The Complex Dynamics of Real World Markets

Georgios Piliouras

Traditional mechanism design

- Aims to design institutions or systems that achieve desired outcomes despite the self-interest of participating individuals.

Desired outcomes: Efficient allocation of valuable resources (e.g. allocation of blockchain space or ad slots, etc.)

Key tools: Implementation via a strongly stable game theoretic equilibrium concept (e.g. dominant strategy incentive compatibility)

+ **Conceptually simple and clean**
- **Not always perfectly applicable in practice**

The Dynamics of

i) Ethereum's Transaction Fee Market

ii) Google's Auto-bidding Ad-Auctions

The Dynamics of Ethereum's Transaction Fee Market

Leonardos, Stefanos, Barnabé Monnot, Daniël Reijsbergen, Efstratios Skoulakis, and Georgios Piliouras. "Dynamical analysis of the EIP-1559 ethereum fee market." *ACM Conference on Advances in Financial Technologies*, 2021.

Reijsbergen, Daniël, Shyam Sridhar, Barnabé Monnot, Stefanos Leonardos, Stratis Skoulakis, and Georgios Piliouras. "Transaction fees on a honeymoon: Ethereum's EIP-1559 one month later." *IEEE International Conference on Blockchain (Blockchain)*, 2021.

Leonardos, Stefanos, Daniël Reijsbergen, Barnabé Monnot, and Georgios Piliouras. "Optimality despite chaos in fee markets." In *International Conference on Financial Cryptography and Data Security*, 2023.

Prior to EIP-1559 Fee Market: Generalized First Price Auction

- Strategic bidding
- Volatile transaction fees
- Transaction delays

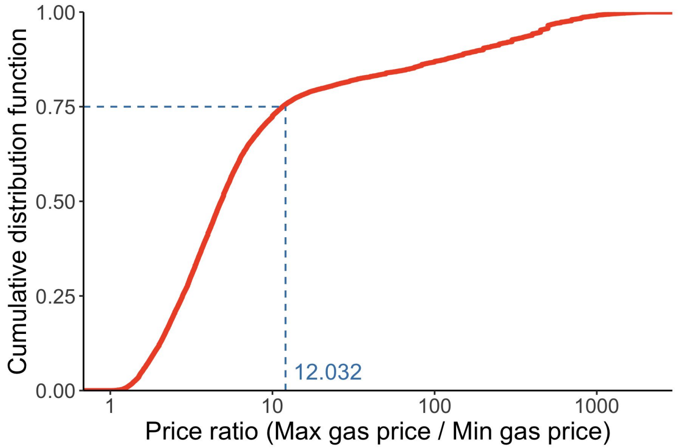

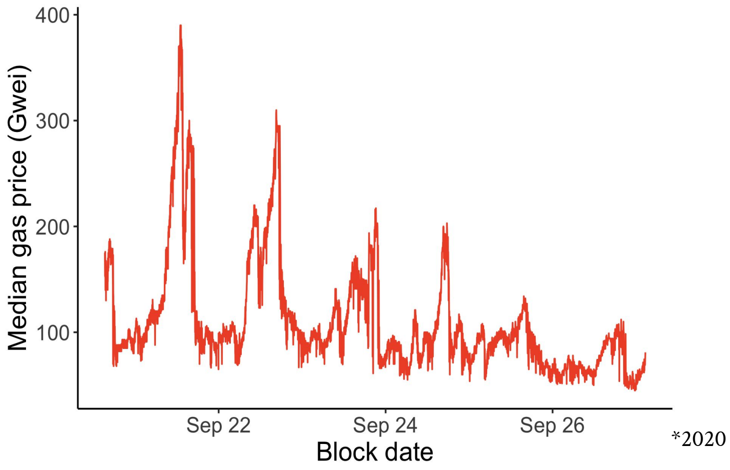

- Second-price auctions will not work: they are prone to collusions and gaming by miners (non-credible)

Goals of EIP 1559

- Price discovery for the transaction fee (quickly match demand)
- Incentive-compatible for both miners and users (under "most conditions")*
- Price-taking behavior (transparency, efficiency): next-block inclusion

Not a goal: always the same fee

Not a goal: lower fees (but yes, lower variance of fees)

*T. Roughgarden, *Transaction Fee Mechanism Design*, EC'21

EIP-1559

Protocol parameters

- target block size: T/2
- base fee: a dynamical gas price which is **burnt** and aims to control congestion

User bids

- max fee: the highest fee the user is willing to pay overall
- max priority fee: the maximum tip the user is willing to pay to the miner

Protocol	User bids	Miners	Transaction
T/2 = target block size b = base fee	f = max fee p = max priority fee	e = min cost	✅ or ❎

Transaction Fees

- Base fee is **burnt**
- user pays transactions fees at most equal to max fee
- miner's tip at most the max priority fee set by the user

