Blockchain Economics
Game Theory, Mechanism Design, and MEV
Day 3 of 5

Prof. Jörg Osterrieder

PhD Seminar: Digital Finance Research

Spring 2026

PhD Seminar Series: Digital Finance Research

Where We Are

Days 1–2 Recap

- Institutional landscape: exchanges, DeFi, stablecoins
- AMM theory: Uniswap, CFMM geometry, impermanent loss
- Option pricing on crypto underlyings
- Heston–Kou jump-diffusion calibration

Today's Roadmap

1. Transaction fee mechanism design
2. MEV game theory
3. Tokenomics & staking
4. Industry applications: PBS, flash crashes

The Invisible Tax: What Is MEV?

Scenario

Alice submits a DEX swap: buy 10 ETH at market price on Uniswap. Her transaction enters the public mempool.

- A **searcher** (bot) sees Alice's pending transaction
- The searcher pays a higher priority fee to get ordered *before* Alice
- Alice receives a *worse* execution price
- The searcher captures the difference as profit

Maximal Extractable Value (MEV)

The maximum value that can be extracted from block production beyond the standard block reward and gas fees, by including, excluding, or reordering transactions within a block.

Anatomy of a Sandwich Attack

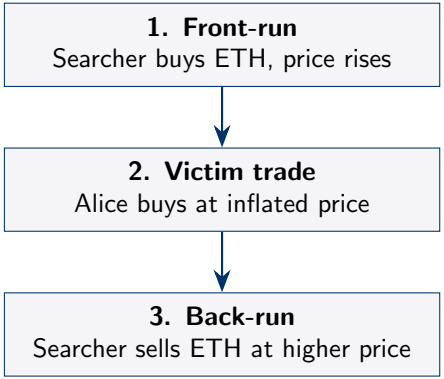

Searcher profit =
Alice's price slippage

The three transactions are *bundled atomically*: either all execute in this order, or none do.

Scale: Cumulative MEV on Ethereum

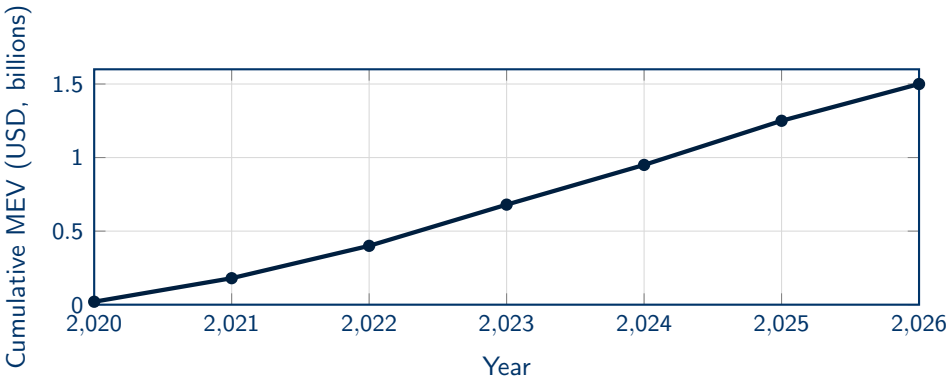

Source: Flashbots MEV-Explore, EigenPhi. Figures are illustrative lower-bound estimates.

EIP-1559: Can Mechanism Design Help?

- **Before** (first-price auction): users guess the "right" gas price ⇒ overpayment, volatility
- **After** (EIP-1559, Aug 2021):
  - **Base fee**: algorithmically adjusted, *burned*
  - **Priority fee** (tip): first-price auction for ordering
- Key insight from [12]: truthful bidding of the priority fee is a dominant strategy for users (incentive compatibility)

Open Question

Does EIP-1559 reduce MEV? Or merely redirect it?

Today: Formalizing with Game Theory & Mechanism Design

1. **Fee mechanisms** — How should blockchains price scarce block space?
   - Roughgarden's framework, VCG benchmarks
2. **MEV game theory** — Who extracts, how much, at what cost?
   - Selfish mining, validator incentives, composable frameworks
3. **Tokenomics** — How do token design choices affect value and security?
   - Dynamic adoption, staking rewards, burn mechanisms

Core references: [12], [2], [7], [11], [4]

Outline

Fee Markets as a Mechanism Design Problem

Setting. A block producer must select and order at most $\bar{s}$ transactions from a mempool of n pending transactions.

