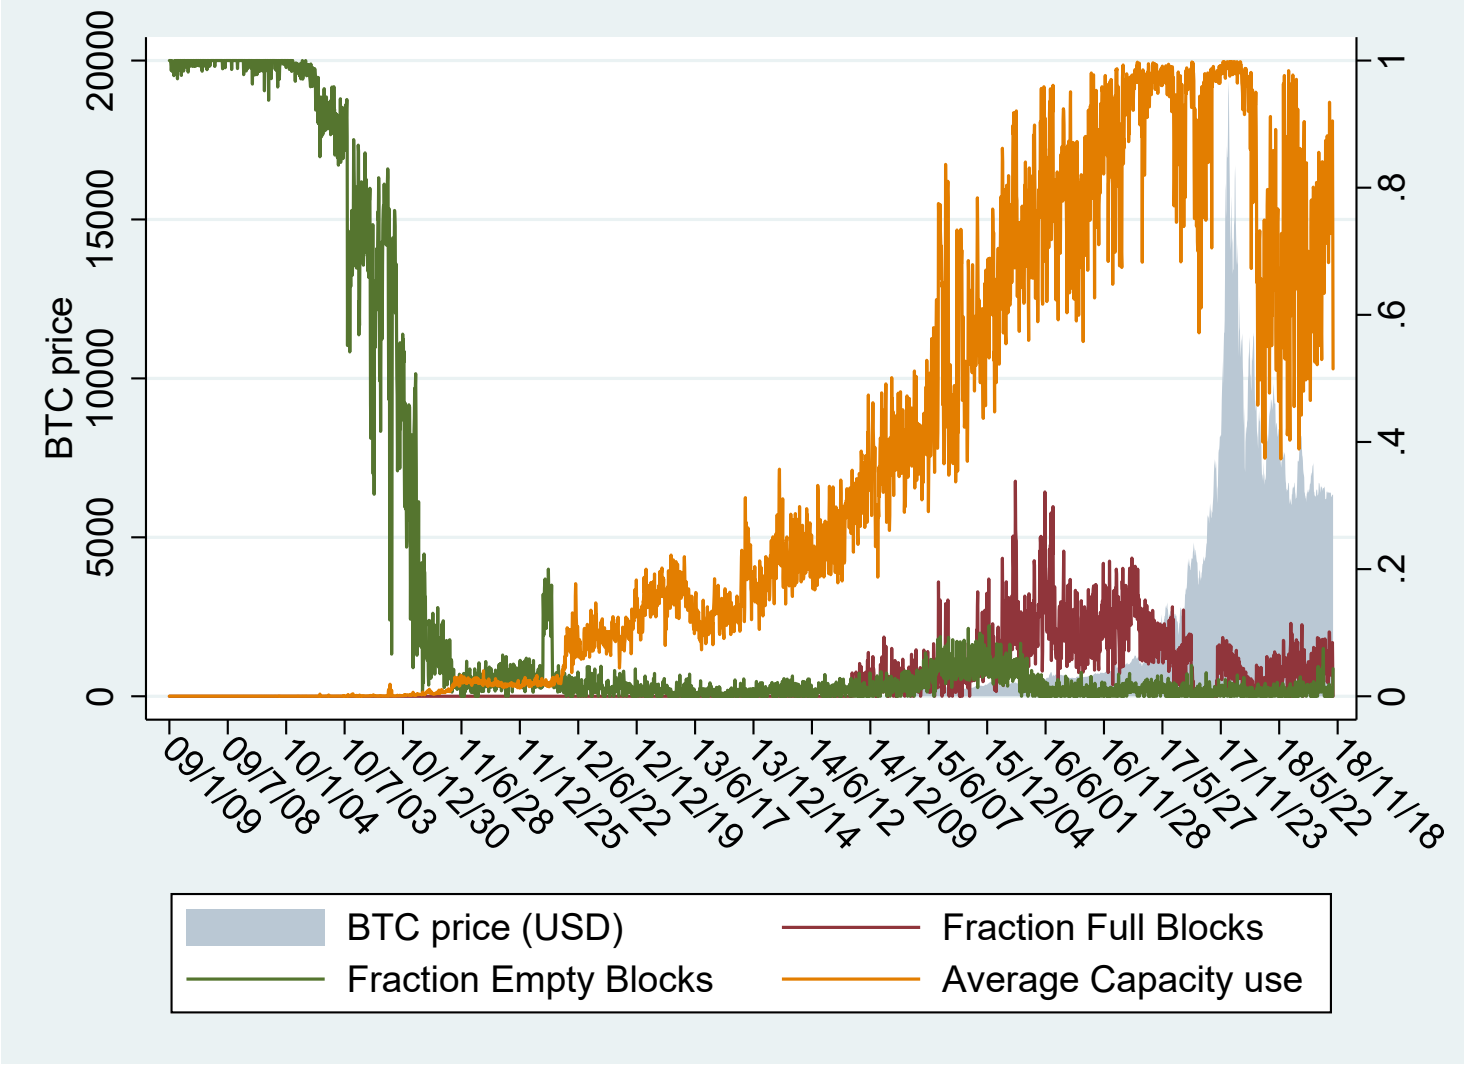

Figure: **Number of blocks and fraction of empty and full blocks per day. Days are defined over UTC.**

Excess Capacity: Empty Blocks

- On Dec 17, 2017, for example, when ₿ was trading at over USD 19,000, two empty blocks, were mined.
  - Block 499704 and 499763 were mined empty by BTTC pool.
  - On the same day the same pool also mined 5 non-empty blocks making a technical problem unlikely.
- During the peak period, 40 empty blocks were mined,
  - 39 were mined by 11 different identifiable mining pools.
  - The largest pool in terms of share of mined blocks, AntPool, also mined the largest number of empty blocks.
- For all blocks mined in or after 2016 about 1.1% of all blocks mined by a given pool are empty.