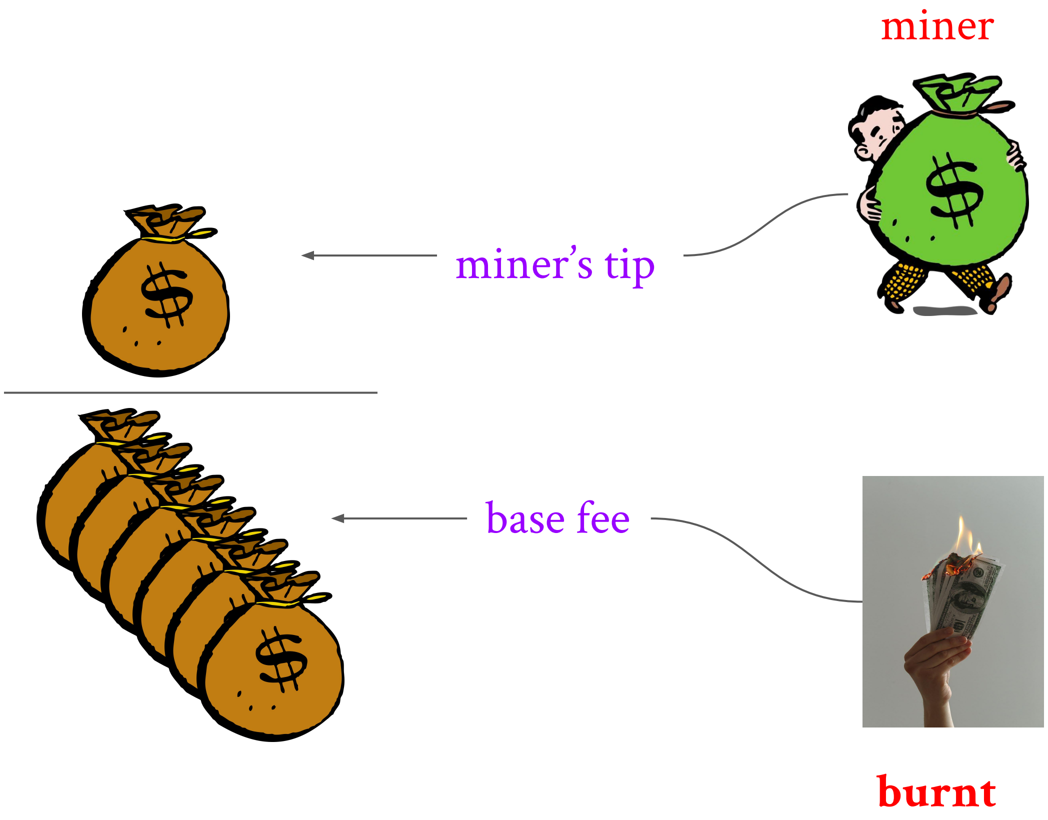

Fees and Tips

- Examples

Protocol	User	Miner	Transaction	Why?
b = 6	f = 10, p = 3	e = 2	✅ included	f>b+e, p>e
b = 6	f = 10, p = 1	e = 2	❎ not included	f>b+e but p<e
b = 6	f = 7, p = 3	e = 2	❎ not included	p>e but f<b+e

- Condition for inclusion by miner:

$$\min(f - b_t, p) > e$$

- Left side: miner's tip

EIP-1559 Base fee Dynamics

Base fee

The base fee, b_{t+1}, at block height t+1 is

$$b_{t+1} = b_t \left(1 + d \times \frac{g_t - T/2}{T/2}\right)$$

- d = adjustment parameter or step size (think d = 12.5%)
- T/2 = target block size (think T = 952)
- g_t = number of included transactions in block B_t, t>0 given that the base fee is b_t.

Goals of this work

- Under which conditions do the base fee dynamics self-stabilize?
- What happens outside the stability regime?
- Stress-test the system by pushing it past its stability regime.

EIP-1559: Evolution of the Dynamics

User demand described by

- Stationary distribution of **valuations:** each user has a valuation
- Optimal **base fee:** b^*

Will EIP-1559 **find/converge to** b^* for different design choices?

Oscillations

- Example: 3000 users, equally spaced valuations in 200-230 Gwei.
- Many users with similar valuations: base fee jumps (by a step of d) from below to above

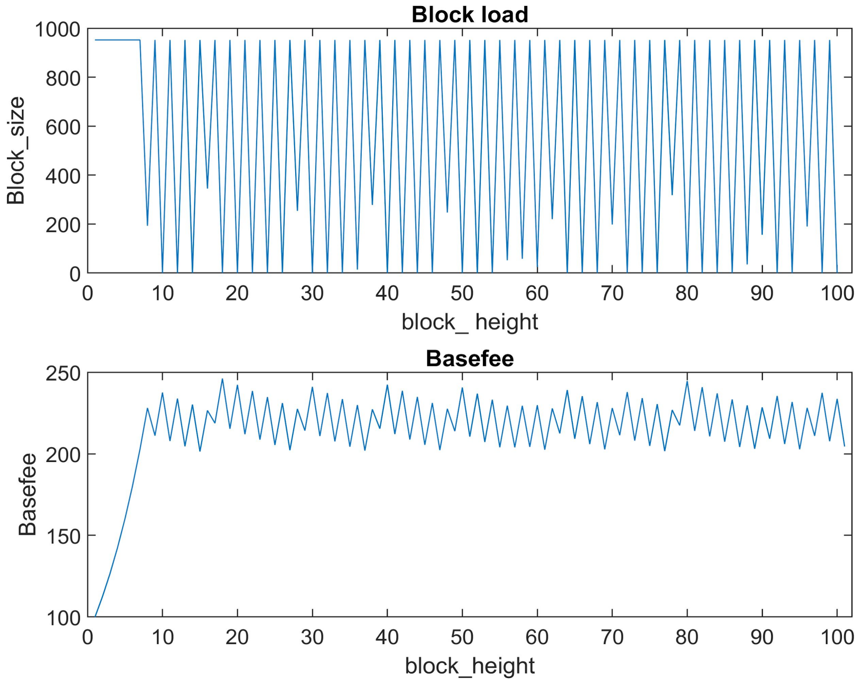

Convergence

Theorem (Informal)
The base fee dynamics always converge to a theoretically predictable neighborhood of the "correct" value. Moreover, for any reasonable distribution of user valuations, there exists a low enough adjustment parameter (step-size), for which the base fee dynamics precisely converge to that (theoretically predictable) "correct" value.

Proof Hint: **Lyapunov function** arguments

One meta-theorem for all convergence results

[Losert, Akin 1983]

Theorem

Suppose a continuous-time dynamical system obtained from a differential equation on compact space X (or a discrete time dynamical system obtained by iterating a continuous map $F:X\to X$) admits a ***Lyapunov function*** *$L:X\to$, i.e.,*

$$dL/dt \leqslant 0 \quad (\text{resp. } L(F(p))\leqslant L(p))$$

with equality at p only when p is an equilibrium.

THEN

the ***limit set*** *of an orbit $\{p(t)\}$ is a compact connected set consisting entirely of* ***equilibria*** *and upon which the Lyapunov function L is constant.*

Convergence

Theorem (Informal)
The base fee dynamics always converge to a theoretically predictable neighborhood of the "correct" value. Moreover, for any reasonable distribution of user valuations, there exists a low enough adjustment parameter (step-size), for which the base fee dynamics precisely converge to that (theoretically predictable) "correct" value.