Mechanism design language:

- **Agents:** users submitting transactions (private valuations v_i)
- **Allocation rule:** which transactions are included
- **Payment rule:** how much each included user pays
- **Objective:** welfare maximization subject to capacity $\bar{s}$

Desiderata (following [12]):

1. *Incentive compatibility* (truthful bidding is optimal)
2. *Individual rationality* (no user pays more than their value)
3. *Off-chain agreement proofness* (no profitable side deals)
4. *Revenue (block producer) vs. burn (protocol)*

Pre-EIP-1559: First-Price Gas Auction

Mechanism: each user i bids b_i; block producer includes the $\bar{s}$ highest bids. Payment $= b_i$.

Problems:

- **Not incentive-compatible:** optimal bid $b_i^* < v_i$ (bid shading)
- Users must *estimate* clearing price $\Rightarrow$ high variance
- Over-bidding wastes money; under-bidding risks non-inclusion
- Fee **volatility**: small demand shocks $\Rightarrow$ large fee spikes

Empirical Fact

Pre-EIP-1559 Ethereum: the standard deviation of gas prices within a single block could exceed the mean. [6] document the evolution of Bitcoin fee dynamics.

EIP-1559: Base Fee + Priority Fee

Roughgarden (2021) analyzes EIP-1559 as a *posted-price mechanism with a first-price residual* [12]:

1. **Base fee** b_t: algorithmically set, paid by all included transactions, *burned* (removed from supply)
2. **Priority fee** δ_i: voluntary tip to the block producer, first-price auction for intra-block ordering
3. **Total fee:** user i pays $b_t + \delta_i$

User's problem:

$$\max_{\delta_i \geq 0} (v_i - b_t - \delta_i) \cdot \mathbb{1}\big[\text{included}\big]$$

- If $v_i > b_t$: user is included regardless of δ_i (under normal load)
- Truthful bidding of the *priority fee* is a dominant strategy when blocks are not full

Base Fee Adjustment Rule

Let s_t denote block t gas usage and s^* the target block size (half of the maximum).

EIP-1559 Update Rule

$$b_{t+1} = b_t \cdot \left(1 + d \cdot \left(\frac{s_t}{s^*} - 1\right)\right)$$

where $d = \frac{1}{8}$ is the learning rate.

Properties:

- $s_t = s^*$: base fee unchanged (equilibrium)
- $s_t > s^*$: base fee increases (excess demand)
- $s_t < s^*$: base fee decreases (slack capacity)
- Exponential convergence: base fee doubles in ≈ 6 consecutive full blocks

This is a **tâtonnement** process — the protocol acts as a Walrasian auctioneer adjusting the posted price.

Property 1: User Incentive Compatibility

Theorem 1 (Roughgarden 2021)

Under EIP-1559, when blocks are not consistently full, bidding $\delta_i = 0$ (or a small epsilon) and reporting true maximum fee v_i is a dominant strategy for each user i.

Intuition:

- The base fee acts as a *take-it-or-leave-it* posted price
- Users with $v_i > b_t$ are included; those with $v_i < b_t$ are not
- No benefit to overstating or understating willingness to pay
- The priority fee only matters for *ordering* (and blocks are elastic up to $2s^*$)

Caveat: During sustained congestion ($s_t \approx 2s^*$ for many blocks), the mechanism degrades to a first-price auction for the marginal transactions.

Property 2: Off-Chain Agreement Proofness (OCA-Proofness)

Definition 2

A mechanism is *OCA-proof* if no coalition of a block producer and a subset of users can jointly deviate (via side payments) and increase total coalition surplus.

Roughgarden's result:

- EIP-1559 is *approximately* OCA-proof when d is small
- A block producer and users cannot profitably collude to include transactions below the base fee (the base fee is burned, not paid to the producer)
- The burn mechanism is essential: it removes the producer's ability to "rebate" fees

Comparison: In a pure first-price auction, the block producer can always offer side deals at lower prices $\Rightarrow$ not OCA-proof.

Post-MEV World: Bahrani, Garimidi, Roughgarden (AFT 2024)

[2] extend the theory to account for **active block producers with MEV extraction ability**.