Excess Capacity: Partially filled blocks

- Excluding empty blocks on August 22, 2017 (before Segwit) there was room for extra 625 transactions.
- On December 17, when ₿ peaked, there was room for 4957 Segwit compliant transactions or 1175 non-Segwit transactions.

- In December 2017, on average room for 25,839 Segwit compliant transactions was left empty.
- For the broader sample since Jan 1 2014 on average there was empty space for 598,190 transactions per day.

Most Transactions pay fees

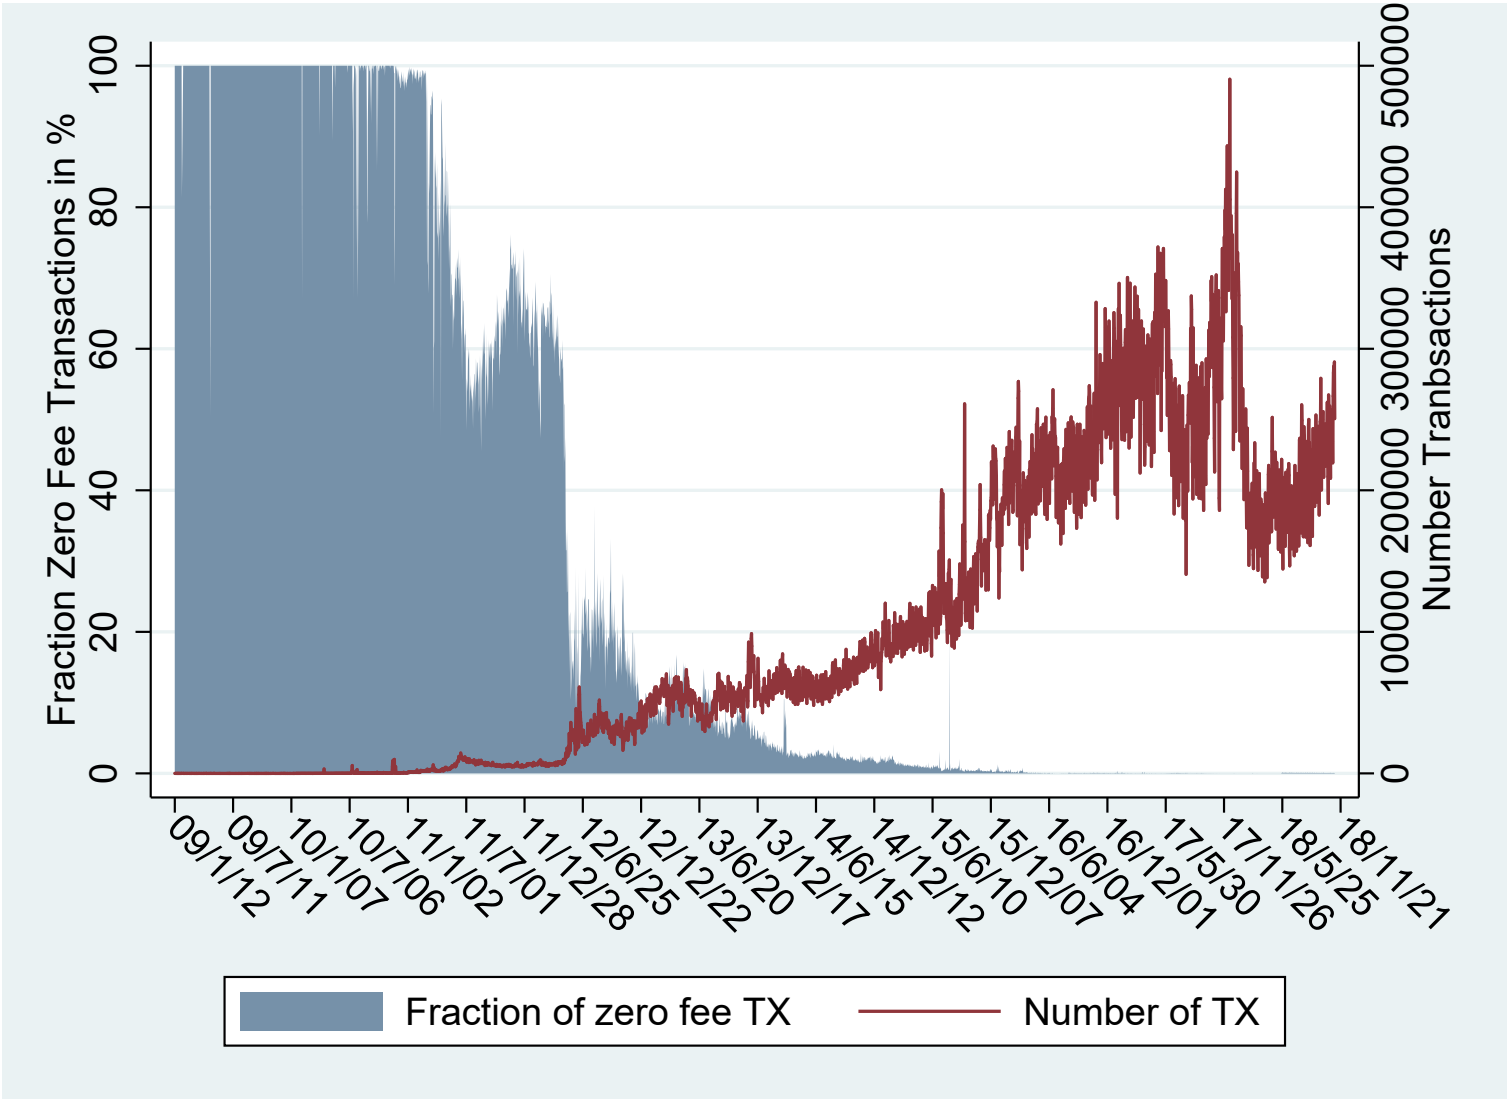

Money Left on the Table

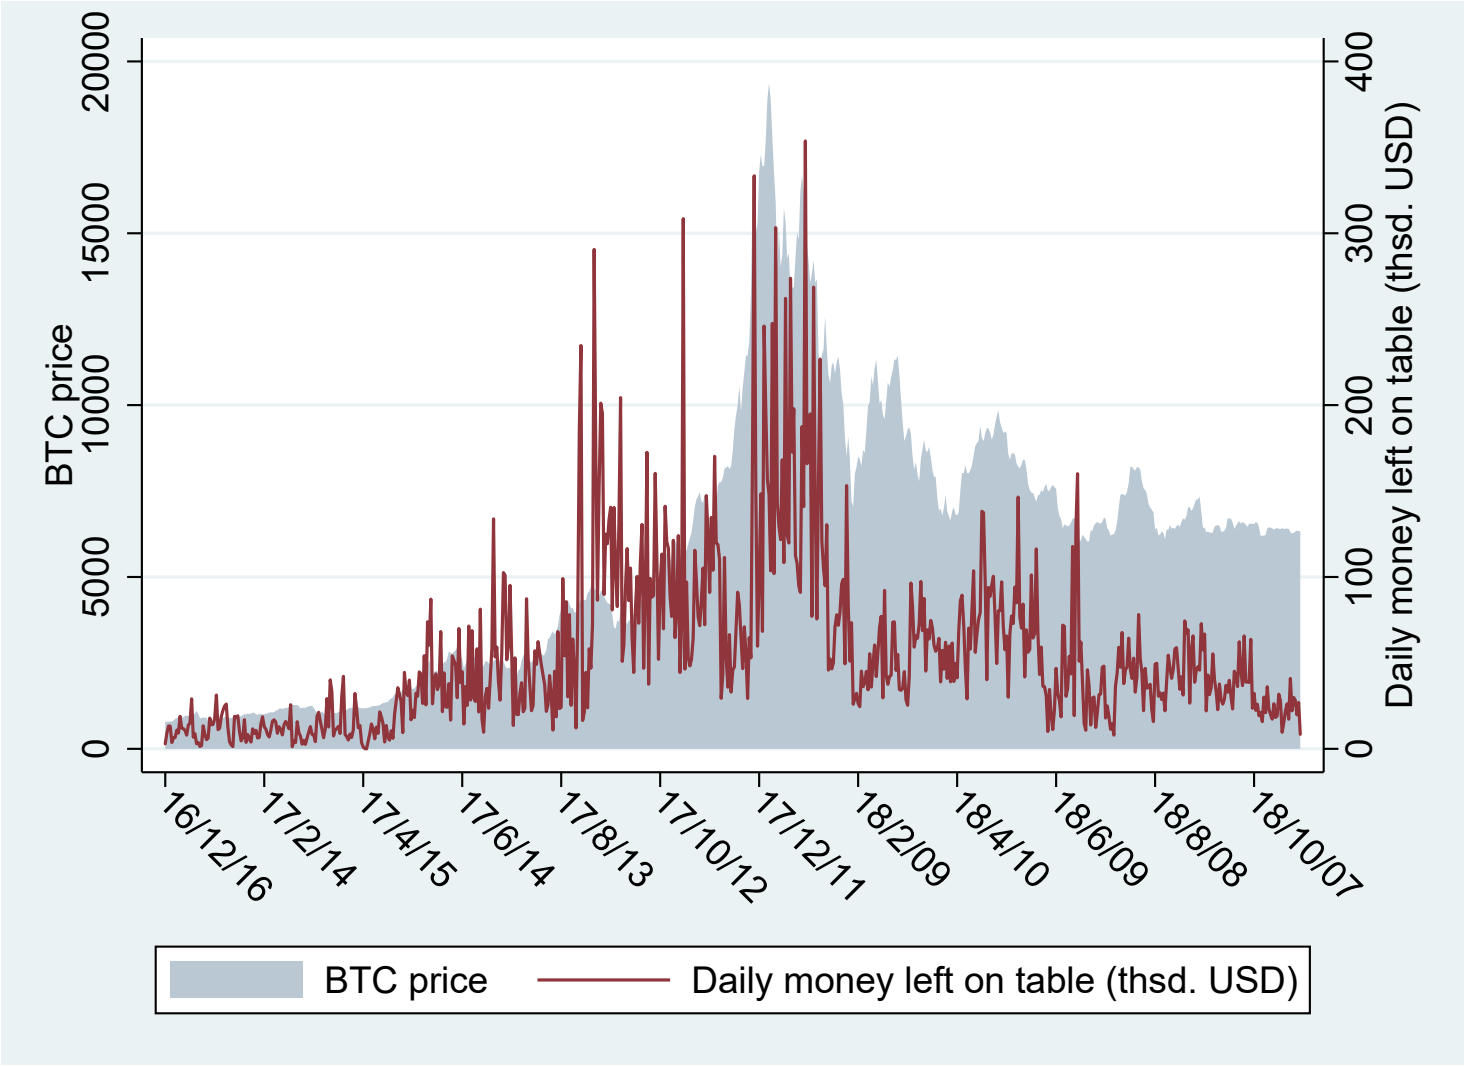

