Economic Determinants of Ethereum Transaction Fees in the Priority Fee and Proof of Stake Periods

Alexander Karaivanov* and Shayan Zarifian*

May 2024

Abstract

We analyze the economic determinants and dynamics of transaction fees in the Ethereum blockchain before and after two significant platform updates. The first is the August 2021 "London" upgrade, a switch from user-bid *gas price* (transaction fee per unit of complexity) to a fee model in which the gas price is the sum of an algorithmically determined base fee and an optional priority fee (tip) chosen by the user. The second update ("the Merge") is the switch from proof-of-work to proof-of-stake transactions validation in September 2022. We estimate the impact on Ethereum transaction fees of both demand factors (block utilization, transaction type, ETH price in USD) and algorithmic supply-side factors (the block gas limit and base fee). Using data from nearly 900 million blockchain transactions, we find that the gas price is statistically significantly positively associated with the block utilization rate. A larger share of contract call transactions or legacy (user-bid gas price) transactions is linked with higher gas prices on average. On the supply side, a higher block gas limit is statistically significantly associated with lower gas prices.

Keywords: transaction fees, blockchain, Ethereum, gas price, supply and demand factors, time series analysis

*Department of Economics, Simon Fraser University. Corresponding author email: akaraiva@sfu.ca.
We thank Anil Donmez, Bertille Antoine, Marie Rekkas and participants at the Canadian Economics Association conference for useful comments and suggestions. Karaivanov gratefully acknowledges financial support from the Social Sciences and Humanities Research Council of Canada, grants 435-2018-0111 and 435-2024-0317.

1 Introduction

Ethereum, established in 2015, is the second-largest blockchain platform by market capitalization, with a valuation of $452 billion as of May 22, 2024. Unlike other digital platforms mainly focused on cryptocurrency investment or trading (e.g., Bitcoin), Ethereum distinguishes itself with a broad range of applications and flexibility, by enabling users to create and deploy automatized digital "smart contracts" for diverse purposes including decentralized applications (DApps), decentralized financial services (DeFi), non-fungible tokens (NFTs), etc.

Transferring digital funds or deploying and interacting with digital smart contracts on the Ethereum blockchain requires paying a transaction fee determined in terms of the platform's internal resource unit called *gas* and the price per unit of gas, *gas price*. Each use of the blockchain (transaction) has an algorithmically specified gas requirement which depends on the transaction's computational complexity and number of operations.[1] The transaction fee that a user pays equals the transaction's gas requirement multiplied by the gas price.

Before August 5, 2021 the Ethereum gas price was determined in a way similar to a first-price auction – users submit gas price bids and the transactions of the users with the highest bids were prioritized for validation and execution, that is, included in more immediate data blocks. As of August 5, 2021, the economic mechanism for determining Ethereum transaction fees was changed significantly with the implementation of Ethereum Improvement Proposal EIP-1559, also known as 'the London upgrade'.[2] This platform code upgrade was introduced with the goal to reign in rising transaction fees and to reduce network congestion and delays in transaction processing.

Combined with a doubling of the gas supply (the gas limit per block), the major change introduced in the London upgrade was a switch from gas prices fully bid by the users to a new transaction fee model in which the gas price is formed as the sum of two components: a 'base fee' and a 'priority fee'. Specifically, the base fee of each block is algorithmically determined by the blockchain code (see Appendix B) as a function of two variables: the base fee of the immediately preceding block and the block utilization (gas used) in the immediately preceding block. Higher utilization (above 50% of the block gas limit) in block $t-1$ automatically triggers a proportional increase (capped at 12.5%) of the base fee in the following block t, while lower block utilization in block $t-1$ (below 50%) results in a proportional decrease of the base fee in block t, again by at most 12.5%. By design, the base fee therefore acts to stabilize Ethereum transaction costs, since the block-on-block gas price growth rate is capped and depends on the gas usage (that is, demand) in the preceding block. Paying the base fee is mandatory – a transaction which does not provide

[1] For instance, the gas requirement for a simple bilateral ETH transfer between two addresses is 21,000 gas. More complex transactions, for example deploying a digital contract, require substantially larger amounts of gas.

[2] See https://github.com/ethereum/EIPs/blob/master/EIPS/eip-1559.md for complete details.

sufficient Ether (ETH, Ethereum's digital currency unit) to pay the base fee would be rejected by the platform.

In addition to the base fee, users can optionally pay a so-called "priority fee" (akin to a tip), chosen in a way similar to the pre-London gas price bid. All submitted transactions must pay the base fee and transactions with higher priority fees are more likely to be included on the blockchain and thus executed sooner. Users are also allowed to continue to adhere to the old pre-London/EIP-1559 transaction format, i.e., set the gas price for their transaction directly, however, such transactions would be valid only if the bid gas price equals or exceeds the current block base fee. We refer to such transactions as "legacy transactions". In contrast, users who adopt the new EIP-1559 fee mechanism (we refer to such transactions as "non-legacy transactions"), choose two parameters ('max fee per gas' and 'max priority fee per gas') and the resulting gas price they pay is determined as the smaller value of 'max fee per gas' and the sum of the base fee and 'max priority fee per gas'.

A second major update to the Ethereum platform ('the Merge upgrade') was implemented on September 15, 2022 by switching from *proof of work* (also known as "mining") – the original and most commonly used consensus algorithm for verifying and recording transaction data securely on a blockchain,[3] see Nakamoto (2008) – to *proof of stake*, a protocol in which major ETH holders, who have posted collateral in a smart contract deployed on the blockchain, validate submitted transactions.[4] Unlike the computationally and electricity consumption heavy proof-of-work mechanism, proof-of-stake validation requires much lower computer processing power. For example, Kapengut and Mizrach (2023) estimate that the transition to proof of stake in Ethereum reduced electric energy consumption by 99.98%.[5] We focus on the Merge's impact on transaction fee determination and find that the switch to proof-of-stake resulted in an approximately 12%–15% increase in the daily gas supply by reducing and stabilizing the time interval between consecutively created blocks.[6]

To illustrate the impact of the 2021 London upgrade and the 2022 Merge upgrade on Ethereum gas prices, in Figure 1, we display the daily-aggregated average values of our main variable of

[3] In 'proof of work', miners (special users contributing computing power to operate the platform) compete to solve a computationally-hard cryptographic problem. The first miner to find a solution wins the right to add a new block of transactions to the blockchain and is rewarded with newly-minted ETH currency and the transaction fees (priority fees post-London) of all transactions included in the block.

[4] To qualify as a validator, a user must deposit 32 ETH into a designated smart contract. A single validator is randomly selected to serve as the current block builder/proposer and is responsible for constructing and broadcasting the block to the Ethereum platform, along with updating the virtual machine state.

[5] See also De Vries (2023) who estimates that the replacement of proof-of-work with proof-of-stake validation in Ethereum reduced electricity use by 99.84% to 99.9996%, a magnitude similar to the electric power requirement of a country as large as Ireland or Austria.

