Cryptocurrencies, Mining & Mean Field Games

Ronnie Sircar

Department of Operations Research and Financial Engineering
Bendheim Center for Finance
Program in Applied & Computational Mathematics
Andlinger Center for Energy & the Environment
Princeton University

http://sircar.princeton.edu

Joint work with: Zongxi Li & Max Reppen

Bitcoin Mining

- Bitcoins, created Jan 2009, limited to 21 million: $>$ 17 million in circulation now.
- Independent "miners" compete for the right to record the next transaction block on the blockchain. They follow proof-of-work protocol and solve math puzzles.
- Once a miner obtains a solution, the corresponding block is added on top of the blockchain and the miner obtains the reward.
- The math puzzle is designed such that there is no known better way of solving it than brute force calculation: the chance of getting the reward is proportional to the computational power or the hash rates that miners can provide.
- The difficulty of the puzzle varies to maintain a consistent solving time, for example 10 minutes.

Context

We hope to better understand incentives for participants in a proof-of-work system.

Context

We hope to better understand incentives for participants in a proof-of-work system (like Bitcoin).

Context

We hope to better understand incentives for participants in a proof-of-work system (like Bitcoin).

Contributes to the growing area of cryptocurrency research.

Our work most closely relates to

- Arnosti and Weinberg 2018 consider a one-block asymmetric costs mining model and show that lower cost leads to higher market share.
- Alsabah and Capponi 2020 explore a two-stage mining game consisting of research and development, followed by competition.
- Cong, He, and Li 2019 examine mining pools and the impacts of their risk sharing.

Context

We hope to better understand incentives for participants in a proof-of-work system (like Bitcoin).

Contributes to the growing area of cryptocurrency research.

Our work most closely relates to

- Arnosti and Weinberg 2018 consider a one-block asymmetric costs mining model and show that lower cost leads to higher market share.
- Alsabah and Capponi 2020 explore a two-stage mining game consisting of research and development, followed by competition.
- Cong, He, and Li 2019 examine mining pools and the impacts of their risk sharing.

We propose a continuous time (repeated) mining game in which miners optimize terminal utility.
Our focus is on how factors such as competition, cost advantages, and resource endowment affect profit distribution, mining power, and centralization.

Results

- Agents with more resources have a stronger incentive to mine.
- This leads to preferential attachment and initial wealth imbalances are exacerbated.
- Under liquidity constraints, low-wealth miners are effectively blockaded by wealthier miners.
- Cost advantages in mining lead to significant shares of the mining market, unaffected by competition.
- These effects serve as explanations for the concentration of mining.

Cryptocurrency Issues

- What one hears most about cryptocurrencies, particularly Bitcoin, concerns the wildness of their prices, which seem to follow perpetual cycles of speculative mania and manic depression.
- This talk is not about their prices: we are interested in understanding and modeling the interaction of bitcoin miners and the consequent evolution of wealth inequality among participants.
- Are cryptocurrencies currencies or commodities? CFTC in the US classifies them as commodities, and their electronic structure of production mirrors the uncertainty and language of mining resources in finite supply.
- Connection to game theoretic models of energy production from various sources many of which, like oil, are exhaustible.

Cryptocurrency Issues II

- Much of the buzz around crypto comes from novice investors scoring a quick profit off an astounding price soar
- Data privacy concerns could be allayed by a payment and banking system founded on the underlying blockchain technology.
- A largely unregulated network could have myriad unintended benefits for trafficking and laundering.
- The hype may parallel that of the liberating internet a quarter century ago: anyone would be able to communicate whatever they want to everyone, and now of course we do just that.
- Time will tell the future of cryptocurrencies in society.

Mining model

Let $D = 1/10\text{min}$.

For a miner $i \in 1, \dots, M$ producing α_t^i hashes per $\mathrm{d}t$, the block rate of miner i is

$$\lambda_t^i = \frac{1}{D} \frac{\alpha_t^i}{\sum_{i=1}^M \alpha_t^i}.$$

Mining model

Let $D = 1/10\text{min}$.