Figure: **Money left on the table by miners and BTC price.** We fill up empty capacity on the blockchain with unmined mempool transactions offering the highest fee/byte.

Huge cross sectional variation in fees paid

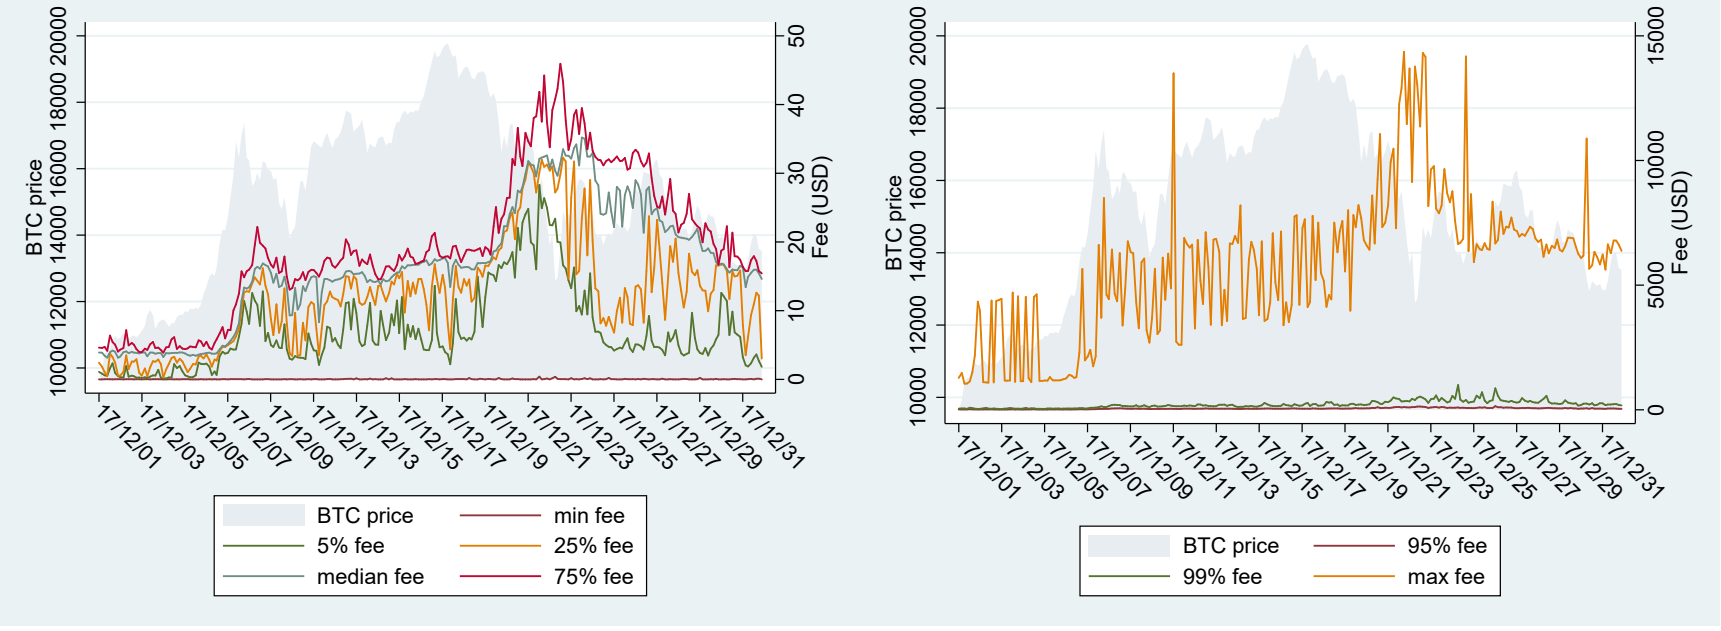

Fees in USD (right axis) and Bitcoin price in USD (left axis) in December 2017.