[6] On average, an Ethereum block is created approximately every 12 seconds.

interest – the Ethereum gas price, in Gwei and in US dollars (USD), over our period of analysis January 1, 2021 to December 31, 2022.[7] The block median gas price reached more than 200 Gwei in early 2021, was about 60 Gwei on August 5, 2021, the date of transition to base fee plus priority fee gas price ('London upgrade'), rose again to about 180 Gwei in September-October 2021, then fell to about 18 Gwei on September 15, 2022, the date of transition to proof-of-stake ('Merge upgrade') and stayed at similarly low values until the end of 2022.[8]

Figure 1: Gas Price

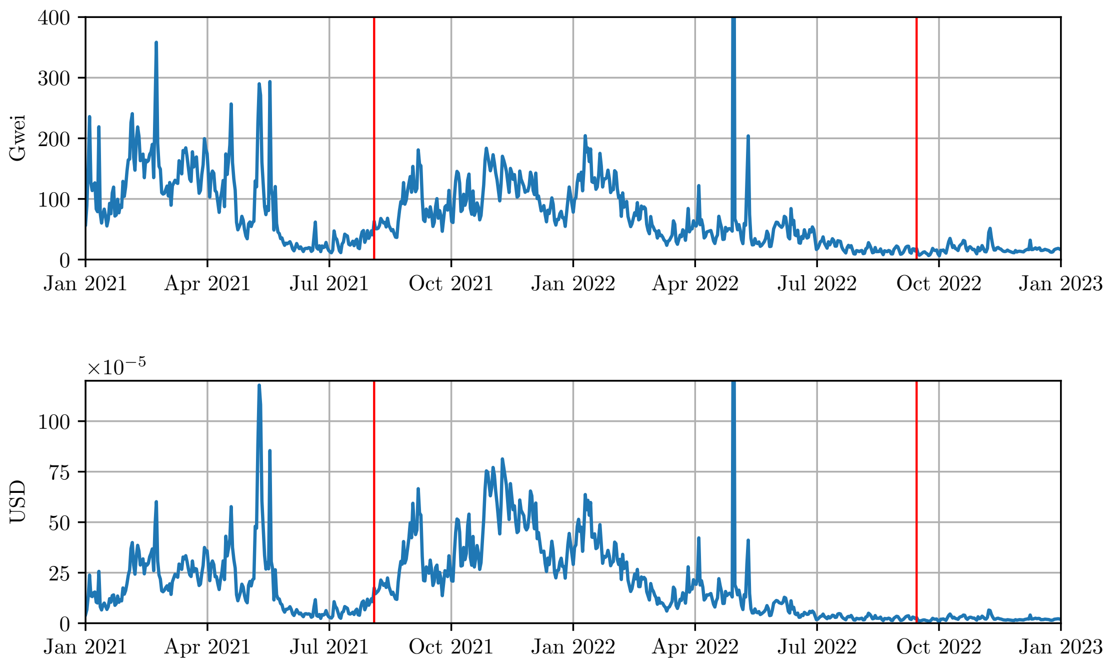

Notes: The Figure plots the daily average of the block median gas price over the studied period, measured in Gwei (10^{-9} ETH) in the top panel, and in US dollars (USD) in the bottom panel. The vertical red lines mark the August 5, 2021 London upgrade and the September 15, 2022 Merge upgrade. The May 1, 2022 'Bored Ape' NFT release caused a gas price spike to 797 Gwei (truncated in the Figure).

In Figure 2, we display the daily average *block utilization rate*, defined as the total gas used by all transactions included in a given block divided by the block gas limit. The block utilization rate proxies the demand for Ethereum services (in units of gas) relative to the available gas supply (the block gas limit) at a given time. The Figure shows that the block utilization rate underwent a substantial and sustained decrease and stabilization immediately after the London upgrade, as intended by the algorithmic fee determination in which deviations above and below

[7] One Gwei equals 10^{-9} Ether (ETH). 1 ETH was worth about $2,700 on August 5, 2021 and about $1,600 on September 15, 2022 (see Figure A1).

[8] On May 1, 2022 the Bored Ape Yacht Club released 100,000 non-fungible tokens (NFTs), a type of digital asset hosted on Ethereum, which caused a significant surge of transaction demand by crypto investors and drove the median gas price from 47 Gwei on April 30, 2022 to 797 Gwei on May 1, 2022.

50% utilization trigger upward or downward changes in the block base fee (see Appendix B1). Prior to the August 5, 2021 London upgrade most Ethereum blocks exhibited extremely high congestion, with utilization rates close to 100%. The block utilization dropped sharply to 53% after the London upgrade. Following the proof-of-stake (Merge) upgrade on Sep. 15, 2022, the block utilization rate dropped further and stabilized at around 51% on average until the end of our period of analysis.

Figure 2: Block Utilization

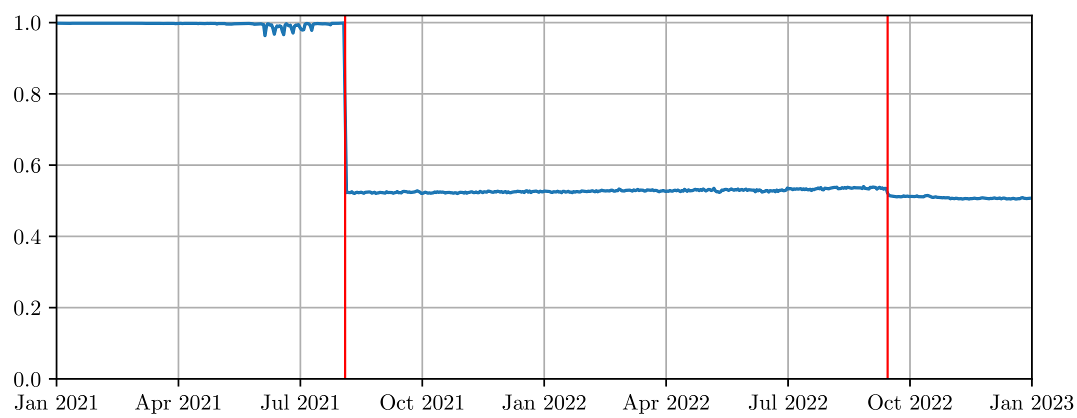

Notes: The Figure plots the daily average of the block utilization rate, defined as the total gas used by all transactions included in a block divided by the block gas limit (maximum capacity).

Figure 3 plots the daily average of the *block gas limit*, the algorithmically determined gas supply level, during our period of study. The sum of the gas requirements of all transactions included in a block cannot exceed the block gas limit. The transaction gas requirement and the gas limit are essential for the operation of Ethereum's Turing-complete virtual machine code and prevent potential issues such as infinite loops, coding errors, network disruptions, and more. Figure 3 shows that the gas limit was increased from 12.7 million gas per day to 14 million gas on April 21, 2021 and further to 15 million gas by the end of April 2021. It was then doubled to 30 million on August 5, 2021 (the London upgrade) and stayed at that level until the end of the sample period.

We use the observed variation of the Ethereum gas price (transaction fee per unit of complexity) over the studied period, including the major London and Merge platform upgrades, to identify and analyze its demand-side and supply-side determinants using a time-series regression model. Importantly for identification and for separating demand and supply factors, we use that changes in the gas supply in Ethereum are exogenous, as they arise from algorithmic adjustments to the block gas limit and/or the block creation rate.