Proof Hint: **Lyapunov function** arguments

- "Low enough" depends on the distribution of valuations.
- A low enough step size may be impractical.

This raises the question: what happens if this is not the case?

Chaos and Instabilities

Theorem (Informal)
For every positive adjustment parameter, there exist a (reasonable) distribution of valuations, so that the base fee dynamics become Li-Yorke chaotic.

Li-Yorke chaos:

- Uncountably many pairs of trajectories get arbitrarily close together (but never intersect) and move apart indefinitely.
- We cannot tell which of the two trajectories will be realized in the future: this is exactly what unpredictable means.

Definition of Li-Yorke Chaos

Scrambled set: Given a dynamical system with update rule f, a pair of points x and y is called scrambled if the trajectories get arbitrarily close, i.e.,

$$\lim_{n\to\infty} \inf |f^n(x) - f^n(y)| = 0$$

and then the trajectories move apart infinitely often

$$\lim_{n\to\infty} \sup |f^n(x) - f^n(y)| > 0$$

A set S is called scrambled if **all pairs of points** in S are **scrambled**.

Li-Yorke chaos: A discrete time dynamical system is called chaotic if (a) for each natural number k>0 there exists a periodic point of period k and (b) there is an uncountably infinite set that is scrambled.

How to prove chaos?

- As long as the dynamical system is continuous, discrete-time and 1-dimensional, i.e. the system has only one degree of freedom then there exist several efficiently checkable conditions that imply Li-Yorke chaos.

Specifically, **period 3 implies chaos**. [Li, Yorke 1975]

- Prototypical example:
  Logistic map f(x) = rx(1-x), r constant, 0<x<1

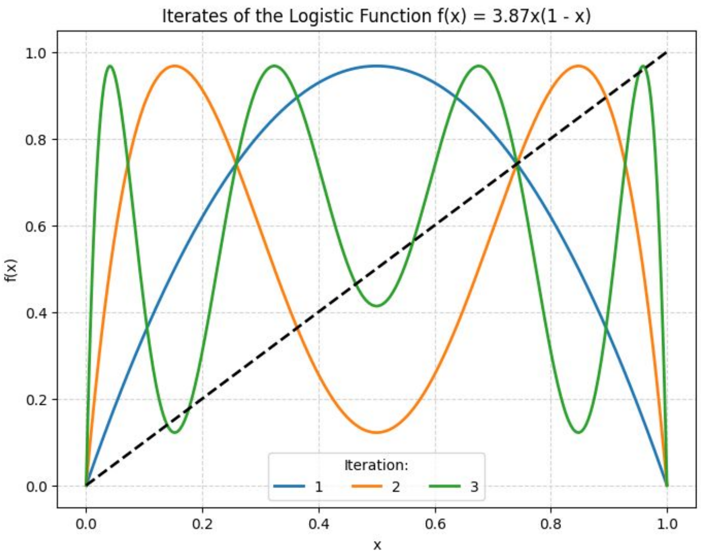

How to showcase chaos?

Logistic map F(x)= r x(1-x)

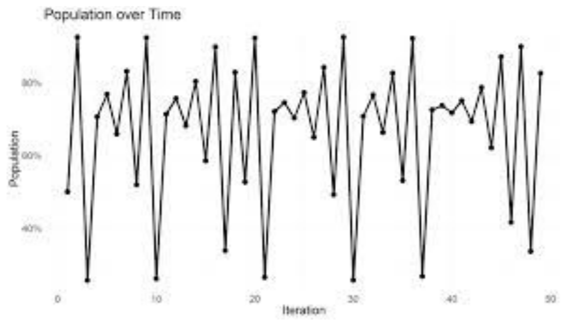

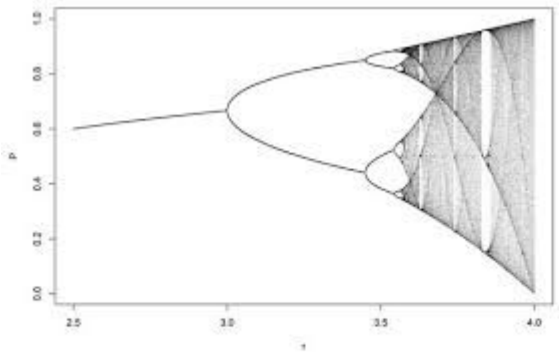

Time-series plots

Bifurcation plots

How to showcase chaos?
EIP-1559

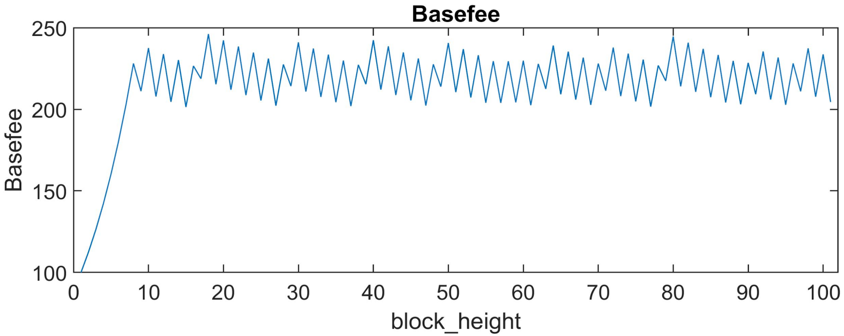

Time-series plots

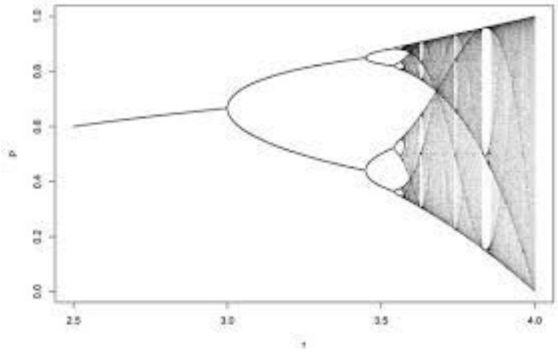

Bifurcation plots

Bifurcation plots

- Bifurcation plot with respect to the adjustment parameter (learning rate)
- Uniform distribution of valuations with range = 20Gwei.
- As the step size increases, the base fee becomes chaotic