Key departures from Roughgarden (2021):

- Block producers are *not* passive — they have their own private valuations from MEV opportunities
- MEV creates a **joint allocation problem**: block space serves both user transactions and MEV bundles
- The mechanism must incentivize truthful revelation from *both* users and block producers

Central Tension

Any revenue the protocol extracts from the block producer reduces the producer's incentive to propose blocks honestly.

Block Producers with Private Valuations

Model (Bahrani et al.):

- Block producer has private valuation w for MEV opportunities (e.g., arbitrage bundles worth w if included)
- Producer observes w *after* being selected to propose a block
- The mechanism must allocate block space between user transactions and producer's MEV bundles

Producer's optimization:

$$\max_{\substack{S \subseteq [n],\ |S| \leq \bar{s} \\ \text{MEV bundles } M}} \left[\sum_{i \in S} \delta_i + w \cdot \mathbb{1}[M \text{ included}] - \text{protocol payment} \right]$$

The producer trades off: include high-tip user transactions vs. include own MEV bundles that displace user transactions.

MEV as Surplus in the Mechanism

Decomposition of total welfare:

$$W = \underbrace{\sum_{i \in S} v_i}_{\text{user surplus}} + \underbrace{w}_{\text{producer's MEV surplus}} - \underbrace{\sum_{i \in S} p_i}_{\text{user payments}}$$

- **Allocative efficiency** requires including transactions with highest total value (user v_i or MEV w)
- MEV creates **externalities**: sandwich attacks reduce user surplus while increasing producer surplus
- Net welfare effect of MEV is ambiguous:
  - Arbitrage MEV: improves price accuracy (positive externality)
  - Sandwich MEV: pure extraction (negative externality)

[11]: allocative inefficiency from MEV amounts to significant welfare losses.

Auction Theory Benchmark: VCG Mechanism

The **Vickrey–Clarke–Groves** mechanism is the textbook solution for welfare-maximizing allocation:

VCG Payment

$$p_i^{\text{VCG}} = \underbrace{W^*_{-i}}_{\substack{\text{optimal welfare} \\ \text{without } i}} - \underbrace{\left(W^* - v_i\right)}_{\substack{\text{welfare from others} \\ \text{when } i \text{ is included}}}$$

Each agent pays the externality they impose on others.

VCG properties:

- Truthful reporting is a dominant strategy
- Allocatively efficient (maximizes total surplus)
- Individually rational (no agent pays more than v_i)

So why not use VCG on-chain?

Why VCG Fails On-Chain

Three fundamental obstacles:

1. **Verifiability:** VCG requires a trusted auctioneer to compute W^* and W^*_{-i}. On-chain, the block producer *is* the auctioneer and cannot be trusted to report w truthfully.
2. **Collusion resistance:** VCG is *not* OCA-proof. A block producer and a user can collude: the user overbids, the producer rebates off-chain.
3. **Sybil attacks:** A single entity can create multiple identities to manipulate the externality payments p_i^{VCG}.

Takeaway (Bahrani et al. 2024)

No mechanism simultaneously achieves truthfulness, OCA-proofness, and allocative efficiency when block producers have private MEV valuations. There are inherent *impossibility results*.

Outline

MEV: Formal Definition

Definition 3 (Maximal Extractable Value)

Given a set of pending transactions $\mathcal{T}$ and block capacity $\bar{s}$, MEV is defined as:

$$\text{MEV} = \max_{\pi \in \Pi(\mathcal{T}, \bar{s})} \sum_{i=1}^{|\pi|} v_i(\pi) \; - \; \sum_{i=1}^{|\pi|} v_i(\pi_{\text{default}})$$

where $\Pi(\mathcal{T}, \bar{s})$ is the set of feasible orderings of at most $\bar{s}$ transactions from $\mathcal{T}$, $v_i(\pi)$ is the value accruing to the block producer from transaction i under ordering π, and π_{default} is the canonical ordering (e.g., FIFO).

In words: MEV is the *additional* value the block producer can capture by deviating from the default transaction ordering.