We find that in the pre-London period, the median gas price tends to be higher on average when the platform experiences higher block utilization rate, which aligns with economic theory. After the London upgrade, when the predominant part (about 92%) of the gas price is the algorithmically

Figure 3: Gas Limit

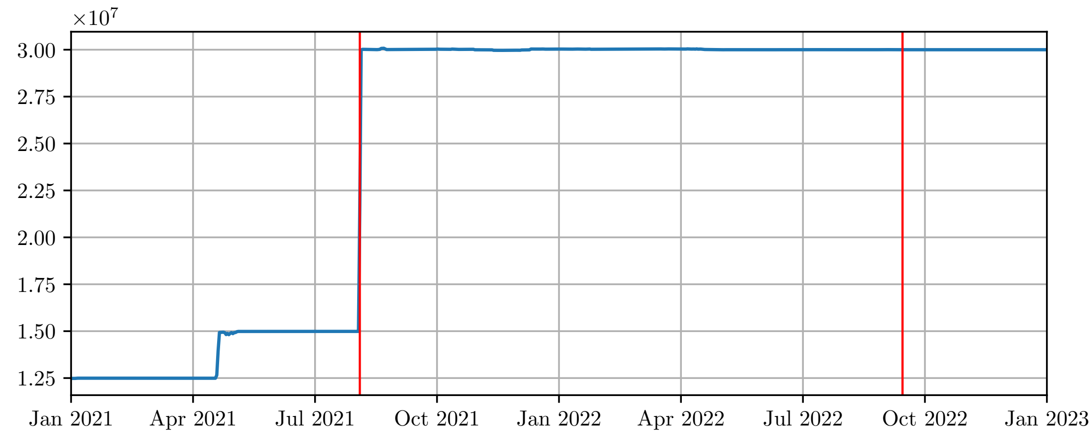

Notes: The Figure plots the daily average of the block gas limit over the studied period. The block gas limit is algorithmically determined and defines the total gas available to be used by all transactions included in a block.

determined block base fee, lagged block utilization plays a similar role, that is, a higher utilization rate in the previous block (which leads to a higher base fee in the current block, all else equal) is associated with higher gas prices (see Appendix B2). Second, we find that gas prices are statistically significantly positively associated with the proportion of contract calls in a block, both before and after the London upgrade, with a mitigated effect in the post-London period. Higher average gas prices for contract-call transactions suggest that contract calls may have higher time sensitivity or urgency. Third, we find that in the post-London period, users of legacy transactions pay higher gas prices on average compared to the users of non-legacy transactions. Fourth, the gas price decreases on average as the gas supply within each block (the gas limit) is increased. Overall, we find that the London and Merge upgrades resulted in a sizable reduction of median gas prices as of the end of 2022 and significantly stabilized the variability of Ethereum transaction fees over time.

Related literature

Donmez and Karaivanov (2022) use daily-averaged data to analyze the determinants of the Ethereum gas price and transaction fees in the period Nov. 2017 – Jan. 2019, prior to the 2021 London upgrade, and find that changes in demand are the key gas price determinant. Specifically, when there is high daily platform utilization the gas price increases on average, with a strong positive nonlinear effect above 90% utilization. The authors also find that a larger daily fraction of direct transfers between users is associated with a higher gas price on average.

Liu et al. (2022) study the effect of the London upgrade on Ethereum transaction fees and waiting times using a regression discontinuity approach. Similar to our findings, they document a decrease in the volatility of fees in the post-London period. However, our results further reveal a significant impact on the overall gas price level through changes in the gas price determination mechanism – specifically, the algorithmically determined base fee anchoring the gas prices – along with the importance of the block utilization rate as a measure of demand and the block gas limit (gas supply). Brown et al. (2021) empirically show that the Segwit (segregated witness) protocol, which increased the block size on the Bitcoin platform, reduces the congestion rate and transaction fees. Their findings indicate that congestion causes users to bid higher fees to have their transactions processed sooner, consistent with our results for Ethereum. The authors also find that higher Bitcoin price in USD and its volatility lead to higher transaction fees.

Jain et al. (2023) study the impact of network congestion, the six-month return of the platform currency, and the Merge upgrade on transaction fees in Bitcoin and Ethereum using daily-level data. The authors find that transaction fees in USD sharply increase with congestion (measured by the 'mempool' count of waiting transactions), and that the Merge upgrade led to lower transaction fees.[9] Our results, using block-level data and a different measure of congestion (the block utilization rate), characterize the dynamics of transaction fees across three different time periods (pre-London, London, and Merge) and demonstrate the effect of additional demand and algorithmic/supply factors on gas prices. Also using daily-level Ethereum data and first-difference quantile regressions, Koutmos (2023) shows that the number of daily transactions (a measure of demand or congestion) is the factor most consistently associated with transaction fees across all gas price quantiles, in contrast to other variables such as the ETH price or block size. This is consistent with our findings about the positive association between the block utilization rate (a more direct measure of congestion at the block level) and the Ethereum gas price.

Our paper contributes to this literature in several ways. First, unlike previous papers which use data aggregated at the daily level, we use granular block-level data constructed from the universe of more than 900 million transactions directly downloaded from the Ethereum blockchain over the two-year period January 1, 2021 to December 31, 2022. Using block-level data is especially important in the post-London upgrade period, since key gas price determinants such as the base fee and the block utilization rate are determined and vary at the block level. Second, we analyze the changes in the Ethereum gas price and its determinants preceding and following both major recent platform changes, the 2021 London upgrade and the 2022 Merge upgrade, thus spanning a longer time horizon than the papers which focus on the short-term effects of these upgrades. Third, we use time-series regressions and analyze the effect on the gas price of multiple economic factors, both from the demand and the supply side.

[9] See also Pierro and Rocha (2019) who examine the effect of the number of pending transactions, the ETH price in USD, and the number of miners, on Ethereum transaction fees in the pre-London period and conclude that the number of pending transactions and the number of miners significantly impact transaction fees compared to other variables.

2 Data

2.1 Data sources

We downloaded the complete block-level and transaction-level data directly from the Ethereum blockchain for the period January 1, 2021 to December 31, 2022. The block-level data include the block timestamp, block number, miner, size, gas limit (representing the block capacity), gas used (the total gas consumed by all transactions included in the block), difficulty level (pre-Merge), number of transactions, and the block base fee (after August 5, 2021).

The transaction-level data contain information about each transaction, including the transaction index within its block, a " nonce" (count of the number of previous transactions originated from the sender's address), the sender's address, the receiver's address, the transaction value in Wei (10^{-18} ETH), the maximum gas for the transaction, the gas price, the transaction type, the maximum fee per gas, the maximum priority fee per gas, and legacy (pre EIP-1559 format) or non-legacy indicator. We combine the blockchain data with hourly-level data on the USD price of Ether (ETH), the internal digital currency of Ethereum in which the transaction fees are paid, obtained from Cryptocompare.

The main data variables which we use in our empirical analysis are the gas price, the priority fee, the block base fee, the block gas limit, total gas used per block, the transaction type (one of: 'regular' transfer between two addresses, 'contract call', and 'contract creation'), and the ETH price in USD. Our sample period includes 4,657,807 blocks and 870,350,631 transactions in total. We aggregate the raw transaction data at the block level and perform all statistical analysis with block-level data.

2.2 Variables and definitions

Our main variable of interest is the median *block gas price*, defined as the median of the gas prices of all transactions recorded in a given block. We define the *block utilization rate* as the ratio of the total gas used by all transactions in a block to the *block gas limit*. The block utilization rate serves as an indicator of demand or congestion at the block level, i.e., at a specific moment of time. The block gas limit is a measure of gas supply and is mostly algorithmically fixed, with discrete increases in April 2021 and August 2021.

We construct two variables capturing the transaction type. We define the *block contract call share* as the ratio of all contract call transactions to the total number of transactions in a given block.[10] This variable captures the potential effect of changes in the composition of transactions

[10] A contract call is a type of transaction on the Ethereum blockchain that triggers the execution of a pre-defined

within a block, (e.g., more vs. less urgent) on the gas price.[11] We also define the *block legacy share* as the proportion of legacy transactions among all transactions in that block. This variable captures the possible effect of the EIP-1559 transaction format adoption rate on gas prices following the London update (all transactions are 'legacy' prior to August 5, 2021).