For a miner $i \in 1, \ldots, M$ producing α_t^i hashes per dt, the block rate of miner i is

$$\lambda_t^i = \frac{1}{D} \frac{\alpha_t^i}{\sum_{i=1}^{M} \alpha_t^i}.$$

With costs c per unit of hash, the net reward is

$$-c\alpha_t^i \, \mathrm{d}t + r \, \mathrm{d}N_t^i,$$

where N_t is a process for which

$$P[N_{t+\Delta t}^i - N_t^i = 1] = \lambda_t^i \Delta t + o(\Delta t) \quad \text{and} \quad P[N_{t+\Delta t}^i - N_t^i \geq 2] = o(\Delta t).$$

Continuum approximation

For large M, this interaction becomes very complex.

Continuum approximation

For large M, this interaction becomes very complex.

Unfortunately, the interaction is not of mean-field structure:
If $M \to \infty$, either $\alpha_t^i \to 0$ or $\sum_{i=0}^{M} \alpha_t^i \to \infty$.

Continuum approximation

For large M, this interaction becomes very complex.

Unfortunately, the interaction is not of mean-field structure:
If $M \to \infty$, either $\alpha_t^i \to 0$ or $\sum_{i=0}^{M} \alpha_t^i \to \infty$.

However, observe that

$$\frac{\text{pl. } i\text{'s hash rate}}{\text{total hash rate}} = \frac{\text{pl. } i\text{'s hash rate}}{\#\text{players} \times \text{mean hash rate}} \approx \frac{\text{pl. } i\text{'s hash rate}}{\text{pl. } i\text{'s hash rate} + (\#\text{players} - 1) \times \text{mean hash rate}}.$$

Continuum approximation

For large M, this interaction becomes very complex.

Unfortunately, the interaction is not of mean-field structure:
If $M \to \infty$, either $\alpha_t^i \to 0$ or $\sum_{i=0}^{M} \alpha_t^i \to \infty$.

However, observe that

$$\frac{\text{pl. } i\text{'s hash rate}}{\text{total hash rate}} = \frac{\text{pl. } i\text{'s hash rate}}{\#\text{players} \times \text{mean hash rate}} \approx \frac{\text{pl. } i\text{'s hash rate}}{\text{pl. } i\text{'s hash rate} + (\#\text{players } -1) \times \text{mean hash rate}}.$$

Let

$$\lambda_t^i \overset{N=M}{=} \frac{1}{D} \frac{\alpha_t^i}{\alpha_t^i + M \frac{1}{N-1} \sum_{\substack{j=1,\\ j \neq i}}^{N} \alpha_t^j} \xrightarrow{N\to\infty} \frac{1}{D} \frac{\alpha_t^i}{\alpha_t^i + M\bar{\alpha}_t},$$

where $\bar{\alpha}$ is interpreted in the mean field games sense.

The miner's problem

Recall that the net earnings from mining at the rate α_t is

$$\mathrm{d}X_t = -c\alpha_t \,\mathrm{d}t + r\,\mathrm{d}N_t.$$

The miner's problem

Recall that the net earnings from mining at the rate α_t is

$$\mathrm{d}X_t = -c\alpha_t \,\mathrm{d}t + r\,\mathrm{d}N_t.$$

Each miner observes the aggregate mining of the population, $M\bar{\alpha}$, and seeks to maximize terminal utility:

$$v(t_0, x; \bar{\alpha}) = \sup_{\alpha} \mathbb{E}[U(X_T)|X_{t_0} = x],$$

The miner's problem

Recall that the net earnings from mining at the rate α_t is

$$\mathrm{d}X_t = -c\alpha_t \,\mathrm{d}t + r\,\mathrm{d}N_t.$$

Each miner observes the aggregate mining of the population, $M\bar{\alpha}$, and seeks to maximize terminal utility:

$$v(t_0, x; \bar{\alpha}) = \sup_{\alpha} \mathbb{E}[U(X_T) | X_{t_0} = x],$$

The HJB equation is

$$\partial_t v + \sup_{\alpha} \left(-c\alpha\partial_x v + \frac{\alpha}{D(\alpha + M\bar{\alpha}_t)} \Delta v \right) = 0, \quad v(T, x) = U(x),$$

where $\Delta v = v(t, x + r; \bar{\alpha}) - v(t, x; \bar{\alpha})$.