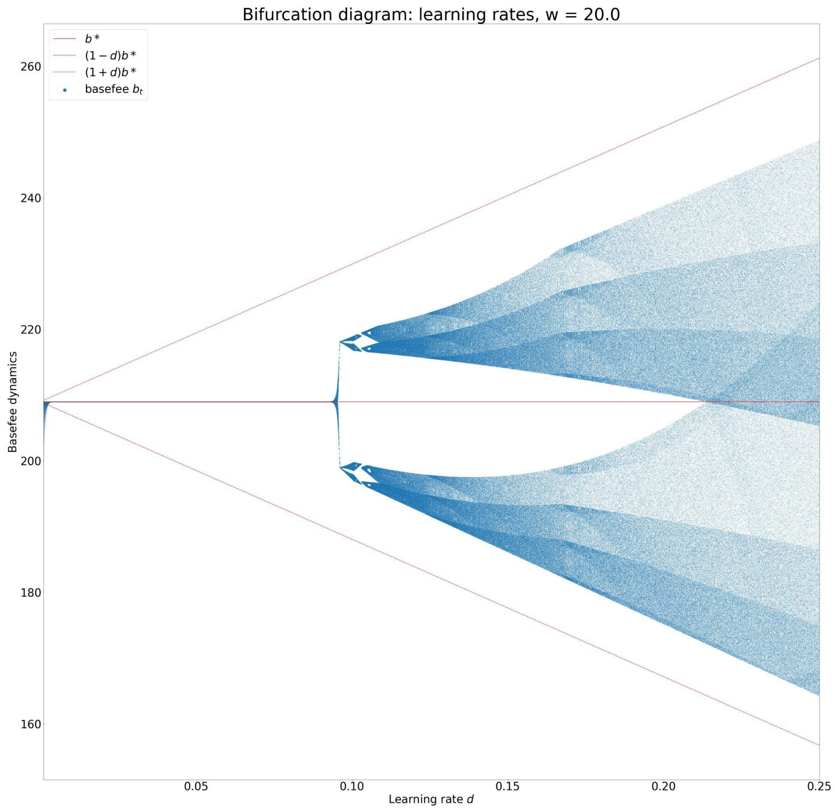

Bifurcation plots

- Bifurcation plot with respect to the range of valuations
- Step size d = 0.125 (default).
- As the range of valuations decreases, the base fee becomes chaotic

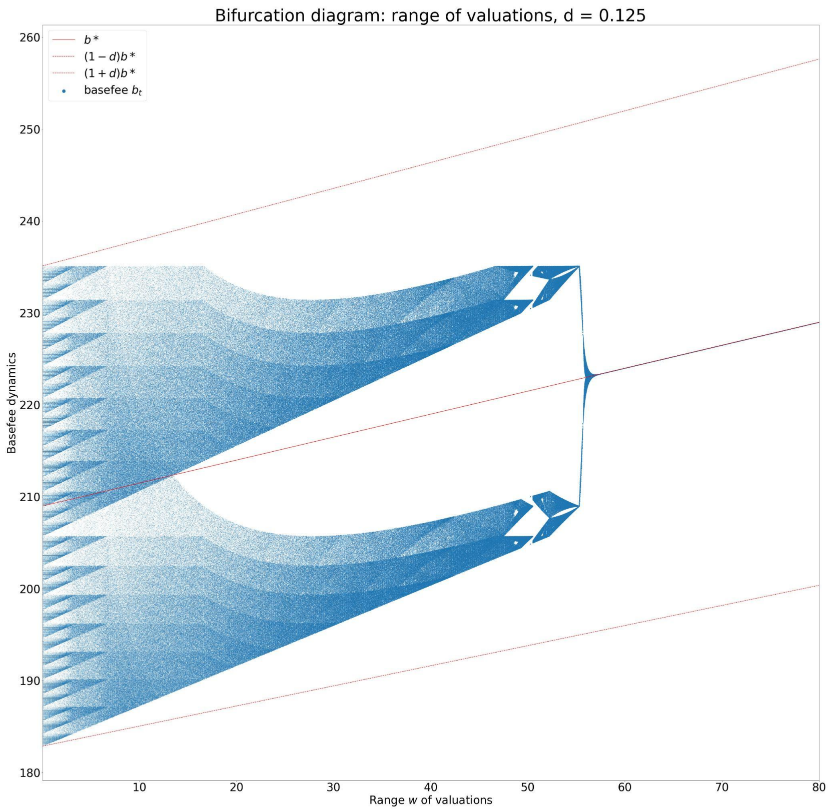

EIP-1559: (Time-average) Performance

EIP-1559: Evolution of the Dynamics

User demand described by

- Stationary distribution of **valuations:** each user has a valuation
- Optimal **base fee:** ***b****

We have seen that the day-to-day dynamics can be **chaotic.**

Will EIP-1559 **find/converge to *b**** ***in a time average sense*** for different design choices?