We also control for the algorithmically-determined *block gas limit*, as a measure of gas supply, and for the *ETH price in USD*, as a measure of the real cost of transactions. We match and merge the hourly data on the ETH price in USD with the block-level data using the blocks' timestamps.

Before the London upgrade on August 5, 2021, all users supplied (bid) a single value for the gas price (transaction fee per unit of complexity) for their transactions. Following the upgrade, the block base fee is automatically determined by the blockchain platform code and the users choose a non-negative *priority fee* (akin to a tip), as well as a maximum fee per gas, which must be greater or equal to the base fee in order for the transaction to be processed. Consequently, following the London upgrade, we analyze the economic demand and supply determinants of the priority fees chosen by the users. Table A1 displays summary statistics of the variables used in our analysis.

2.3 Descriptive findings

Figure 4 illustrates the gas price distribution in the three time periods we analyze: 'pre-London' (January 1, 2021 to August 4, 2021), 'London' (August 5, 2021 to September 14, 2022), and 'Merge' (September 15, 2022 to December 31, 2022), including separately for legacy and non-legacy transactions. We observe that the median block gas price decreases from approximately 79 Gwei (pre-London) to 55 Gwei after the EIP-1559 switch to base fee plus priority fee model (London), and further to 14 Gwei in the Merge period. In addition, the dispersion in gas prices is progressively reduced over the two-year sample period. Overall, these observations indicate a significant reduction in the level and variability of Ethereum transaction fees as of the end of 2022 compared to the two prior years. We also find that users who continue to submit legacy transactions (using the pre-London bid gas price format) after the London upgrade pay higher gas prices (both average and median) compared to the users who submit non-legacy (the new, London format) transactions – 65 Gwei vs. 54 Gwei in the London period and 15.7 Gwei vs. 14 Gwei in the Merge period. In the top-right panel of Figure 4 we display the transaction fee for a 'regular transaction' (ETH transfer between two addresses) in USD. The transaction fee rose on average in the London period, because of the higher price of ETH in USD (see Figure A1) but then decreased in the Merge period as both the ETH price in USD and the gas price fell.

function of a smart digital contract deployed on the platform, e.g., an exchange of token into ETH, etc.

[11] We similarly define 'regular share' and 'contract creation share'.

Figure 4: Gas prices by period

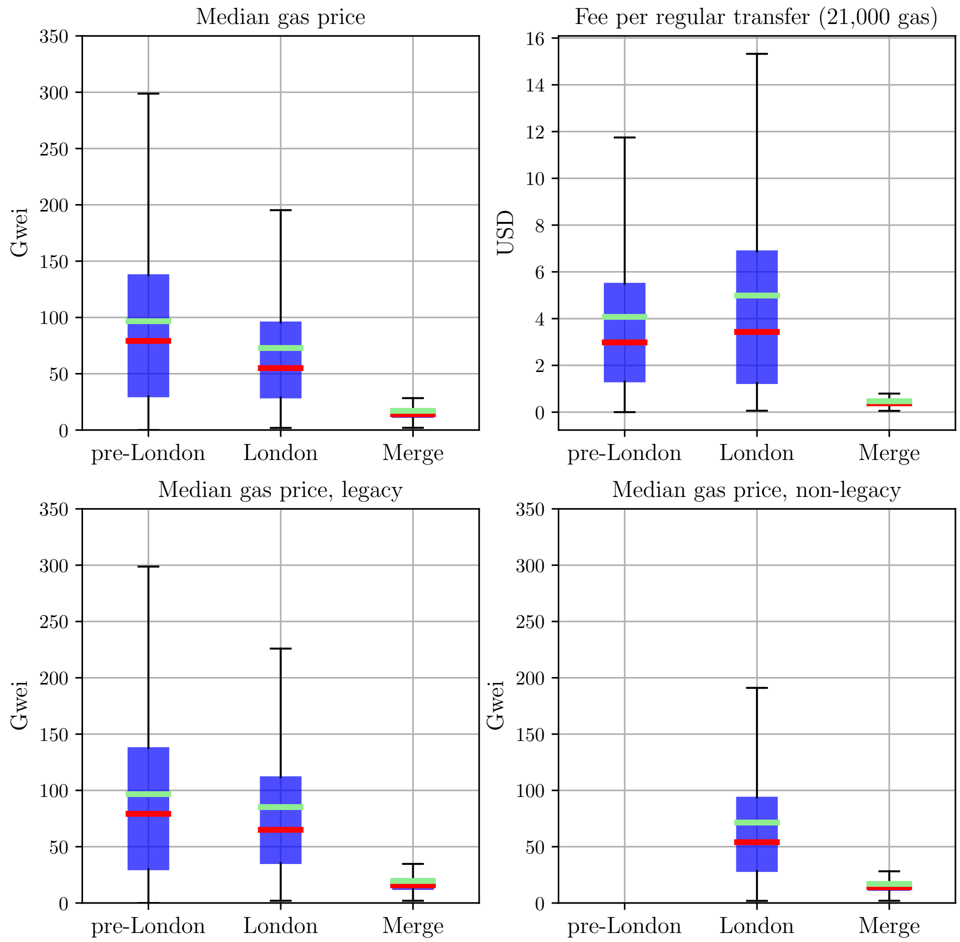

Notes: The top-left panel displays the distribution of the median block gas price in the pre-London, London, and Merge periods. The top-right panel displays the distribution of the US dollar transaction fee per regular transfer, computed by multiplying the median gas price by 21,000 and by the ETH price in USD. The bottom panels display the distribution of the median block gas price for legacy transactions (bottom left) and for non-legacy transactions (bottom right). In each panel the blue boxes display the interquartile range (IQR); the green and red lines denote respectively the mean and median of the plotted data variable; and the whiskers extend to all first and fourth quartile observations that are within 1.5 times the IQR.

Figure 5 displays the distributions of the block base fee, the block utilization rate, the daily block creation rate, and the ETH price in USD, before and after the London and Merge upgrades. Comparing the London and Merge sub-periods in the top left panel, we note a significant decrease in both the magnitude and the dispersion of the block base fee, which is the main factor determining the gas price level.[12] The block utilization rate (top right panel) experiences a substantial reduction after the London upgrade, from more than 99% on average in the pre-London period to 51% on average as of December 2022. The block utilization rate varies significantly around its average value, notably in the London period. We explore this variation in our empirical analysis in Section

[12]There is no algorithmic base fee in the pre-London period; instead the whole gas price is bid by the users.