Optimal hash rate

The optimal response to the population mean hash rate $\bar{\alpha}$ is

$$\alpha^*(t, x; \bar{\alpha}) = \arg\max_{\alpha} \left(-c\alpha\partial_x v + \frac{\alpha}{D(\alpha + M\bar{\alpha}_t)} \Delta v \right)$$
$$= \begin{cases} -M\bar{\alpha}_t + \sqrt{\dfrac{M\bar{\alpha}_t \Delta v(t, x; \bar{\alpha})}{Dc\partial_x v(t, x; \bar{\alpha})}}, & \text{if} \quad \bar{\alpha}_t < \dfrac{\Delta v(t, x; \bar{\alpha})}{(M-1)Dc\partial_x v(t, x; \bar{\alpha})}, \\ 0, & \text{otherwise.} \end{cases}$$

Fokker–Planck equation

With α^* the optimal response to $\bar{\alpha}$, denote by $m(t, x; \bar{\alpha})$ the resulting density of miners.

Fokker–Planck equation

With α^* the optimal response to $\bar{\alpha}$, denote by $m(t, x; \bar{\alpha})$ the resulting density of miners.

Then m is a solution to

$$\partial_t m - \partial_x(c\alpha^* m) - \frac{1}{D}\left(\frac{\alpha^*(t, x-r)}{\alpha^*(t, x-r) + M\bar{\alpha}_t} m(t, x-r) - \frac{\alpha^*(t, x)}{\alpha^*(t, x) + M\bar{\alpha}_t} m(t, x)\right) = 0,$$

with initial distribution $m(t_0, x) = m_0(x)$.

Fokker–Planck equation

With α^* the optimal response to $\bar{\alpha}$, denote by $m(t, x; \bar{\alpha})$ the resulting density of miners.

Then m is a solution to

$$\partial_t m - \partial_x(c\alpha^* m) - \frac{1}{D}\left(\frac{\alpha^*(t, x-r)}{\alpha^*(t, x-r) + M\bar{\alpha}_t} m(t, x-r) - \frac{\alpha^*(t, x)}{\alpha^*(t, x) + M\bar{\alpha}_t} m(t, x)\right) = 0,$$

with initial distribution $m(t_0, x) = m_0(x)$.

This leads to a corresponding mean hash rate

$$\bar{\alpha}'_t = \int_{\mathbb{R}} \alpha^*(t, x; \bar{\alpha}\ )m(t, x; \bar{\alpha}\ )\,\mathrm{d}x, \quad \forall t \in [t_0, T].$$

Fokker–Planck equation

With α^* the optimal response to $\bar{\alpha}$, denote by $m(t, x; \bar{\alpha})$ the resulting density of miners.

Then m is a solution to

$$\partial_t m - \partial_x(c\alpha^* m) - \frac{1}{D}\left(\frac{\alpha^*(t, x-r)}{\alpha^*(t, x-r) + M\bar{\alpha}_t} m(t, x-r) - \frac{\alpha^*(t, x)}{\alpha^*(t, x) + M\bar{\alpha}_t} m(t, x)\right) = 0,$$

with initial distribution $m(t_0, x) = m_0(x)$.