Time averages

- Step size d = 0.125 (default): skip = 400, iterations = 600.
  - Average base fee
  - Average block size

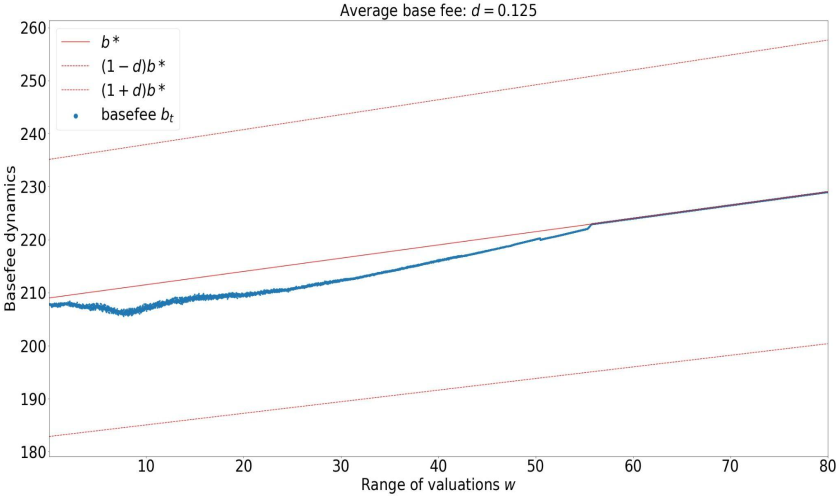

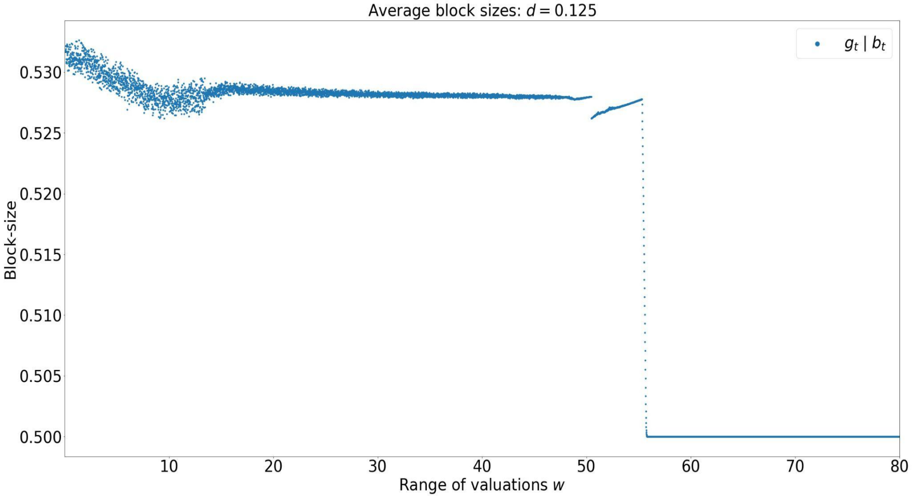

EIP-1559: Chaos and Optimality

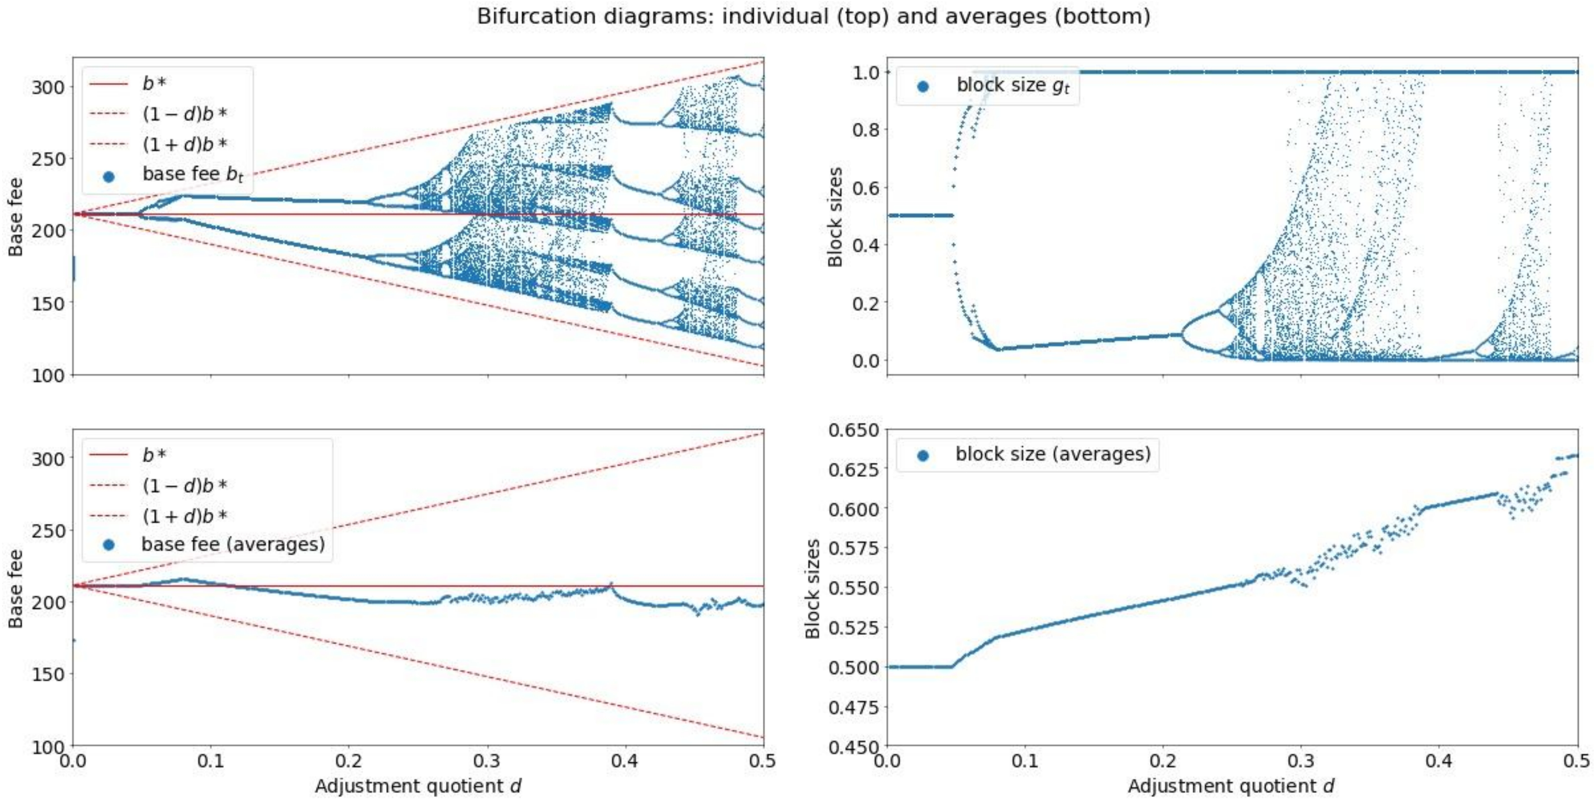

Time-average Performance

Theorem (Informal)
Regardless of the day-to-day behavior of EIP-dynamics, the time average utilization will converge to a narrow range around its optimal/target value.