Huge cross sectional variation in fees paid

- Between Dec. 20^{th} and 24^{th} 2017:
  - 1,674,141 transactions processed
  - 51 Transactions with fees above USD 10,000.
  - 752 tx with no fee
  - 16,191 tx with fee less than USD 5.
- Huge variation in fees within the same block

Fees represent an increasing portion of miners' payoff

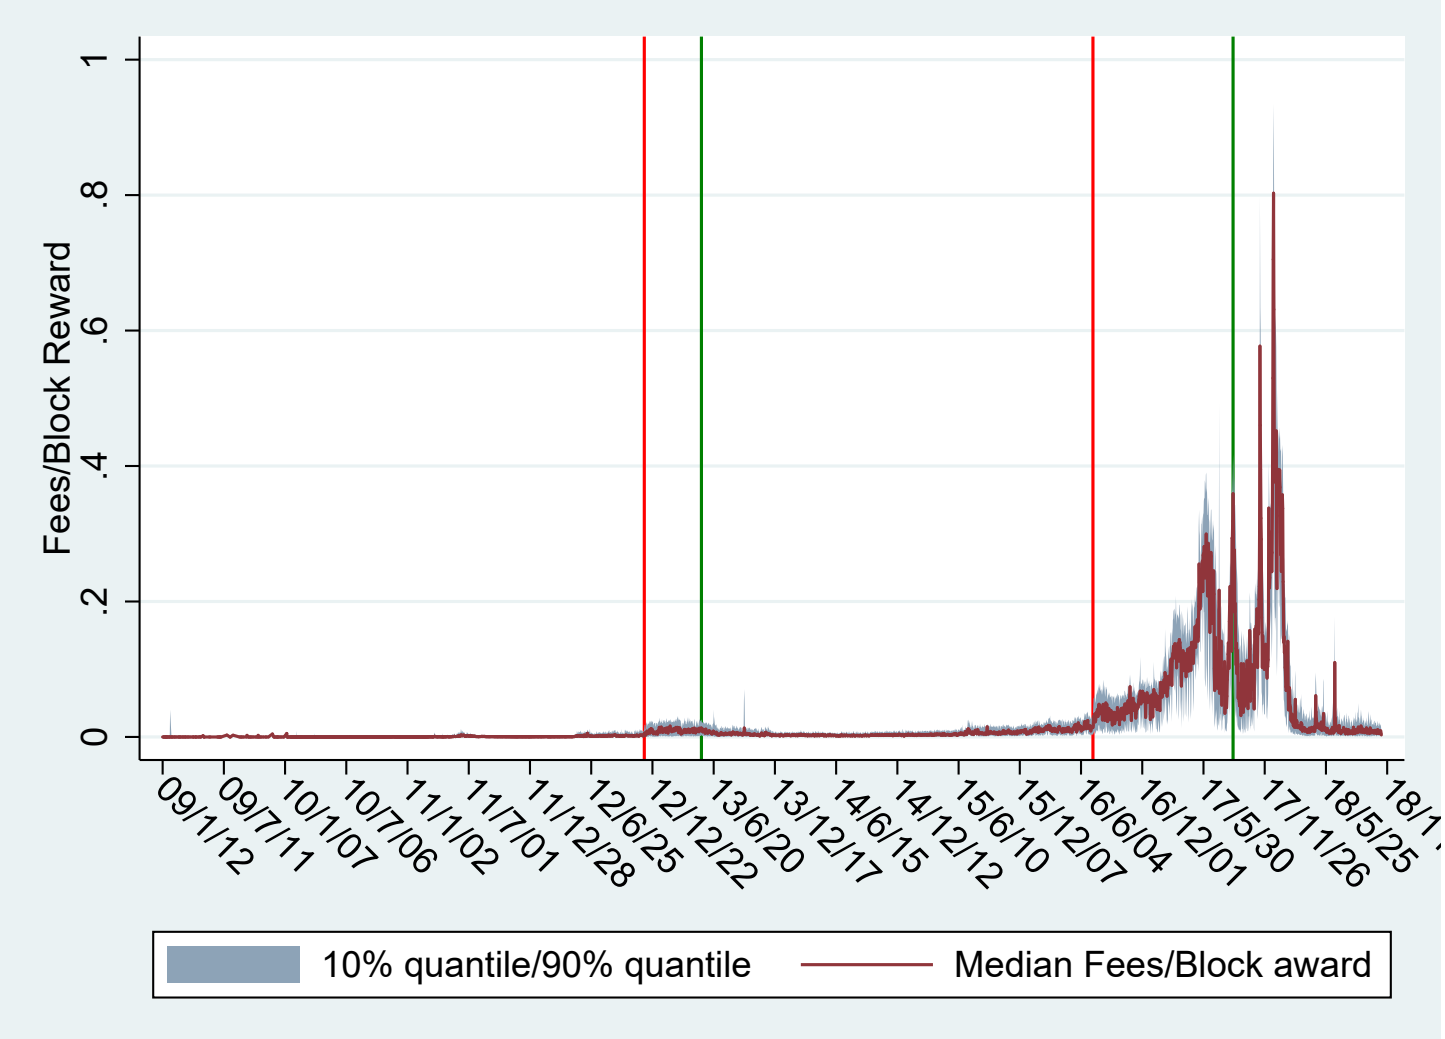

The red lines show when the block reward was cut in half, the green lines indicate when the block capacity was increased.

Fees as Bidding

- 2 types of agents, $j \in \{h, \ell\}$
- Type is (v_j, c_j), a value of settlement $v_h > v_\ell$ and a cost of waiting for the next block $c_h > c_\ell$
- All agents have the same outside option, $\bar{v}$

- Agents arrive at the market and submit a bid b_i,
- In reality, agents cannot revisit the market or change their bids, model this by assuming agents resubmit the same bid.
- They incur the cost for every block that their transaction remains unexecuted.