This leads to a corresponding mean hash rate, and we call $\bar{\alpha}^*$ an *equilibrium mean hash rate* if

$$\bar{\alpha}^*_t = \int_{\mathbb{R}} \alpha^*(t, x; \bar{\alpha}^*) m(t, x; \bar{\alpha}^*) \, \mathrm{d}x, \quad \forall t \in [t_0, T].$$

Mean field game

We are looking for a solution to the coupled system

$$\begin{cases} 0 = \partial_t v + \sup_{\alpha} \left(-c\alpha\partial_x v + \dfrac{\alpha}{D(\alpha + M\bar{\alpha}_t^*)} \Delta v \right), \quad v(T, x) = U(x), \\ 0 = \partial_t m - \partial_x(c\alpha^* m) - \dfrac{1}{D} \left(\dfrac{\alpha^*(t, x-r)}{\alpha^*(t, x-r) + M\bar{\alpha}_t^*} m(t, x-r) - \dfrac{\alpha^*(t,x)}{\alpha^*(t,x) + M\bar{\alpha}_t^*} m(t,x) \right), \end{cases}$$

where α^* is the optimizer in the first equation and

$$\bar{\alpha}_t^* = \int_{E_t} \alpha^*(t,x) m(t,x) dx.$$

Solving the mean field game

Given any $\bar{\alpha}$, we can compute the following steps.

$$\bar{\alpha}$$

Solving the mean field game

Given any $\bar{\alpha}$, we can compute the following steps.

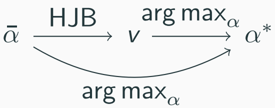

Solving the mean field game

Given any $\bar{\alpha}$, we can compute the following steps.

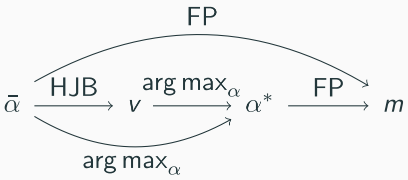

Solving the mean field game

Given any $\bar{\alpha}$, we can compute the following steps.

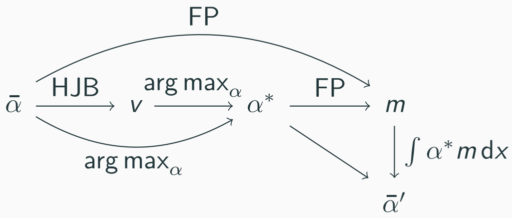

Solving the mean field game

Given any $\bar{\alpha}$, we can compute the following steps.

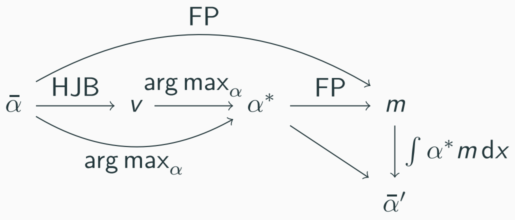

Idea: Start with $\bar{\alpha}$ and iterate this procedure $\bar{\alpha} \mapsto \bar{\alpha}'$ until convergence to $\bar{\alpha}^*$.

Solving the mean field game

Given any $\bar{\alpha}$, we can compute the following steps.

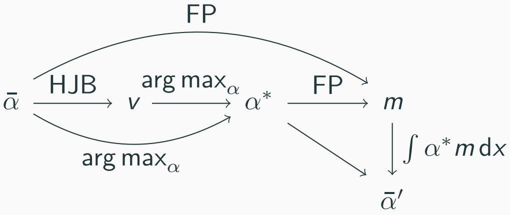

Idea: Start with $\bar{\alpha}$ and iterate this procedure $\bar{\alpha} \mapsto \bar{\alpha}'$ until convergence to $\bar{\alpha}^*$.

This typically fails due to oscillations, especially for large M.

Solving the mean field game

Given any $\bar{\alpha}$, we can compute the following steps.

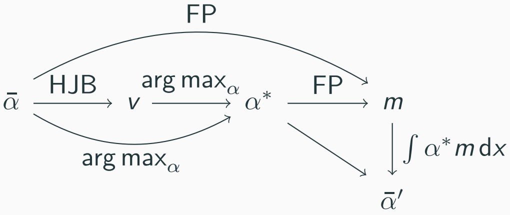

Idea: Start with $\bar{\alpha}$ and iterate this procedure $\bar{\alpha} \mapsto \bar{\alpha}'$ until convergence to $\bar{\alpha}^*$.

This typically fails due to oscillations, especially for large M.