- Our **lower bound** is equal to the **target utilization** level whereas our **upper bound** is approximately **6% higher** than optimal.
- Empirical evidence is shown in great agreement with these theoretical predictions. Specifically, the historical average was approximately 2.9% larger than the target rage under Proof-of-Work and decreased to approximately 2.0% after Ethereum's transition to Proof-of-Stake.

Time-average Performance

Theorem (Informal)
Regardless of the day-to-day behavior of EIP-dynamics, the time average utilization will converge to a narrow range around its optimal/target value.

Proof Hint 1: This is not a regret-based, Price of (Total) Anarchy type of analysis.

Proof Hint 2: It easier to prove this theorem for an (exponential variant) of EIP-1559.

EIP-1559: Average Block Size (Theoretical Bound & Simulations)

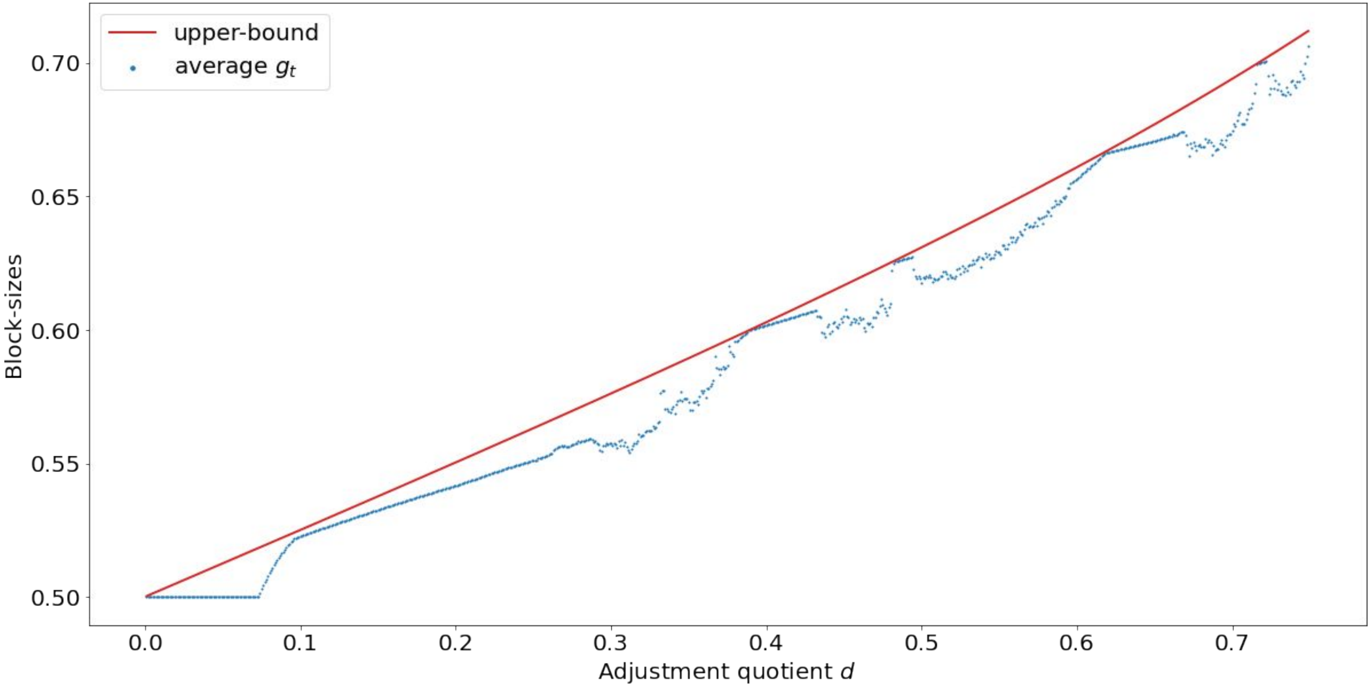

EIP-1559: Market Data

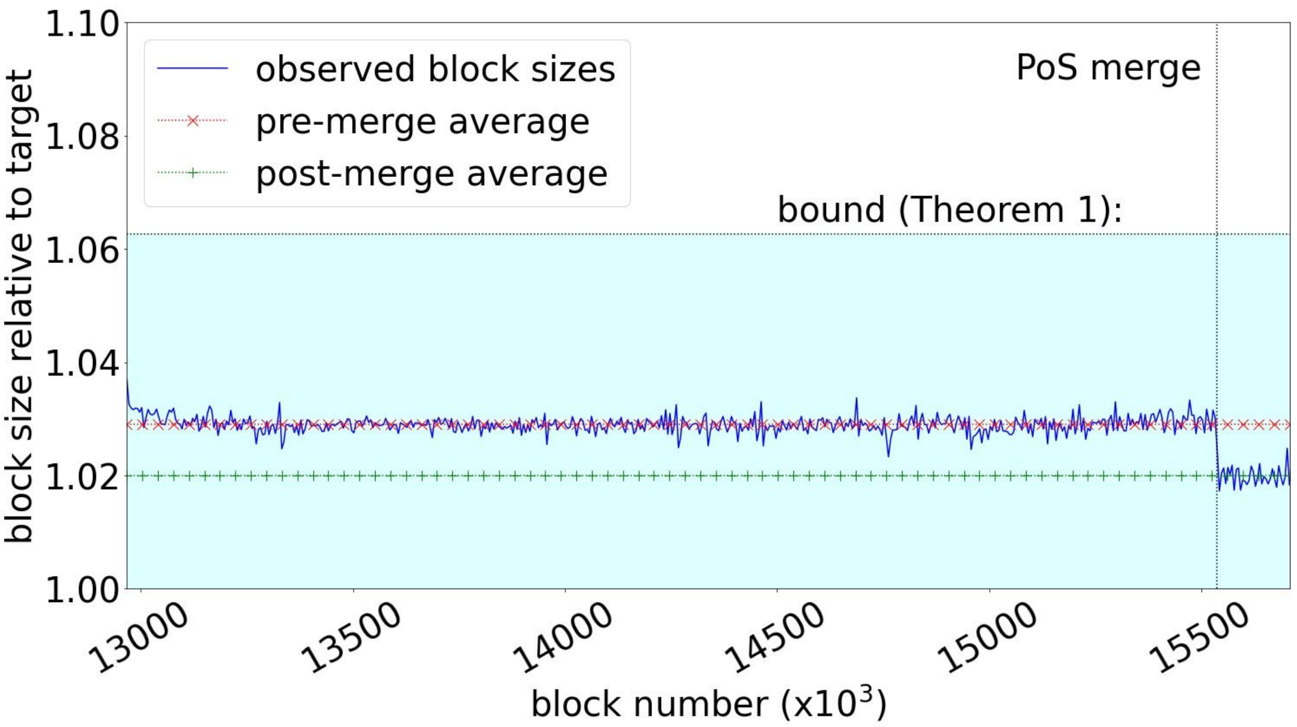

Discussion

- EIP-1559 is **simple** and makes transaction fees more ***predictable***
- Unintended consequences:
  - **Oscillations**
  - **Overshoots** block size target
- Both have lessened since the Merge
- **Optimality** is approximately possible in fee markets with **non-convergent** behavior
- We can still analyze the system even out of **equilibrium**

The Dynamics of

i) Ethereum's Transaction Fee Market

ii) Google's Auto-bidding Ad-Auctions

The Dynamics of Google's Auto-bidding Ad-Auctions

Renato Paes Leme, Georgios Piliouras, Jon Schneider, Kelly Spendlove, Song Zuo
Complex Dynamics in Auto-bidding Systems
ACM Conference on Economics and Computation (2024).

Gagan Aggarwal et al.
Auto-bidding and Auctions in Online Advertising: A Survey
ACM SIGecom Exchanges, 22 (2024)