Sequence of Events within a block mining period

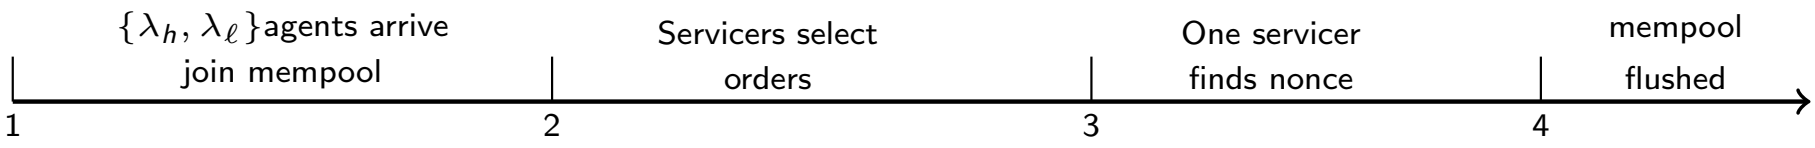

Payoff for an agent with type $\{c_i, v_i\}$, whose order is picked with probability p_i. Here t refers to the block height

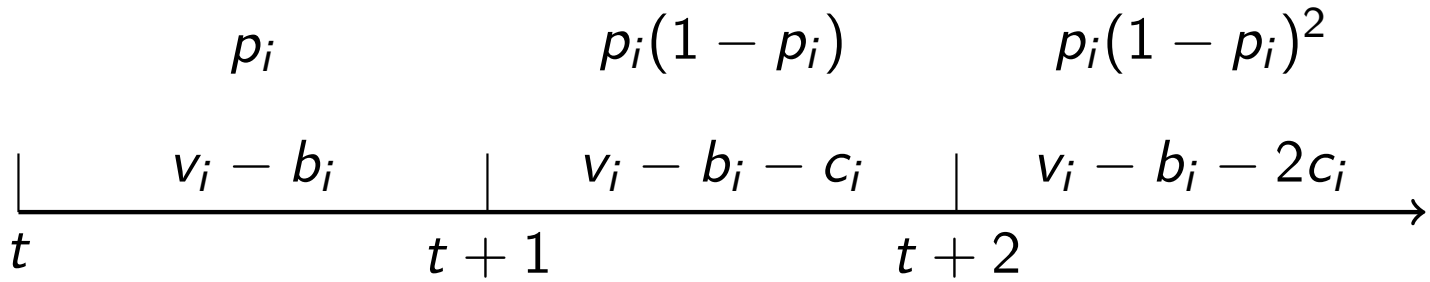

Agent's Utility

- The expected utility of an agent who submits a bid b_i and anticipates an execution probability p_i is

$$\begin{aligned} V(b_i \mid c_i) &= \sum_{k=1}^{\infty} p_i[1-p_i]^{k-1}(v_i - b_i - c_i(k-1)) \\ &= (v_i - b_i) - c_i \frac{1-p_i}{p_i}. \end{aligned}$$

- Clearly, increasing in p_i.

Different Rationing/Settlement Rules

- Each Servicer has discretion over the orders that he chooses to mine.
- Notice, this is different from a standard auction in which the rules are agreed on in advance.
  1. Full Capacity Servicing: The miners start with the highest bids and add more until the block is full.
  2. Discriminatory: The miners strategically manipulate execution probabilities so that bids are as high as possible.

- Because $v_h > v_\ell$ and $c_h > c_\ell$, construct an equilibrium in which all the type h agents execute with probability 1 and the type ℓ are rationed.
- Assume $\lambda_h < \kappa < \lambda_h + \lambda_\ell$.

Competitive Servicing

- If the servicers execute by price priority and exhaust all capacity then the high types retain rents.
- They bid just enough to beat the low types, or $b_h^c = v_\ell - \bar{v}$
  - Relatively few "desperate" agents arrive at any point in time.
  - To get processed they just have to bid slightly higher than the "less desperate" ones.

- Low types are driven down to their reservation utility.

Discriminatory Servicing: If the servicers execute by price priority and ration, then

i. Type h bids $b_h^d = v_h - \bar{v}$, and executes with probability $p_h^d = 1$.

ii. Type ℓ bids $b_\ell^d = v_\ell - \bar{v} - c_\ell\left(\frac{\Delta_v}{\Delta_c}\right)$ and executes with probability $p_\ell^d = \min\left[\frac{\Delta_c}{\Delta_v+\Delta_c}, \kappa - \lambda_h\right]$

iii. The profits of the servicer are $\pi^d = (v_h - \bar{v})\lambda_h + \frac{(v_\ell-\bar{v})\Delta_c - c_\ell\Delta_v}{\Delta_v+\Delta_c}\lambda_\ell$.

Where $\Delta_c = c_h - c_\ell$ and $\Delta_v = v_h - v_\ell$.

Discriminatory Servicing – Things to notice

- Servicers only pick a transaction for sure if the price is sufficiently high, else they leave it.
- "Desperate" agents end up bidding exactly what the transaction is worth to them.

- "Less desperate" have to wait sufficiently long, so that more desperate agents don't try to bid low.

Implications

1. The maximum bid observed under discriminatory pricing is higher than the maximum bid under capacity servicing.

2. In full capacity servicing, bids are increasing in capacity, whereas under discriminatory servicing, changes in capacity will have no effect on bids.

- Under full capacity servicing, the higher the capacity, the higher the utility agents get from using bitcoin.