Figure 5: Data Descriptives

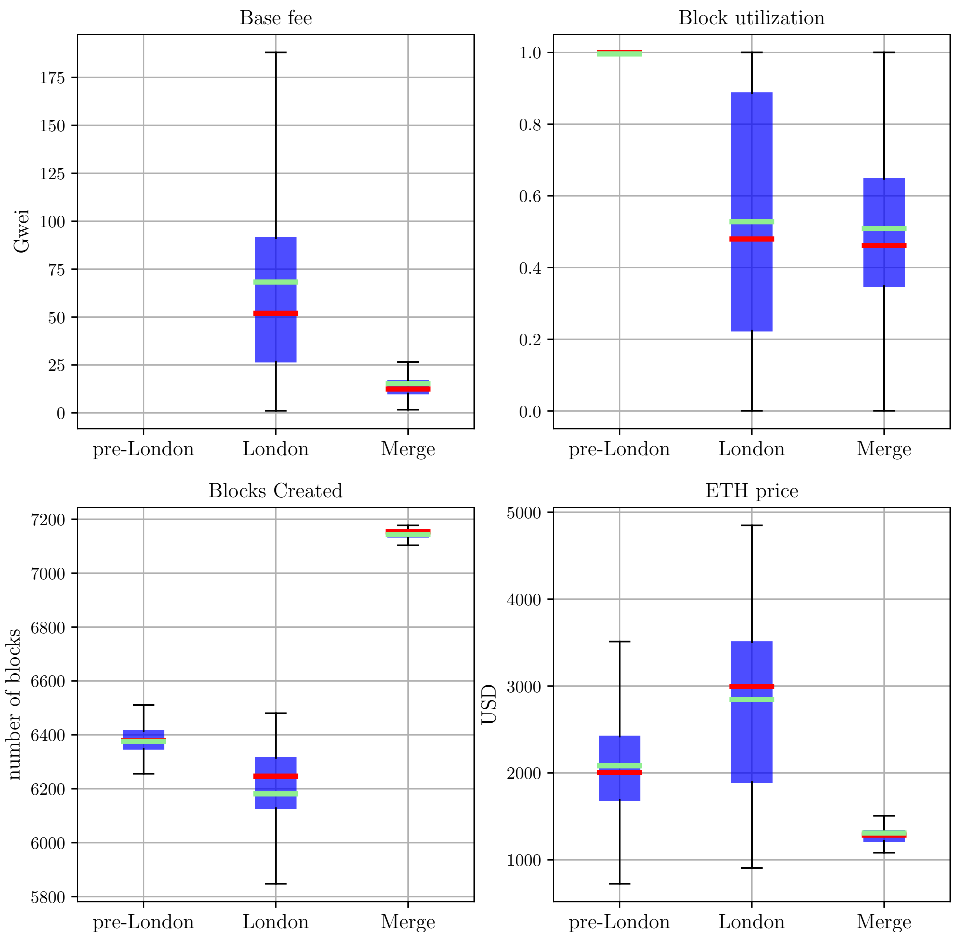

Notes: The Figure plots the distribution of the block base fee, block utilization rate, blocks created per day, and the ETH price in USD in the pre-London, London, and Merge periods. In each panel the blue boxes display the interquartile range (IQR); the green and red lines denote respectively the mean and median of the plotted data variable; and the whiskers extend to all first and fourth quartile observations that are within 1.5 times the IQR.

3. While the block gas limit was not changed in the Merge upgrade, the block creation rate (bottom left panel of Figure 5), defined as the number of new blocks created per day, went up from 6,305 blocks per day on average before the Merge to 7,153 blocks per day after the Merge, representing an increase in gas supply per unit of time. The ETH price in USD (bottom right panel) increased from 2,006 USD on average in the pre-London period to 2,994 USD in the London period and declined to 1,286 USD on average after the Merge upgrade.

Figure 6 displays the distribution of the share of contract calls and the share of legacy transactions out of all transactions in a block. We observe that, on average, the contract call share rises steadily from 66% prior to the London upgrade, to 72% following the Merge upgrade. As expected, the legacy transaction share falls significantly in the London and Merge periods, but still remains at about 18% of all transactions as of December 2022.

Figure 6: Contract Call Share and Legacy Share

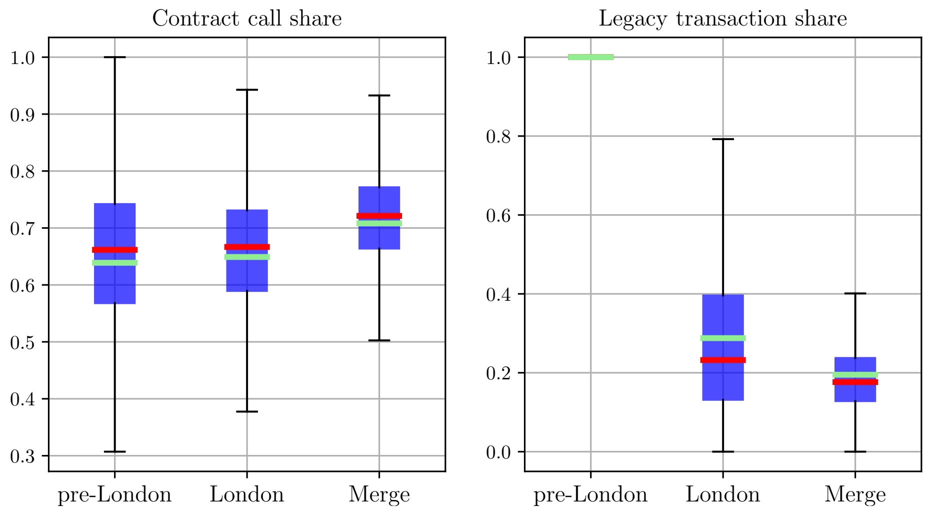

Notes: The Figure displays the distributions of the contract call share (left panel) and the legacy transaction share (right panel) in the pre-London, London, and Merge periods. In each panel the blue boxes display the interquartile range (IQR); the green and red lines denote respectively the mean and median of the plotted data variable; and the whiskers extend to all first and fourth quartile observations that are within 1.5 times the IQR.

Figure 7 illustrates the distribution of the priority fees per unit of gas submitted by the users, overall and for different transaction types. Priority fees did not exist in the pre-London period. We observe that the median block priority fee decreases from around 2 Gwei in the London period to 1.5 Gwei in the Merge period across all transactions, from 2 Gwei to 1.7 Gwei for regular (ETH transfer) transactions, from 2 Gwei to 1.5 Gwei for contract calls, and from 1.8 Gwei to 1.3 Gwei for contract creations.

Prior to the London upgrade, users bid the gas price they were willing to pay considering their transaction's complexity and gas requirements. If a miner selected a specific transaction to be included in the current block, the user would ultimately pay a total transaction fee calculated as the *gas price* bid by the user multiplied by the transaction's *gas requirement*. Following the London upgrade, users submitting non-legacy transactions can bid a priority fee, factoring in the mandatory base fee of the current block and transaction gas requirement. The sum of the priority fee and the algorithmically determined base fee forms the actual gas price paid by the user. Figures 4, 5, and 7 indicate that the London upgrade significantly helped users in predicting the going transaction fee / gas price level, thereby preventing both overpayment and underpayment. In addition, the Ethereum transaction fees decreased and stabilized over the studied period.