Autobidding Auction Overview

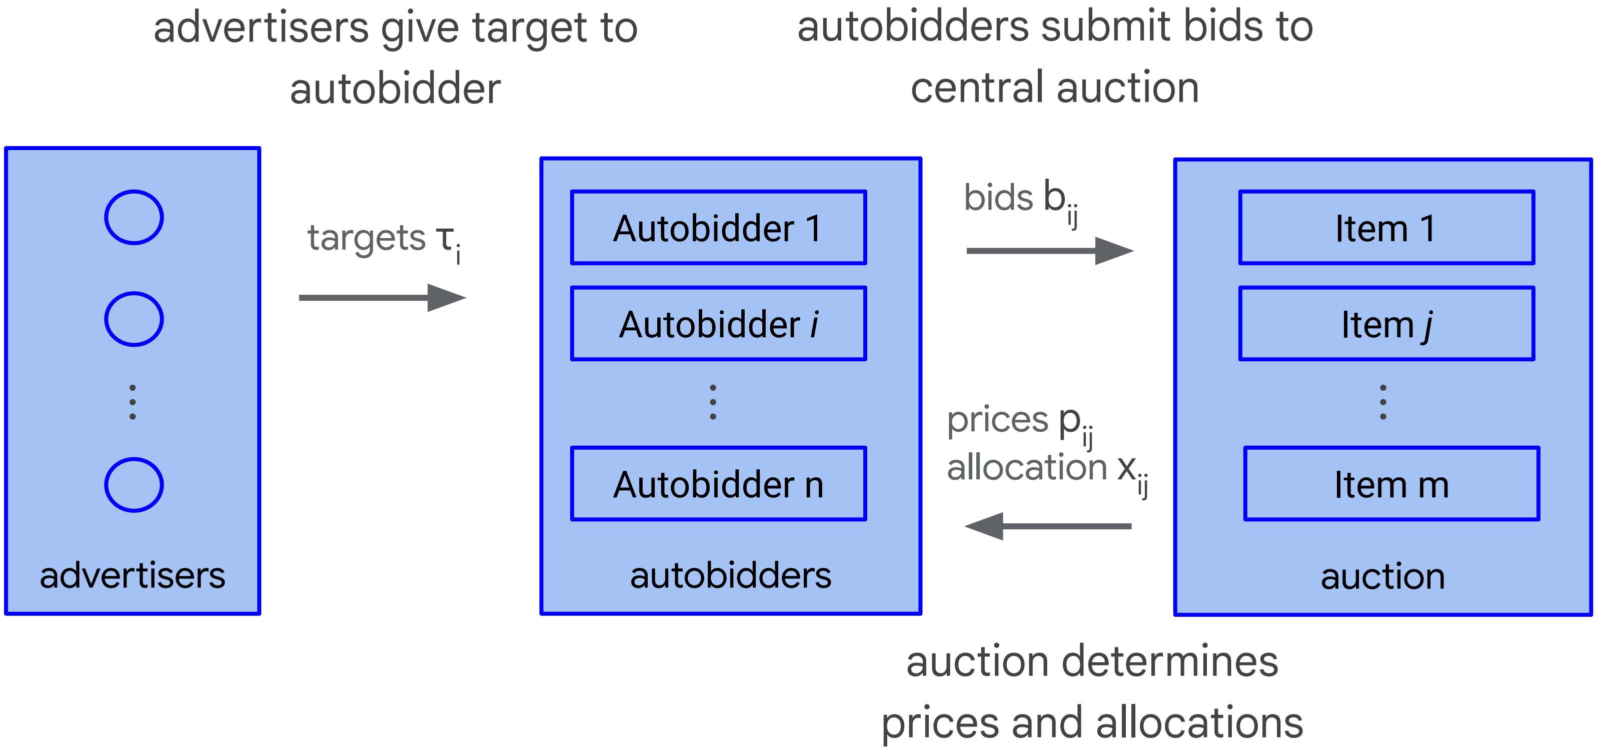

Complex High-Dimensional Auto-bidding Dynamics

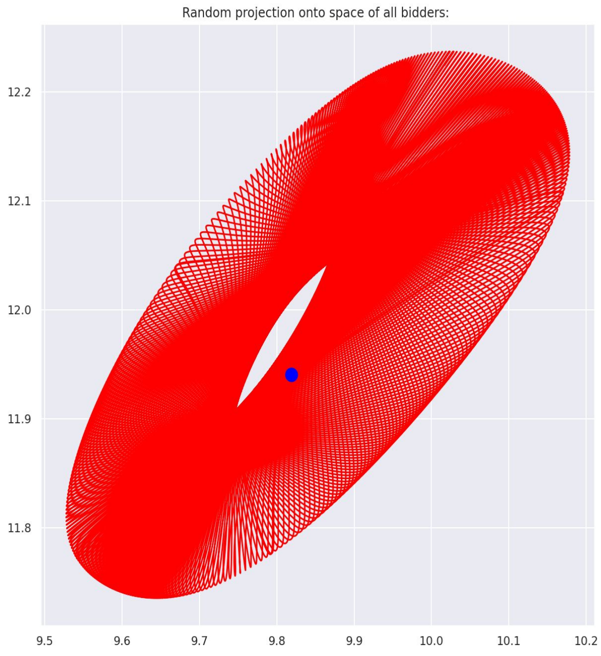

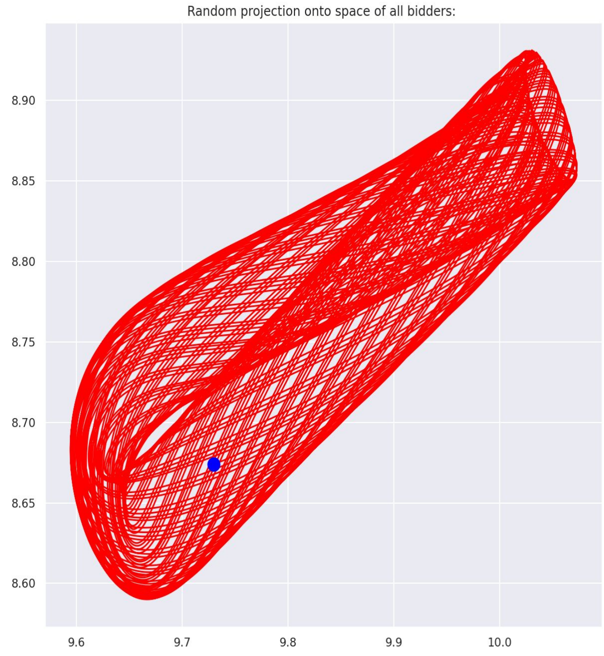

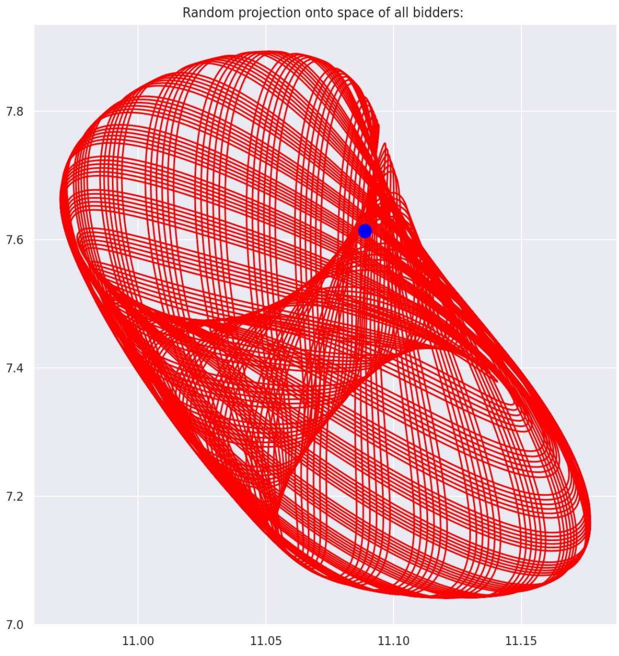

Final Discussion & Open Directions

- **Simple, practical** mechanisms can lead to **complex** behaviors.