Intertemporal Payoffs

- Suppose that there are $i = 1, \ldots N$ servicers of varying size and χ_i denote a servicer's hash rate as the proportion of the total Bitcoin system hash capacity.
- Payoff to servicer i under pricing regime $j = c, d$, a servicer obtains a per block profit at time t of π_t^j.
- Thus, under either regime $j = d, c$, the servicer anticipates from the next block onwards of

$$\Pi_i^j = \sum_{t=1}^{\infty} \delta^{t-1} \chi_i \pi_t^d \qquad j = d, c. \tag{1}$$

Sustaining Collusive Outcomes

- If a servicer sticks to discriminatory pricing then it knows that all other servicers will do the same.
- If a servicer decides to deviate from this strategy, he will mine as much as possible today but knows that in the future, all other miners will also mine at full capacity.
- Suppose that the per block profits from the discriminatory strategy is higher than the per block profits of full capacity mining.

Intertemporal Payoffs

- Payoff to cooperating in discriminatory pricing

$$\Pi_i^{coop} = (v_h - \bar{v})\lambda_h + \left((v_\ell - \bar{v}) - c_\ell \frac{1 - p_\ell^d}{p_\ell^d}\right)\lambda_\ell p_\ell^d + \delta \Pi_i^d$$

- Payoff to deviating and mining at capacity:

$$\Pi_i^{dev} = (v_h - \bar{v})\lambda_h + (\kappa - \lambda_h)\left((v_\ell - \bar{v}) - c_\ell \frac{1 - p_\ell^d}{p_\ell^d}\right) + \delta \Pi_i^c$$

Sustaining Collusion

$$\begin{aligned} \Pi_i^{coop} &\geq \Pi_i^{dev} \\ \delta(\Pi_i^d - \Pi_i^c) &\geq \left((v_\ell - \bar{v}) - c_\ell \frac{1 - p_\ell^d}{p_\ell^d}\right)\left(\kappa - \lambda_h - \lambda_\ell p_\ell^d\right) \end{aligned}$$

- Fewer participants is equivalent to a higher probability that a particular miner or mining pool finds the nonce and is successful.
- Future profits matter more if there are fewer participants:

Mining pools make collusive equilibria easier to sustain

Current Transaction Processors/Mining Pools

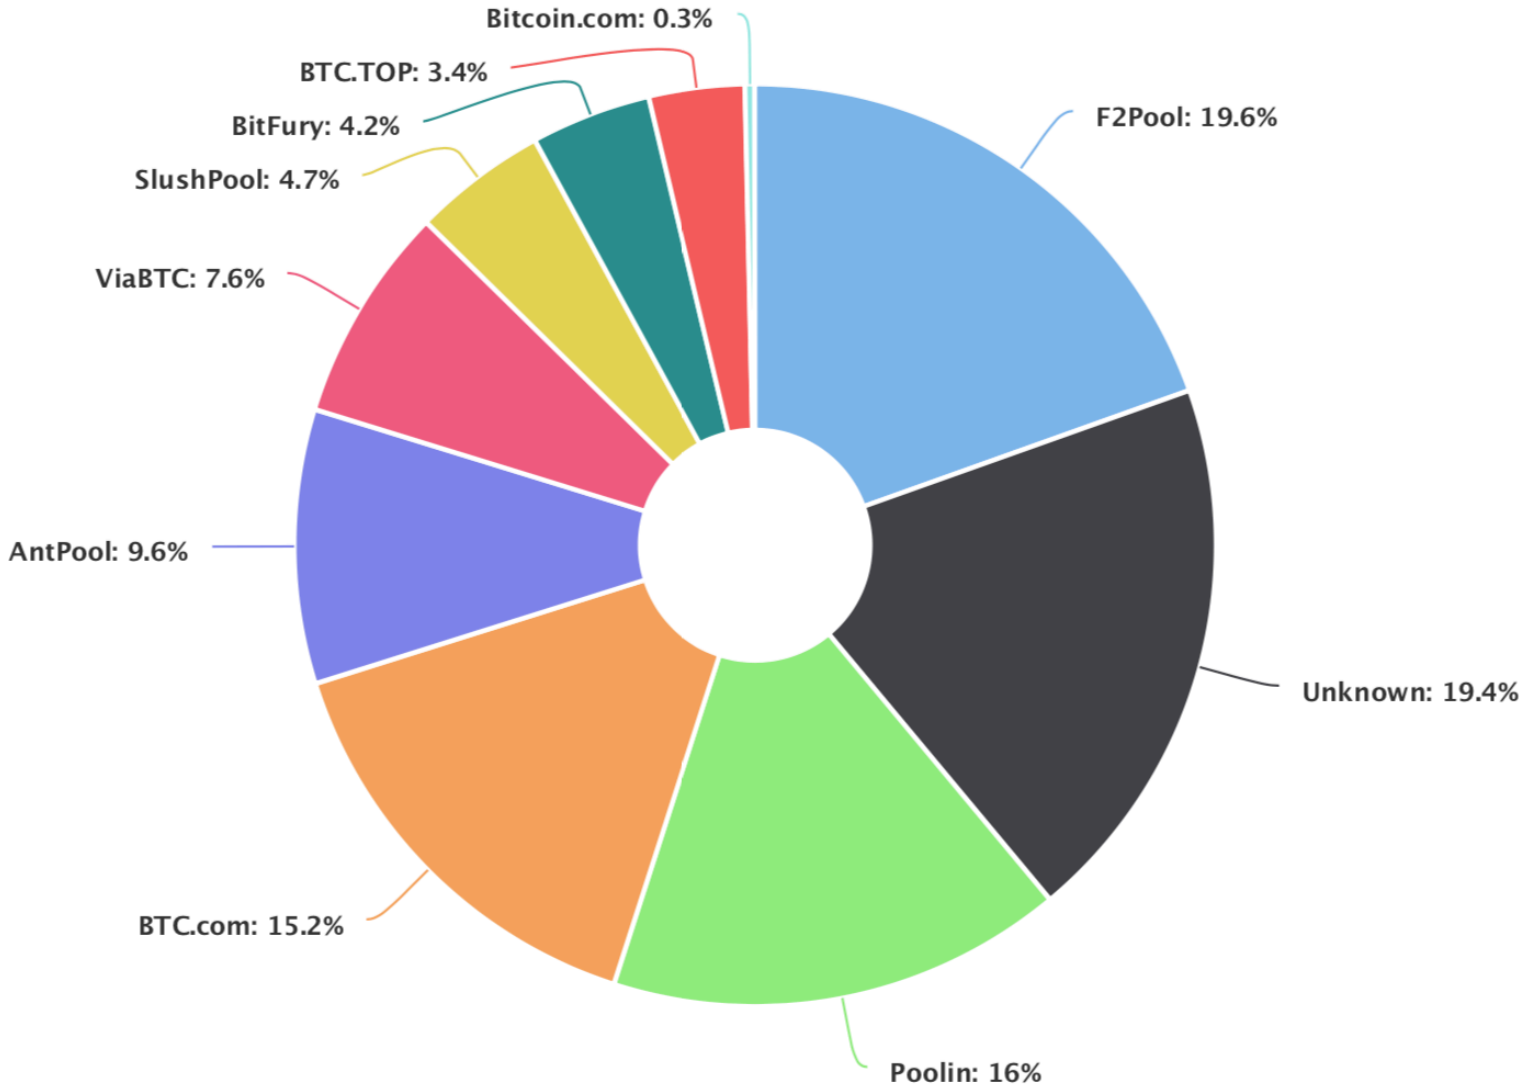

Source:BlockInfo.com

So Far

- Unused capacity can be used to price discriminate
- Mining pools can easily implement this with empty blocks.

- Note, it is does not matter if empty blocks arise for other reasons, because of the equilibrium effect they have on agents payoffs $\Longrightarrow$ increase costs.

- If successfully price discriminate $\Longrightarrow$ fees as proxy for willingness to pay for liquidity.

Time Series Evidence

- Collusion easier to sustain if market share of miners is high.
- Collect miners' signatures from each block's coinbase transaction.
- We compute the daily Hirschman Herfindahl index (HHI) of mining concentration as the sum of the squared shares of each mining pool computed over the day where the block is mined.
- We find that the feespread increases in both the HHI and the mining pool's aggregate share of mining activity.

HHI	3.1684***	3.1502***		
	(0.0373)	(0.0421)		
Agg Pool Share			0.8188***	1.1161***
			(0.0134)	(0.0163)
Mpool	0.0144***	0.0022	-0.0150***	-0.0115**
	(0.0047)	(0.0048)	(0.0049)	(0.0048)
Segwit	0.2025***	0.2943***	0.1859***	0.0899***
	(0.0027)	(0.0095)	(0.0027)	(0.0097)
Blocksize		-2.04e-09***		-2.02-07***
		(0.0000)		(0.0000)
TXSize		0.0000		0.0000
		(0.0000)		(0.0000)
Sum Inputs		-1.62e-07*		-1.51e-07*
		(0.0000)		(0.0000)
OP-Ret dummy		1.0195***		1.0220***
		(0.0459)		(0.0462)
Minimum Outtime		0.0002***		0.0002***
		(0.0000)		(0.0000)
Num blocks 1h		0.0117***		0.0119***
		(0.0011)		(0.0011)
Num Blocks 24h		-0.0005***		-0.0014***
		(0.0001)		(0.0001)
Capacity use 1h		0.2621***		0.2538***
		(0.0116)		(0.0119)
Capacity use 24h		-0.0323		-0.4090***
		(0.0212)		(0.0219)
constant	0.8455***	0.7578***	0.4995***	0.5419***
	(0.0057)	(0.0175)	(0.0117)	(0.0208)
R^2	0.0702	0.1095	0.0557	0.0992
N	198,111	198,105	198,111	198,105

Table: **Regression results of fee spread per block defined as the difference of the 90% and 10% quantile over the average fee. HHI is the Hirschman Herfindahl index of mining pool activity and Agg Pool Share is the share of mining done by mining pools. One, two, and three stars indicate significance at the 10%, 5%, and 1% level, respectively.**

Economic Costs of Collusion

- The total amount of fees paid in all transactions of our sample is USD 844,537,503.30.
- With no congestion, equilibrium fees should be zero.
- As a conservative estimate of maximum fees under competitive mining we approximate the fees that the lowest type would be willing to bid with the 10%-quantile of the fee distribution per block.
- We then define excessive fees as the sum of all fees above the 10%-quantile.

- For the whole sample these excessive fees sum to: USD 557,568,542.21

Conclusion

- BitCoin blockchain allows us to examine completely unregulated markets.
- Empirically, there is consistently underused capacity
- Relationship between fees and capacity consistent with price discrimination.
- Mining Pools have extracted close to $200 million in excess fees per year.