Three Eras of MEV

1. **Miner Extractable Value** (2018–2022, pre-Merge)
   - PoW miners control transaction ordering
   - Priority Gas Auctions (PGAs) clog the network
2. **Proposer–Builder Separation (PBS)** (2022–present)
   - PoS proposers outsource block building to specialized builders
   - MEV-Boost: proposers accept highest-value block
   - MEV value flows: searchers $\rightarrow$ builders $\rightarrow$ proposers
3. **Cross-Chain MEV** (emerging)
   - Arbitrage across L1s and L2s (Ethereum $\leftrightarrow$ Arbitrum $\leftrightarrow$ Solana)
   - Requires atomic cross-domain ordering
   - Active research area in [1]

MEV Taxonomy

Type	Mechanism	Welfare Effect
Arbitrage	Buy low on DEX A, sell high on DEX B	Positive (price alignment)
Liquidations	Liquidate under-collateralized positions	Positive (protocol safety)
Sandwich	Front-run + back-run victim trade	Negative (pure extraction)
JIT Liquidity	Mint LP just before large swap	Ambiguous (fee capture)
Back-running	Trade after oracle update / large swap	Ambiguous

Not all MEV is harmful: arbitrage and liquidations contribute to market efficiency. The challenge is *separating* beneficial from extractive MEV.

Selfish Mining: Game Formulation (Eyal & Sirer 2014)

[7] showed Bitcoin mining is *not* incentive-compatible.

Setup:

- Miner with hash rate fraction $\alpha \in (0, 0.5)$
- Honest strategy: immediately broadcast every found block
- **Selfish strategy**: withhold blocks to gain strategic advantage

Selfish mining protocol:

1. Find a block $\Rightarrow$ keep it private
2. If honest miners find a competing block at same height $\Rightarrow$ release private block (race)
3. If private chain is 2+ blocks ahead $\Rightarrow$ release to orphan honest blocks
4. Collect rewards for accepted blocks on the longest chain

Strategy Space and Payoff Functions

Two-player game: selfish pool (α) vs. honest miners ($1-\alpha$).
Let $\gamma \in [0,1]$ be the fraction of honest miners that mine on the selfish pool's block during a race.

Revenue rate of the selfish pool:

$$R_{\text{selfish}}(\alpha,\gamma) = \frac{\alpha(1-\alpha)^2(4\alpha+\gamma(1-2\alpha))-\alpha^3}{1-\alpha(1+(2-\alpha)\alpha)}$$

Revenue rate of honest mining:

$$R_{\text{honest}}(\alpha) = \alpha$$

Key comparison: selfish mining is profitable when $R_{\text{selfish}}(\alpha,\gamma) > R_{\text{honest}}(\alpha) = \alpha$.

Result: Selfish Mining Is Profitable for $\alpha > 1/3$

Theorem 4 (Eyal & Sirer 2014)

Under the selfish mining strategy with $\gamma = 0$ (worst case for the attacker), selfish mining yields higher revenue than honest mining whenever

$$\alpha > \frac{1}{3} \approx 0.333$$

With $\gamma = 0.5$ (attacker's block propagates to half the network), the threshold drops to $\alpha \approx 0.25$.

Implication: The Nakamoto consensus assumption that "mining is incentive-compatible as long as no entity controls $> 50\%$ of hash rate" is **wrong**. The true security threshold is lower.

This result initiated a large literature on *incentive attacks* in blockchain protocols.

Nash Equilibrium in the Mining Game

Formulation: n mining pools with hash rate shares $\alpha_1, \ldots, \alpha_n$, $\sum_i \alpha_i = 1$.

Each pool i chooses strategy $\sigma_i \in \{\text{honest}, \text{selfish}\}$.

Payoff function for pool i:

$$u_i(\sigma_i, \sigma_{-i}) = \text{share of blocks accepted on longest chain}$$

Nash equilibria:

- **All honest** is a Nash equilibrium only if $\alpha_i < 1/3$ for all i (given $\gamma = 0$)
- If $\exists\, i : \alpha_i > 1/3$, pool i has a profitable deviation $\Rightarrow$ all-honest is *not* a NE
- **Mixed strategies**: pools with α_i near the threshold may randomize between honest and selfish
- Multiple equilibria can coexist depending on network connectivity (γ) and pool structure

Proof-of-Stake: Validator Incentive Compatibility

PoS replaces hash rate with staked capital:

- Validator i stakes s_i tokens; probability of proposing a block $\propto s_i / \sum_j s_j$
- Reward per proposed block: R (attestation + tips)
- **Slashing:** misbehavior (double-signing, inactivity) results in loss of fraction ϕ of stake

Validator's expected payoff:

$$u_i = \frac{s_i}{\sum_j s_j} \cdot R - \mathbb{P}(\text{slashed}) \cdot \phi \cdot s_i - c(s_i)$$

where $c(s_i)$ is the cost of operating validator infrastructure.