- Embracing complexity can help us design **more robust** solutions.
- Could instability, chaos be seen an **asset** instead of a problem?

  (e.g. chaotic dynamics as pseudo-random generators)
- The advent of **AI** and/or **decentralization** will probably create

  increasingly complicated real world markets.

  New ideas are needed…

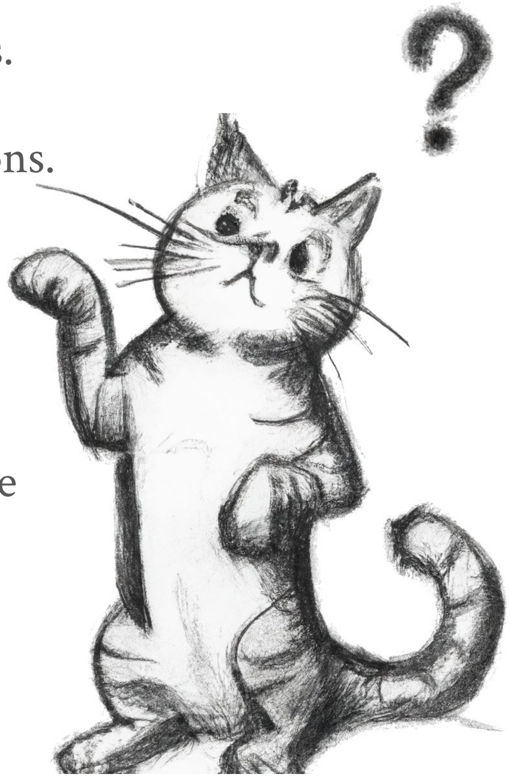

Thank you