Figure 7: Priority fees

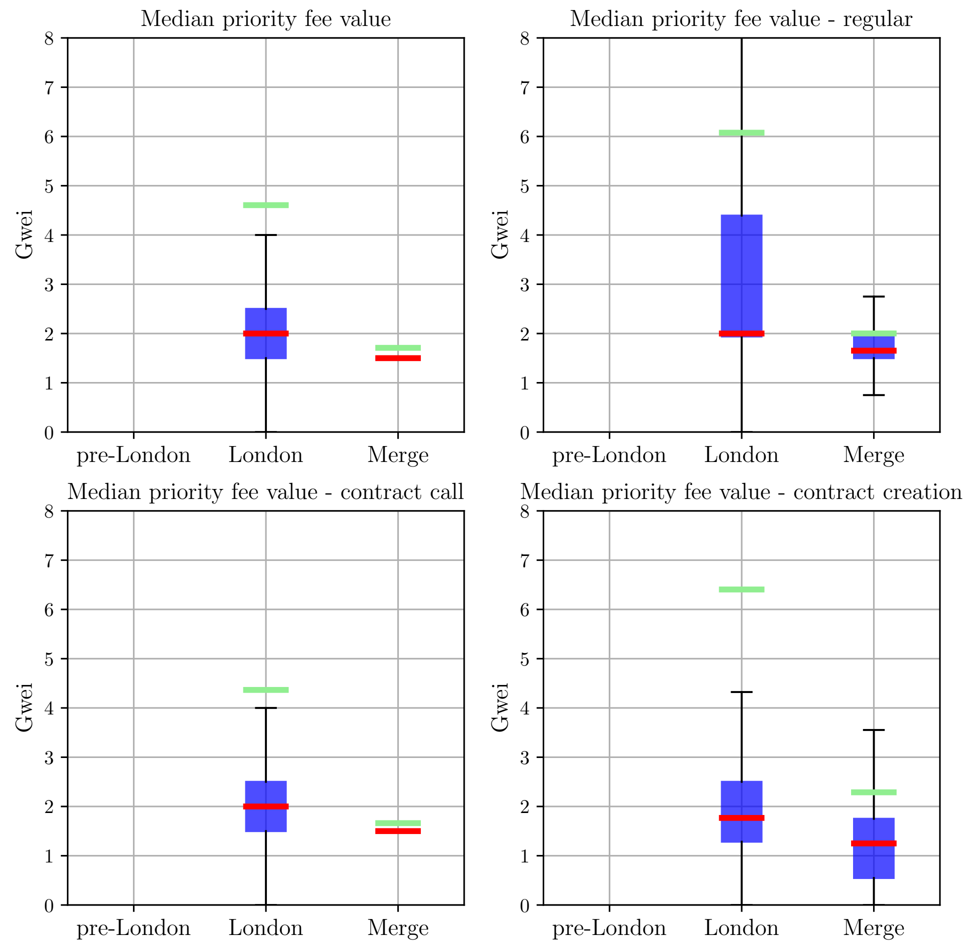

Notes: The Figure displays the distribution of priority fees in the pre-London, London, and Merge periods. The top left panel displays the median block priority fee; the top right panel displays the median block priority fee for regular transactions (direct ETH transfer between users); the bottom left panel displays the median block priority fee for contract calls, and the bottom right panel displays the median block priority fee for contract creations. In each panel the blue boxes display the interquartile range (IQR); the green and red lines denote respectively the mean and median of the plotted data variable; and the whiskers extend to all first and fourth quartile observations that are within 1.5 times the IQR.

3 Empirical analysis

3.1 Preliminaries

We analyze the demand and supply side determinants of Ethereum transaction fees before and after the London and Merge platform upgrades. On the demand side, in the pre-London period users bid the transaction fee per unit of computational complexity (the gas price) that they are willing to pay and bid the priority fee after the London upgrade. Transactions with higher gas prices are prioritized for execution since the miners or validators receive the fees for all user transactions included in the created block and hence maximize their profits by sorting submitted but not yet processed transactions in descending order of their bid gas price. When the platform utilization rate is high there is greater competition among users to have their transactions recorded in the current block, which implies a higher gas price or priority fee. Therefore, the predicted effect of the block utilization rate, as an indicator of demand conditions, on the gas price is positive.

On the supply side, the number and complexity of transactions that can be included in a block is constrained by the block gas limit (total supply of gas per block, see Figure 3), which is algorithmically set by the platform code and determines the service rate capacity. A higher block gas limit allows recording more transactions in a block, all else equal. Therefore, the predicted effect of the block gas limit size on the gas price is negative.

In addition to the block utilization rate and the block gas limit, we also consider and analyze the effect of other variables that can affect the demand for blockchain transactions and the resulting transaction fees (gas price). One such variable is the transaction type. We distinguish between: 'regular' transactions (a simple ETH transfer between two blockchain addresses), 'contract call' transactions (submitting data to or interacting with a smart digital contract deployed on the Ethereum platform), and 'contract creation' transactions (deploying a smart contract on the platform). Specifically, we study the effect of the share of contract calls among all transactions in a block as an indicator of transaction urgency. Our hypothesis is that contract calls are relatively more urgent on average as they often involve automated time-sensitive interactions with smart contracts, e.g. token purchases, trading and exchange. The predicted effect of the block contract call share on the gas price is therefore positive.

We also construct the share of 'legacy' (pre-London format) transactions in each block. Figure 4 suggests that users who continue to submit such transactions after the London upgrade (possibly because of not having updated their clients) tend to overpay in fees. Thus, we expect a positive association between the 'legacy share' and the gas price. Finally, since the users pay the gas price in ETH (the platform's internal digital currency), a higher ETH price in US dollars costs the users more in real monetary terms, which would suggest a negative association between the gas price and the ETH price in USD. On the other hand, however, a higher ETH price in USD may also indicate higher interest / demand for Ethereum transactions and thus higher gas prices. The effect of the ETH price in USD on the gas price is therefore undetermined ex-ante and will be estimated in the empirical analysis.

3.2 Estimation

We estimate two main empirical specifications. Prior to the August 2021 London upgrade, the gas price (transaction fee per unit of complexity) is the main choice variable for the users submitting transactions to the Ethereum platform. We therefore examine the economic factors affecting the users' bidding behavior and gas price choice by estimating the specification:

$$\log(g_t) = \beta_0 + \beta_1 u_t + \beta_2 l_t + \beta_3 c_t + \beta_4 x_t + \epsilon_t \tag{1}$$

where g_t denotes the median gas price in block t.[13] The variable u_t is the utilization rate in block t, defined as the ratio of total gas used by all transactions in the block to the block gas limit. The block utilization rate u_t captures the level of congestion or service demand at the time the block is created. The variable l_t is log of the algorithmically determined block gas limit and measures the available gas supply (capacity) at the block level. The variable c_t denotes the contract call share in block t, as defined in Section 2.2. It captures variation in the within-block composition of transaction types, i.e., the mix/ratio of higher vs. lower urgency transactions. Finally, x_t denotes other controls that we include, such as the hourly ETH price in USD and the block legacy transactions share.

After the London upgrade the gas price is constructed as the sum of an algorithmically determined base fee and a user-bid priority fee. The block base fee accounts for about 92% of the block gas price on average in our data and is a deterministic function of the utilization rate and the base fee of the immediately preceding block (see Appendix B). The user has no direct control over the base fee component of the gas price. Therefore, after the London upgrade, the users' main choice variable is the priority fee (tip) instead of the full gas price. We thus estimate the following specification for the post-London period:

$$\log(\pi_t) = \gamma_0 + \gamma_1 \log(b_t) + \gamma_2 u_t + \gamma_3 l_t + \gamma_4 c_t + \gamma_5 x_t + \nu_t \tag{2}$$

where π_t is the median priority fee in block t, b_t denotes the base fee in block t, and ν_t is the error term. The definitions of u_t, l_t and c_t are the same as in specification (1).

In the post-London period we also estimate a version of specification (1) using the first lag of the block utilization rate, u_{t-1} (see Table 2). The reason is the dependence of the algorithmically determined base fee (which constitutes the major part of the gas price in the post-London period) on the utilization rate of the preceding block (for details, refer to Appendix B2). Thus, by using

[13] To account for observations with zero median gas price (4,535 observations, or 0.3% of the pre-London period sample) we add 1 Wei to all block median gas prices before taking log.

the lagged block utilization rate we can account for the variation in the base fee resulting from past demand conditions.