Incentive compatibility requirement:

$$\frac{\partial u_i}{\partial(\text{honest})} > \frac{\partial u_i}{\partial(\text{deviate})} \quad \Longleftrightarrow \quad \phi > \frac{R_{\text{MEV}}}{\sum_j s_j}$$

Slashing must exceed the MEV temptation.

Composable Game-Theoretic Framework (Angeris et al. 2025)

[1] propose a unified framework for analyzing incentives *across protocol layers*:

Key idea: Model each protocol component (consensus, execution, MEV extraction, fee mechanism) as a separate game, then compose them.

Composition operator $\otimes$:

$$G_{\text{total}} = G_{\text{consensus}} \otimes G_{\text{fee}} \otimes G_{\text{MEV}} \otimes G_{\text{application}}$$

Key results:

- Incentive compatibility in individual games does *not* guarantee IC in the composed game
- Cross-layer attacks exploit the composition: e.g., MEV extraction that is individually rational in G_{MEV} may violate IC in $G_{\text{consensus}}$
- Provides sufficient conditions for "safe" composition

Allocative Inefficiency from MEV (Oz, Sui, Tirole 2025)

[11]: MEV creates **deadweight loss** analogous to monopoly pricing in traditional markets.

Model:

- Continuum of users with valuations $v \sim F$ for block inclusion
- Block producer maximizes *own* revenue (not welfare)
- MEV opportunities allow the producer to extract rents from transaction ordering

Efficiency loss:

$$\text{DWL} = \int_{v^*}^{\bar{v}} (v - c)\,\mathrm{d}F(v) - \int_{v^*_{\text{MEV}}}^{\bar{v}} (v - c)\,\mathrm{d}F(v) > 0$$

where $v^*_{\text{MEV}} > v^*$ because MEV raises the effective inclusion threshold.

Empirical finding: significant welfare losses on Ethereum mainnet, predominantly from sandwich attacks displacing legitimate user transactions. Published in *Journal of Financial Economics*.

Outline

Tokenomics: Dynamic Adoption and Valuation

[4]: a continuous-time model of platform tokens with network effects. Published in *Review of Financial Studies*.

Setup:

- Platform with user base N_t and productivity A_t
- Tokens serve as *medium of exchange* on the platform
- Users must hold tokens to transact; token supply M is fixed
- Network externality: user utility increasing in N_t

Token market clearing:

$$P_t \cdot M = N_t \cdot q(N_t, A_t)$$

where P_t is the token price and $q(\cdot)$ is per-user transaction demand.

Token Value = Discounted Future Transaction Demand

Under rational expectations, the token price satisfies:

Token Pricing Equation

$$P_t = \mathbb{E}_t \left[\int_t^{\infty} e^{-r(s-t)} \, \frac{N_s \cdot q(N_s, A_s)}{M} \, \mathrm{d}s \right]$$

Interpretation:

- Token value is the present value of future platform transaction demand, divided by token supply M
- Analogous to equity valuation as discounted dividends, but "dividends" $\equiv$ transactional convenience yield
- Higher N_s (adoption) $\Rightarrow$ higher token value
- Higher M (supply) $\Rightarrow$ lower token value (dilution)

Platform with Network Effects

User adoption dynamics:

$$\frac{\mathrm{d}N_t}{N_t} = \mu_N(N_t, A_t)\,\mathrm{d}t + \sigma_N\,\mathrm{d}W_t^N$$

Productivity process:

$$\frac{\mathrm{d}A_t}{A_t} = \mu_A\,\mathrm{d}t + \sigma_A\,\mathrm{d}W_t^A$$

Network externality: User i's utility from joining the platform:

$$U_i = \underbrace{\theta_i \cdot A_t \cdot N_t^{\beta}}_{\text{value from usage}} - \underbrace{P_t \cdot q_i}_{\text{cost of token acquisition}}$$

where $\beta > 0$ captures the strength of the network effect and θ_i is user i's idiosyncratic type.