To temper the oscillations, we introduce inertia:

$$\bar{\alpha} \mapsto \left(1 - \frac{1}{M}\right)\bar{\alpha} + \frac{1}{M}\bar{\alpha}'.$$

Liquidity constrained mining

Miners are required to keep their wealth non-negative.

A miner's activity ceases when its wealth reaches 0.

Let

$$U(x) = \frac{1}{1-\gamma} x^{1-\gamma} \quad \text{for } \gamma \in (0,1).$$

Liquidity constrained mining

Miners are required to keep their wealth non-negative.

A miner's activity ceases when its wealth reaches 0.

Let

$$U(x) = \frac{1}{1-\gamma} x^{1-\gamma} \quad \text{for } \gamma \in (0,1).$$

Lemma

For any time t and equilibrium hash rate $\bar{\alpha}_t^* > 0$, there exists $x_b(t) > 0$ such that zero rate mining is optimal, i.e., $\alpha^*(t,x) = 0$ for $x \leq x_b(t)$.

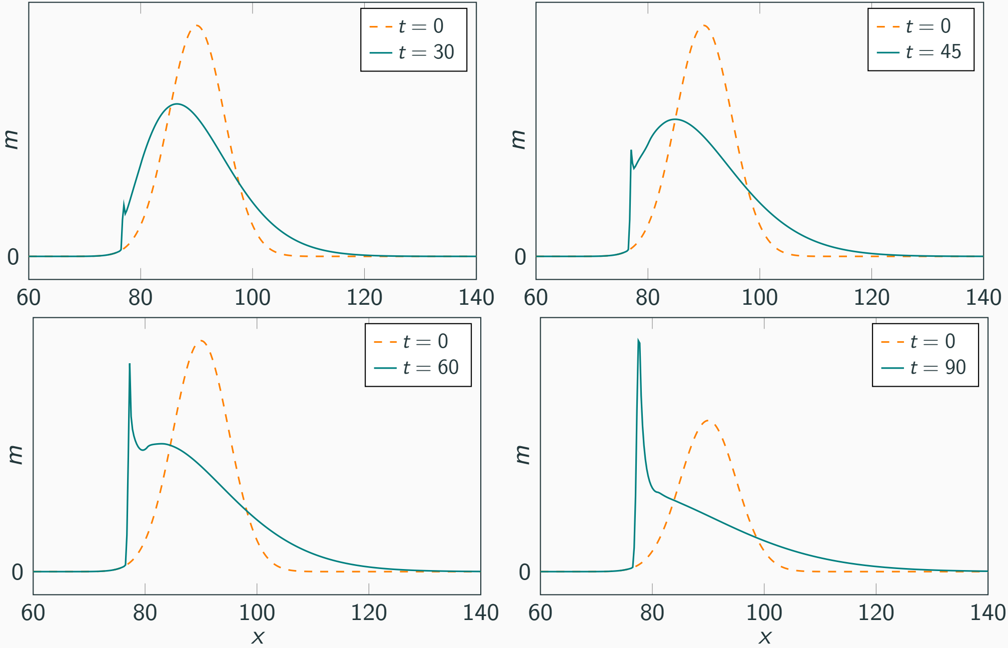

t = 0
t = 30
t = 0
t = 45
t = 0
t = 60
t = 0
t = 90
m
0
60
80
100
120
140
x

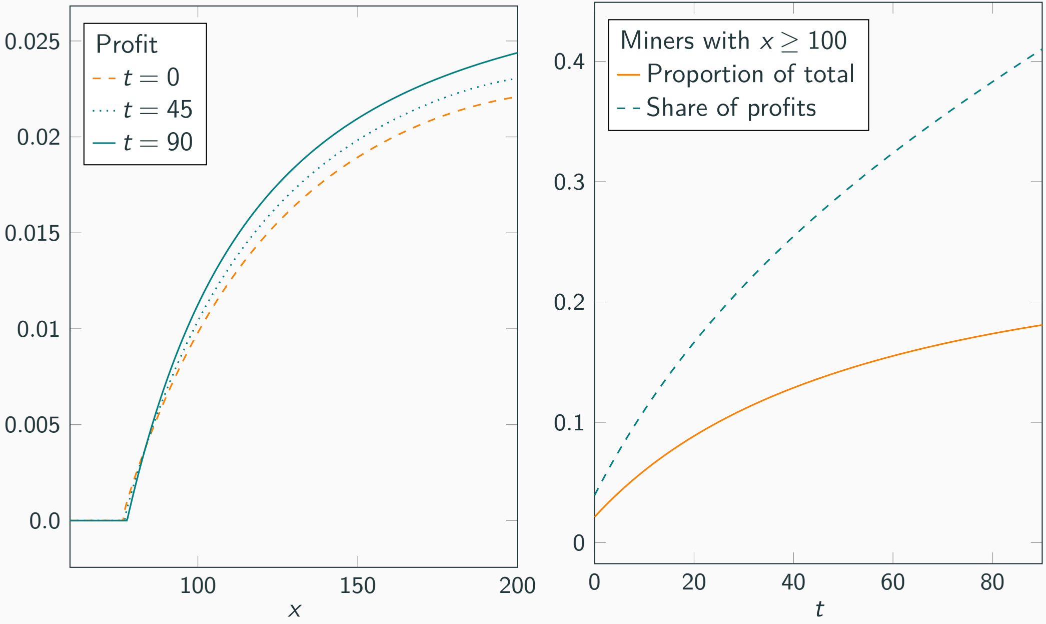

Profit
$t = 0$
$t = 45$
$t = 90$
0.025
0.02
0.015
0.01
0.005
0.0
100
150
200
x
Miners with $x \geq 100$
Proportion of total
Share of profits
0.4
0.3
0.2
0.1
0
0
20
40
60
80
t

Cost-advantaged miner

Cost-advantaged miner hashes at the rate β, pays $k_c c \beta_t$ for $k_c > 0$, and mines blocks at the rate

$$\lambda_t^1 = \frac{\beta_t}{D(\beta_t + M\bar{\alpha}_t)}.$$

For simplicity, this miner is assumed wealthy and risk-neutral:

$$\sup_{\beta_t \geq 0} \mathbb{E}\left[\int_0^T -k_c c \beta_t dt + p dN_t^1\right].$$

The optimal hash rate is

$$\beta^*(t; \bar{\alpha}) = \begin{cases} -M\bar{\alpha}_t + \sqrt{\dfrac{pM\bar{\alpha}_t}{k_c c D}}, & \text{if} \quad \bar{\alpha}_t < \dfrac{p}{k_c c M D}, \\ 0, & \text{otherwise.} \end{cases}$$

The other miners

Taking into account the advantaged miner,

$$\lambda_t = \frac{\alpha_t}{D(\alpha_t + (M-1)\bar{\alpha}_t + \beta_t)},$$

New maximizer

$$\alpha^*(t, x; \bar{\alpha}, \beta) = \begin{cases} -(M\bar{\alpha}_t + \beta_t) + \sqrt{\dfrac{(M\bar{\alpha}_t + \beta_t)\Delta v}{Dc\partial_x v}}, & \text{if} \quad M\bar{\alpha}_t + \beta_t < \dfrac{\Delta v}{Dc\partial_x v}, \\ 0, & \text{otherwise.} \end{cases}$$

Explicit solution

Proposition

Suppose the individual miners have exponential utility $U(x) = -\frac{1}{\gamma}e^{-\gamma x}$ *and no liquidity constraints, suppose the relative cost efficiency satisfies*

$$k_c < \frac{\gamma r}{1 - e^{-\gamma r}} \frac{M}{M - 1},$$

and let

$$\kappa_1 = \frac{1 - e^{-\gamma r}}{Dc\gamma}, \quad \kappa_2 = \frac{Mr}{Dk_c c}.$$

Then, in equilibrium, all miners are active with

$$\alpha^*(t, x) \equiv \bar{\alpha}_t^* \equiv \frac{\kappa_1^2 \kappa_2}{(\kappa_1 + \kappa_2)^2} > 0, \quad \beta_t^* \equiv \frac{\kappa_1 \kappa_2 (\kappa_2 - M\kappa_1)}{(\kappa_1 + \kappa_2)^2} > 0,$$

for all $t \in [t_0, T]$ *and* $x \in \mathbb{R}$.