3.3 Results

In Table 1 we present estimation results for the pre-London period, using specification (1). The dependent variable is the natural logarithm of the median gas price in each block, measured in Wei (10^{-18} ETH). We find that the block utilization rate (total gas used in the block divided by the block gas limit), which is a proxy for the demand for Ethereum services, is positively and statistically significantly associated with the median gas price, consistent with economic theory. Quantitatively, the estimated coefficient of 4.384 in Table 1, column (1), implies that a 1 percentage point increase in block utilization is associated with a 4.48% increase in the median gas price in Wei, calculated as $100(e^{0.01\times4.384}-1)$. Considering that the average value of the block median gas price in our sample is 96.6 Gwei (see Table A1), this implies that a user submitting a transaction with gas requirement (complexity) of 100,000 gas would pay, on average, $96.6\times(100,000)\times(4.48\%)=$ 432,768 Gwei (about 1 USD) more following a 0.01 increase in block utilization in the pre-London period.

We also find that the block gas limit, defined as the total gas supply per block (an algorithmically determined quantity) is negatively and statistically significantly associated with the median gas price. The estimate of -11.10 in column (1) of Table 1 indicates that a 1% increase in the block gas limit is associated with an 11.10% lower median gas price on average.

The block contract call share, defined as the fraction of all contract call transactions among all transactions in the block, has a positive and statistically significant coefficient estimate. This suggests that contract call transactions are associated with higher gas prices bid by the users, on average. The estimate of 1.347 in column (1) means that a 1 percentage point increase in the contract call share is associated with a 1.36% increase in the median gas price. The estimated coefficient on the ETH price in USD (the ETH/USD exchange rate) is also positive and statistically significant, consistent with the hypothesis that higher Ether prices may be associated with higher demand for blockchain transactions. For example, an increase in the dollar price of ETH may attract more traders or investors to Ethereum, leading to a higher transaction volume. The resulting increase in demand for block space would drive gas prices up, as the users compete by offering higher fees to have their transactions processed quicker.

To ensure robustness against possible non-stationarity issues, in Table 1 column (2) we also report estimation results from a specification in which all variables are the first differences of the respective block-level variables from column (1), excluding the ETH price in USD which is

Table 1 Pre-London Results, Gas Price
January 1, 2021 to August 4, 2021

	log median gas price, Wei	
	(1)	(2)
	levels	first differences
block utilization	4.384***	3.394***
	[0.127]	[0.138]
log block gas limit	-11.10***	-27.42***
	[0.059]	[2.496]
block contract call share	1.347***	1.835***
	[0.037]	[0.041]
log ETH price in USD	1.344***	n.a.
	[0.018]	–
adj. R-squared	0.187	0.032
sample size	1,377,279	1,377,278

p-values: $^{*}\,p < 0.1$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$

Notes: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets (Newey and West (1987)). In column (1) the dependent variable 'log median gas price' is the natural logarithm of the median gas price in Wei (10^{-18} ETH) of all transactions within a block; 'block utilization' is the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log block gas limit' is the natural logarithm of the block gas limit; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions in that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD. In column (2) the dependent and independent variables are the first differences of the respective block-level variables from column (1). To account for the observations with zero median gas price (4,535 observations or 0.3% of the sample) we add 1 Wei to all block median gas prices before taking logs.

measured hourly.[14] Our results from column (1) on the relationship between block utilization, the block gas limit, the contract call share, and the median gas price in the pre-London period remain robust.

In Table 2 we estimate a modified version of specification (1) using data from the time period after the London upgrade. We divide this period into two: "London period", spanning from August 6, 2021 to September 14, 2022 and "Merge period", from September 16, 2022 to December 31, 2022.[15] According to the algorithmic base fee formula (see Appendix B for details), the base fee of block t is determined by the utilization rate and the base fee of the immediately preceding block, $t-1$. Since the base fee is very persistent and makes up for about 92% of the gas price on average, in our post-London analysis in Table 2, instead of the time-t block utilization rate which does not affect the time-t base fee, we use the first lag (time $t-1$) of block utilization

[14]We also performed the augmented Dickey-Fuller (ADF) test (Dickey and Fuller (1979)) for stationarity with the time series variables used in Tables 1 and 2. We reject the null hypothesis that a unit root is present for all series at the 95% or higher confidence level, except for the ETH price in USD in the post-London period.

[15]We omit from the sample the date of the London upgrade, August 5, 2021 and the date of the Merge upgrade, September 15, 2022 since there is a mix of regimes on these dates.

Table 2 Post-London Results, Gas Price
August 6, 2021 – December 31, 2022

	log median gas price, Wei			
	London		London + Merge	
	(1)	(2)	(3)	(4)
	levels	first differences	levels	first differences
L.block utilization	0.108***	0.131***	0.110***	0.129***
	[0.001]	[0.0001]	[0.001]	[0.0001]
log block gas limit	-8.875***	-4.247***	-7.979***	-4.294***
	[0.987]	[0.078]	[1.008]	[0.078]
block contract call share	0.174***	-0.026***	0.150***	-0.025***
	[0.009]	[0.0006]	[0.009]	[0.0007]
log ETH price in USD	1.410***	n.a.	1.379***	n.a.
	[0.007]	–	[0.0077]	–
block legacy share	0.084***	0.162***	0.095***	0.149***
	[0.008]	[0.0007]	[0.008]	[0.0008]
adj. R-squared	0.418	0.292	0.557	0.281
sample size	2,503,393	2,503,392	3,267,660	3,267,659

p-values: $^{*}\,p < 0.1$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$

Note: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In columns (1) and (3) the dependent variable 'log median gas price' is the natural logarithm of the median gas price in Wei (10^{-18} ETH) of all transactions within a block; 'L.block utilization' is the first lag of block utilization, defined as the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log block gas limit' is the natural logarithm of the block gas limit; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions within that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD; 'block legacy share' is the number of transactions in a block using the pre-London legacy (bid gas price) format divided by the total number of transactions in that block. In columns (2) and (4), the dependent and independent variables are the first differences of the respective block-level variables from columns (1) and (3). To control for the algorithmic increase in the block creation rate in the Merge period, in columns (3) and (4) we also include a dummy variable which equals 0 before Sep. 15, 2022 and 1 from Sep. 15 onward. Columns (1) and (2) use data from the period Aug. 6, 2021 to Sep. 14, 2022; columns (3) and (4) use data from the period Aug. 6, 2021 to Dec. 31, 2022.

(L.block utilization) in equation (1) to capture the impact of demand conditions on the gas price. In Table 3 we additionally estimate the impact of the time-t block utilization rate on the priority fee component of the gas price.

In Table 2, column (1) we find that the (lagged) block utilization rate is positively and statistically significantly associated with the median gas price in the post-London period. The estimate of 0.108 in column (1) implies that a 1 percentage point increase in the previous block's utilization is associated with a 0.11% higher median gas price. A larger block gas limit (larger gas supply) is associated with a lower median gas price, in line with economic theory and with the pre-London results in Table 1. A larger share of contract call transactions and larger ETH price in USD are each associated with a higher gas price in columns (1) and (3), as in our pre-London results.

In Table 2 we additionally control for the 'block legacy share', defined as the fraction of legacy (pre-London format) transactions in the block. This captures the post-London change in the composition of legacy format (bid gas price) vs. new format (base fee plus priority fee) transactions submitted by users. We find that a larger legacy share is associated with a higher median gas price on average, consistent with our findings in Figure 4 that legacy users tend to submit higher gas prices in the post-London period compared to non-legacy users. The estimate of 0.084 in column (1) suggests that a 1 percentage point increase in the block legacy share is associated with a 0.084% larger median gas price in Wei. Taking into account that the average block median gas price in this period is 72.8 Gwei (see Table A1), a user executing a transaction with complexity 100,000 gas would pay on average $72.8 \times (100,000) \times (0.084\%) =$ 6,115 Gwei (about 1.7 US cents) more when the block legacy share is higher by 0.01.