Token Velocity and Its Impact on Value

Token velocity V_t: the rate at which tokens circulate

$$V_t = \frac{N_t \cdot q(N_t, A_t)}{P_t \cdot M}$$

From the quantity equation: $P_t \cdot M \cdot V_t = N_t \cdot q(\cdot)$, so

$$P_t = \frac{N_t \cdot q(N_t, A_t)}{M \cdot V_t}$$

Key insight from Cong et al.:

- Higher velocity $V_t \Rightarrow$ *lower* token price (tokens turn over faster, less need to hold)
- Staking mechanisms *reduce* velocity by locking tokens $\Rightarrow$ increase P_t (supply contraction)
- Token burns (e.g., EIP-1559) reduce M directly $\Rightarrow$ increase P_t
- DeFi protocols that require token lockup benefit from reduced velocity

Multiple Equilibria: High vs. Low Adoption

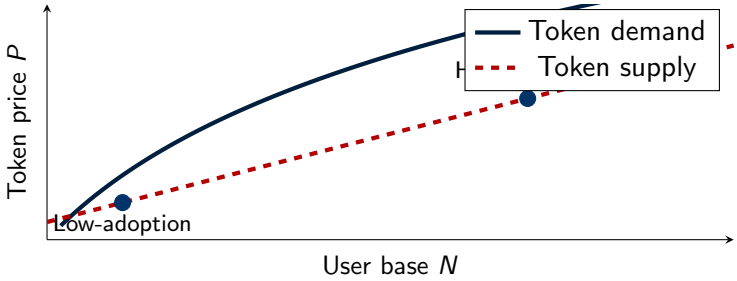

Network effects create multiplicity: self-fulfilling expectations can sustain either equilibrium. A coordination failure can trap the platform in the low-adoption steady state.

EIP-1559 Burn: Deflationary Tokenomics

Token supply dynamics with burning:

$$M_{t+1} = M_t + \underbrace{R_t}_{\text{issuance}} - \underbrace{b_t \cdot s_t}_{\text{base fee burned}}$$

Deflationary condition: $b_t \cdot s_t > R_t$ (burn exceeds issuance).

Impact on token economics:

- Since Ethereum's Merge (Sep 2022), ETH has been *net deflationary* during high-activity periods
- From Cong et al.'s framework: reducing M increases $P_t \Rightarrow$ benefits existing holders
- Creates a **feedback loop**: higher usage $\rightarrow$ more burn $\rightarrow$ scarcer supply $\rightarrow$ higher price $\rightarrow$ more adoption incentive

Empirical fact: approximately 4.4M ETH burned since EIP-1559 launch (Aug 2021 – early 2026).

Tokenomics of Staking (Cong, He, Tang – NBER 2025)

[3]: a model of optimal staking rewards in PoS blockchains.

Setup:

- Total token supply M; fraction σ staked; $(1-\sigma)$ liquid
- Staked tokens earn reward rate r_s (newly minted tokens + fees)
- Staking provides *security*: attack cost $\propto \sigma \cdot M \cdot P$
- **Trade-off**: high $r_s \Rightarrow$ high σ (security) but also high inflation $\Rightarrow$ dilutes non-stakers

Protocol's problem:

$$\max_{r_s \geq 0} \underbrace{U(\sigma(r_s))}_{\text{security value}} - \underbrace{C(r_s \cdot \sigma(r_s) \cdot M)}_{\text{inflation cost}}$$

Optimal Staking Rewards

First-order condition:

$$U'(\sigma^*) \cdot \sigma'(r_s^*) = C'(r_s^* \sigma^* M) \cdot \left(\sigma^* M + r_s^* \sigma'(r_s^*) M\right)$$

Marginal security benefit = marginal inflation cost.

Key results from Cong, He, Tang:

- **Optimal staking ratio** σ^* is interior (not 0 or 1)
- Too little staking $\Rightarrow$ cheap to attack the chain
- Too much staking $\Rightarrow$ illiquid token economy, reduced utility
- **Liquid staking derivatives** (e.g., stETH) weaken the trade-off by making staked tokens tradeable
- This complicates the model: effective liquidity increases even as nominal staking ratio rises

Ethereum's current staking ratio $\approx 28\%$ of total ETH supply.