Market share and profits

For $\gamma \ll p$ and $M \gg 1$,

$$\frac{\beta^*}{\beta^* + M\bar{\alpha}^*} \approx 1 - k_c.$$

In other words, mining competition does not affect the market share of the advantaged miner.

The hash rates β^*, α^* and profits

$$Y^1_{t_0+t} = -k_c c \beta^* t + r N^{1*}_t, \quad Y_{t_0+t} = -c\alpha^* t + r N^*_t$$

satisfy, as $M \to \infty$,

$$\begin{aligned} \beta^*_t &= \mathcal{O}(1) & \alpha^*(t,x) &= \mathcal{O}(1/M) \\ \mathbb{E}(Y^1_{t_0+t}) &= \mathcal{O}(1), & \mathbb{E}(Y_{t_0+t}) &= \mathcal{O}(1/M), \\ \mathrm{Var}(Y^1_{t_0+t}) &= \mathcal{O}(1), & \mathrm{Var}(Y_{t_0+t}) &= \mathcal{O}(1/M). \end{aligned}$$

Liquidity constraints

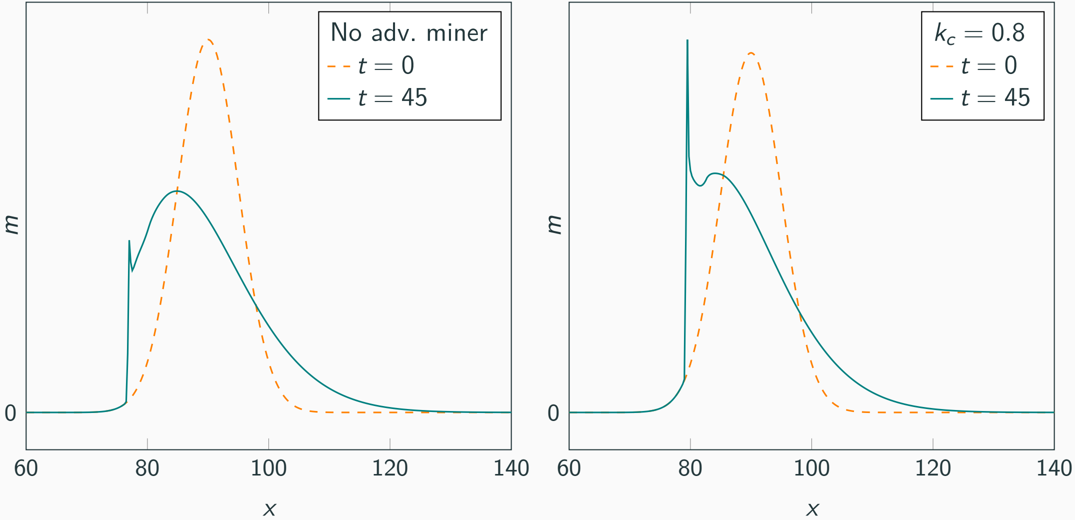

Liquidity constraints

Conclusion

Conclusion

- Concentration of Bitcoin mining enables censorship of transactions.

Conclusion

- Concentration of Bitcoin mining enables censorship of transactions.
- The strong incentives for concentration is fundamental to proof-of-work.

Conclusion

- Concentration of Bitcoin mining enables censorship of transactions.
- The strong incentives for concentration is fundamental to proof-of-work.
- This suggests that the current state of aggregation into a small number of mining pools is not transient.

Conclusion

- Concentration of Bitcoin mining enables censorship of transactions.
- The strong incentives for concentration is fundamental to proof-of-work.
- This suggests that the current state of aggregation into a small number of mining pools is not transient.
- Not only is centralization at odds with the core priciples of Bitcoin; it is also a danger to the system as it enables censorship and possibly fraudulent transactions.