A new report: Transaction Fee Mechanism Design for the Ethereum Blockchain: An Economic Analysis of EIP-1559 http://timroughgarden.org/papers/eip1559.pdf… Top ten takeaways in a thread: 1/14

1. No transaction fee mechanism, EIP-1559 or otherwise, is likely to substantially decrease average transaction fees; persistently high transaction fees is a scalability problem, not a mechanism design problem. 2/14

2. EIP-1559 should decrease the variance in transaction fees and the delays experienced by some users through the flexibility of variable-size blocks. 3/14

3. EIP-1559 should improve the user experience through easy fee estimation, in the form of an "obvious optimal bid," outside of periods of rapidly increasing demand. 4/14

4. The short-term incentives for miners to carry out the protocol as intended are as strong under EIP-1559 as with first-price auctions. 5/14

5. The game-theoretic impediments to double-spend attacks, censorship attacks, denial-of-service attacks, and long-term revenue-maximizing strategies such as base fee manipulation appear as strong under EIP-1559 as with first-price auctions. 6/14

6. EIP-1559 should at least modestly decrease the rate of ETH inflation through the burning of transaction fees. 7/14

7. The seemingly orthogonal goals of easy fee estimation and fee burning are inextricably linked through the threat of off-chain agreements. 8/14

8. Alternative designs include paying base fee revenues forward to miners of future blocks rather than burning them; and replacing variable user-specified tips by a fixed hard-coded tip. 9/14

9. EIP-1559's base fee update rule is somewhat arbitrary and should be adjusted over time. 10/14

10. Variable-size blocks enable a new (but expensive) attack vector: overwhelm the network with a sequence of maximum-size blocks. 11/14

Conclusions: 12/14

Does EIP-1559 offer an improvement over Ethereum's current transaction fee mechanism? The biggest potential benefits of the proposed changes are as advertised: easy fee estimation, in the form of an "obvious optimal bid" outside of periods of rapidly increasing demand (Theorem 6.8); lower variance in transaction fees due to increased flexibility in block size (Section 3.2); game-theoretic robustness to protocol deviations and off-chain agreements, both at the scale of a single block (Theorems 6.4 and 6.12) and of multiple blocks (Section 7); and reduced inflation due to fee burning (Section 9.1).

Most of the major risks in implementing EIP-1559 are the same as those for any major change to the Ethereum protocol: implementation errors; a fork caused by some parties rejecting the changes; extra complexity at the consensus layer; additional parameters to be tweaked with every network upgrade; and the spectre of unforeseeable downstream consequences. Additional risks specific to EIP-1559 include the possibility of a hostile reception by miners (due to lost revenue from burned transaction fees) and a coordinated response (Sections 7.5–7.6); and a new (if expensive) attack vector enabled by variable-size blocks (Sections 8.6.5–8.6.6).

Reasonable people will disagree on whether the benefits of EIP-1559 justify the risks in adopting it. Those who subscribe to a "why fix what isn't (too badly) broken" philosophy may prefer to stick with the status quo. For those who believe that consensus-layer innovation should continue to be a central part of Ethereum's future, however, the arguments in favor of EIP-1559 are strong.

Funding by
@d24nOrg
is gratefully acknowledged. 13/14

For helpful comments and discussions, thanks also to
Maryam Bahrani
@dimahledba
@TimBeiko
@VitalikButerin
Matheus Ferreira
@shemnon
@jamesfickel
@hasufl
@gakonst
Andrew Lewis-Pye
@barnabemonnot
@dmoroz
Mitchell Stern
@ATabarrok
Peter Zeitz! 14/14