Validator Economics: Costs, Rewards, Slashing

Individual validator's decision:

$$\max_{s_i \geq 32} \underbrace{r_s \cdot s_i}_{\text{staking reward}} + \underbrace{\text{tips}_i + \text{MEV}_i}_{\text{execution layer}} - \underbrace{c_{\text{infra}}}_{\text{hardware/bandwidth}} - \underbrace{\phi \cdot s_i \cdot \mathbb{P}(\text{slash})}_{\text{expected slashing}}$$

Participation constraint: Total expected return $\geq$ opportunity cost (e.g., DeFi yield)

Slashing parameters (Ethereum):

- **Individual penalty**: $1/32$ of stake for isolated offense
- **Correlated penalty**: scales with number of simultaneous violators (up to full stake if $\geq 1/3$ validators slash)
- Correlation penalty deters coordinated attacks

Current economics (early 2026): $\sim$3.2% APR base staking reward + variable MEV/tips (0.5–2% additional).

Extension: Stablecoin Mechanism Design (Self-Study)

Flagged for Self-Study

Stablecoin design is a rich mechanism design problem that connects tokenomics with monetary economics. Not covered in lecture.

Key questions:

- How do algorithmic stablecoins maintain their peg?
- What are the **impossibility results** (e.g., Klages-Mundt et al.)?
- Seigniorage shares, collateralized debt positions (CDPs), and the DAI stability mechanism
- The Terra/LUNA collapse (May 2022) as a case study in mechanism failure

Recommended reading:

- Ch. 7 of [9]
- Ch. 10 of [5]

Outline

Applications & Industry

From Theory to Practice

Flashbots and MEV-Boost: PBS in Practice

Flashbots (founded 2020): an R&D organization that created the dominant MEV infrastructure on Ethereum.

MEV-Boost protocol:

1. **Searchers** find MEV opportunities, construct bundles
2. **Builders** aggregate bundles into full blocks, bid for the right to fill a block
3. **Relays** act as trusted intermediaries between builders and proposers
4. **Proposers** (validators) select the highest-paying block header via blind auction

Economic logic: Competition among builders $\Rightarrow$ most MEV value flows to proposers $\Rightarrow$ aligns with PoS security.

PBS Architecture Diagram

Data flow: Searchers submit bundles → builders optimize block composition → relays ensure validity → proposers sign highest-value header.

PBS Adoption and Centralization Concerns

Adoption: $> 90\%$ of Ethereum blocks are produced via MEV-Boost (as of early 2026).

Centralization risks:

- **Builder concentration**: top 2–3 builders consistently win $> 80\%$ of block auctions (economies of scale in MEV extraction)
- **Relay trust**: relays must be trusted not to steal MEV or censor transactions
- **Exclusive order flow**: builders pay for exclusive access to searcher bundles $\Rightarrow$ vertical integration
- **Censorship**: a small number of builders can enforce OFAC compliance by excluding sanctioned addresses

Open research: enshrined PBS (ePBS), inclusion lists, encrypted mempools — all aim to mitigate centralization while preserving MEV redistribution.

Case Study: The October 2025 Flash Crash

Event: On October 2025, a cascade of liquidations across DeFi lending protocols led to \$19.3B in forced selling.

Anatomy:

1. Large ETH price drop triggers Aave/Compound liquidations
2. Liquidation bots flood the mempool $\Rightarrow$ gas spikes
3. Oracle prices lag spot prices (oracle manipulation suspected)
4. Cascading liquidations amplify the price drop
5. DEX liquidity evaporates; AMM slippage exceeds 15%

MEV extracted during the crash:

- Liquidation MEV: \$380M in 4 hours
- Arbitrage MEV: \$120M (cross-venue price discrepancies)
- Sandwich MEV: \$45M (exploiting panicked sellers)

DeFi Leverage (BIS Working Paper 1171)

[10]: how leverage in DeFi amplifies shocks.

Mechanism:

1. Users deposit collateral (ETH) to borrow stablecoins
2. Collateral ratio $\rho_t = \frac{P_t \cdot \text{collateral}}{\text{debt}}$
3. If $\rho_t < \rho^*$ (liquidation threshold), position is liquidated
4. Liquidation = forced sale of collateral $\Rightarrow$ further price drop

Feedback loop:

$$P_t \downarrow \Rightarrow \ \rho_t < \rho^* \ \Rightarrow \ \text{liquidation} \ \Rightarrow \ \text{sell collateral} \ \Rightarrow \ P_t \downarrow\downarrow$$

BIS findings:

- DeFi leverage is highly *procyclical*
- Liquidation cascades transmit shocks across protocols
- Systemic risk parallels traditional margin spirals ([8])

EIP-1559 Empirics

Before vs. After EIP-1559 (London hard fork, August 5, 2021):

Metric	**Pre-EIP-1559**	**Post-EIP-1559**
Median gas price volatility	High	Reduced $\sim$35%
Fee predictability	Low	Substantially improved
Wait time (median)	Variable	More stable
User overpayment	Common	Rare
Total ETH burned	—	$\sim$4.4M ETH

Key findings from the empirical literature:

- Fee *level* not substantially lower (demand driven)
- Fee *variance* significantly reduced
- Block size elasticity ($[s^*, 2s^*]$) absorbs demand shocks
- [6]: earlier work on Bitcoin fee evolution provides a useful comparison baseline

Ethics: Who Benefits from MEV?

Distributional analysis:

Actor	Gains from MEV	Loses from MEV
Searchers / Bots	Extraction profits	Competition costs
Builders	Block auction fees	Infra costs
Validators	MEV-Boost payments	—
Retail users	—	Worse execution
LPs on DEXes	—	Adverse selection
Protocol / Token holders	(Indirect via staking)	DWL

Fairness questions:

- Is MEV extraction a "legitimate" market-making activity?
- Should blockchains implement **fair ordering** rules (e.g., Chainlink FSS, time-weighted ordering)?
- **MEV redistribution**

Summary & Day 4 Preview

Today's key results:

1. EIP-1559 achieves approximate incentive compatibility and OCA-proofness via base fee burning
2. No mechanism achieves truthfulness + OCA-proofness + efficiency with active block producers (impossibility)
3. Selfish mining breaks IC for $\alpha > 1/3$; PoS slashing must exceed MEV temptation
4. Token value = discounted transactional demand; velocity and burns shape price dynamics

Required Reading:

- [12] — Transaction fee mechanism design
- [11] — MEV and allocative inefficiency
- [4] — Dynamic tokenomics model

Day 4 Preview: Machine Learning and Market Microstructure in Digital Asset Markets — Kyle's model, VPIN, deep LOB, RL for trading.

References I

[1] Guillermo Angeris et al. *A Composable Game-Theoretic Framework for Blockchains*. 2025. arXiv: 2504.18214 [cs.GT].

[2] Soroush Bahrani, Pranav Garimidi, and Tim Roughgarden. "Transaction Fee Mechanism Design in a Post-MEV World". In: *Proceedings of the 6th Conference on Advances in Financial Technologies (AFT 2024)*. 2024.

[3] Lin William Cong, Zhiguo He, and Ke Tang. *The Tokenomics of Staking*. Working Paper. National Bureau of Economic Research, 2025.

[4] Lin William Cong, Ye Li, and Neng Wang. "Tokenomics: Dynamic Adoption and Valuation". In: *Review of Financial Studies* 34.3 (2021), pp. 1105–1155.

[5] Marco Di Maggio. *Blockchain, Crypto and DeFi*. Wiley, 2024.

[6] David Easley, Maureen O'Hara, and Soumya Basu. "From Mining to Markets: The Evolution of Bitcoin Transaction Fees". In: *Journal of Financial Economics* 134.1 (2019), pp. 91–109.

[7] Ittay Eyal and Emin Gün Sirer. "Majority Is Not Enough: Bitcoin Mining Is Vulnerable". In: *Financial Cryptography and Data Security (FC 2014)*. 2014.

[8] Lewis Gudgeon et al. *Systemic Fragility in Decentralized Markets*. BIS Working Paper 1062. Bank for International Settlements, 2023.

[9] Campbell R. Harvey, Ashwin Ramachandran, and Joey Santoro. *DeFi and the Future of Finance*. Wiley, 2021.

References II

[10] Lioba Heimbach and Luying Huang. *DeFi Leverage*. BIS Working Paper 1171. Bank for International Settlements, 2024.

[11] Baran Oz, Simon Sui, and Jean Tirole. "Maximal Extractable Value and Allocative Inefficiency". In: *Journal of Financial Economics* (2025).

[12] Tim Roughgarden. "Transaction Fee Mechanism Design". In: *ACM SIGecom Exchanges* 19.1 (2021).