In Table 2, columns (2) and (4), we also report estimation results for a first-differences specification of our empirical model. Our main results remain robust, with the only exception being the estimate of the block contract call share which turns negative in the first-differenced specification. In columns (3) and (4) of Table 2 (the combined London+Merge period) we additionally control for the algorithmic increase in the block creation rate after the Merge upgrade (see Figure 5) by including a dummy variable that equals 0 before Sep. 15, 2022 and 1 from Sep. 15 onward.

In Table 3, we present the results from estimating equation (2) for the post-London period, using the priority fee component of the gas price as the dependent variable.[16] This specification isolates the current demand conditions driven user-bid effect on the gas price. In columns (1) and (3) we find that a larger block utilization rate is associated with a higher priority fee on average, consistent with our earlier results but with lower magnitude. The estimate of 0.094 in column (1) implies that a 1 percentage point increase in the block utilization rate is associated with a 0.094% increase in the median priority fee in Wei, holding the block base fee constant. Given that the average value for the block median priority fee in the 'London' period is 4.6 Gwei (see Table A1), a user executing a transaction with complexity 100,000 gas would pay on average $4.6 \times (100,000) \times (0.094\%) = 432.4$ Gwei higher priority fee following a 0.01 increase in the block utilization rate. Comparing to our pre-London results in Table 1, this suggests a lower sensitivity of the gas price (holding the base fee constant) to the current block utilization rate in the post-London period.

The coefficient estimate for the block gas limit in Table 3 is statistically significantly negative and large in magnitude in all specifications – for example, the value -20.83 in column (1) means that a 1% increase in the block gas limit is associated with a 20.83% larger median priority fee

[16]The priority fee is directly recorded in the blockchain data for all non-legacy transactions. For the legacy (pre-London format) transactions, we compute the implied priority fee as the difference *gas price* - *base fee*, where *gas price* is the user-bid gas price and *base fee* is the base fee of the block in which the transaction is recorded.

Table 3 Post-London Results, Priority Fee
August 6, 2021 to December 31, 2022

	log median priority fee, Wei			
	London		London + Merge	
	(1)	(2)	(3)	(4)
	levels	first differences	levels	first differences
block utilization	0.094***	-0.002	0.066***	-0.032***
	[0.003]	[0.004]	[0.003]	[0.003]
log block gas limit	-20.83***	-33.91***	-22.83***	-34.69***
	[1.012]	[1.755]	[1.025]	[1.793]
block contract call share	2.119***	2.601***	1.837***	2.306***
	[0.037]	[0.044]	[0.032]	[0.040]
log ETH price in USD	-0.282***	n.a.	-0.216***	n.a.
	[0.006]	–	[0.005]	–
block legacy share	2.631***	3.865***	2.429***	3.467***
	[0.011]	[0.024]	[0.011]	[0.024]
log block base fee	0.177***	1.493***	0.159***	1.242***
	[0.003]	[0.025]	[0.002]	[0.023]
adj. R-squared	0.170	0.165	0.165	0.145
sample size	2,503,394	2,503,393	3,267,661	3,267,660

p-values: $^{*}\, p < 0.1$, $^{**}\, p < 0.05$, $^{***}\, p < 0.01$

Note: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In columns (1) and (3) the dependent variable 'log median priority fee' is the natural logarithm of the median priority fee in Wei (10^{-18} ETH) among all transactions within a block (imputed for legacy transactions as the difference between the gas price and the block base fee); 'block utilization' is the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log block gas limit' is the natural logarithm of the block gas limit; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions in that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD; 'block legacy share' is the number of legacy transactions in a block divided by the total number of transactions in that block. 'log block base fee' is the natural logarithm of the base fee on a given block. In columns (2) and (4), the dependent and independent variables are the first differences of the respective block-level variables in columns (1) and (3). To control for the algorithmic increase in the block creation rate in the Merge period, in columns (3) and (4) we include a dummy variable which equals 0 before Sep. 15, 2022 and 1 for Sep. 15 onward. Columns (1) and (2) use data from Aug. 6^{th}, 2021 to Sep. 14^{th}, 2022 while columns (3) and (4) use data from Aug. 6^{th}, 2021 to Dec. 31^{st}, 2022. To account for the observations with zero median priority fee (9,289 observations or 0.3% of the sample) we add 1 Wei to all median priority fees before taking logs.

paid by users. The block contract call share is positively and statistically significantly associated with the median priority fee on average, as in Tables 1 and 2. The estimate of 2.119 suggests that a 0.01 larger contract call transaction share is associated with a 2.14% higher priority fee on average.

We also find that the block median priority fee is negatively and statistically significantly associated with the ETH price in USD, holding all else constant and controlling for the block base fee. A higher ETH price in USD implies larger real costs of paying the transaction fee, the major share of which is the mandatory base fee. Our results suggest that users tend to offer lower priority fees (smaller tip) on average when the USD price of ETH is higher. For example, the estimate of -0.282 in column (1) of Table 3 implies that a 1% increase in the ETH dollar price is associated with a 0.3% lower priority fee on average.

As in our Table 2 results for the gas price in the post-London period, we find that a larger block legacy share is associated with higher priority fees on average and with larger coefficient estimates (higher elasticity). We also find that a higher base fee, which reflects higher demand in the immediately preceding blocks, is associated with higher priority fees on average, in line with the theoretical expectations. The first-difference specifications in columns (2) and (4) of Table 3 are consistent with the results in levels from columns (1) and (3), with one exception of an opposite sign for block utilization in column (4).

4 Robustness analysis

4.1 Alternative supply measure – daily gas supply

In Table A2, for the pre-London period, and in Tables A3 and A4, for the post-London period, we estimate a variant of our main specifications in Tables 1, 2, and 3 where instead of using the block gas limit we use the *daily gas supply*, defined as the sum of the gas limits of all blocks recorded on a given calendar day. This specification accounts for variations in the gas supply on a daily scale by capturing both the impact of the block gas limit (gas supply per block), as in the baseline tables, and in addition the *number of blocks created* per day. Our main results from Tables 1–3 remain robust and we find that the daily gas supply has the expected negative association with the gas price or the priority fee.

4.2 Alternative accounting for zeros

The median bid gas price in the pre-London period equals 0 for 0.3% of the observations in our sample. Since the dependent variable in our empirical specification is log of the gas price in Wei, in our baseline regressions in Table 1 we add 1 Wei (10^{-18} ETH) to all gas prices (in comparison the median gas price over the pre-London period is 79.2 Gwei = 79.2bln Wei, see Table A1), so that we retain these observations in our sample and keep them as close to the raw data as possible. Alternatively, in Table A5 we deal with the zero gas price observations in the pre-London period in a different way, by replacing log(0) with zero. Our main results from Table 1 regarding the positive association between the block median gas price and the block utilization, the block contract call share and the ETH price in USD and the negative association between the block median gas price and the block gas limit remain essentially identical. In Table A6, we perform the same robustness check for the priority fee model in the post-London period where 0.3% of the priority fee observations equal zero. Our baseline results from Table 3 remain unchanged.

4.3 Alternative dependent variable – 5th percentile of the gas price

In Tables A7 and A8 we re-estimate our baseline specifications from Table 1 and Table 2 by using the 5th percentile of the block gas price as the dependent variable instead of the median block gas price. This robustness exercise provides insight into how low-cost transactions in a block are priced, given the demand and supply determinants we analyze. Using the 5th percentile gas price can also help us understand at what price point transactions are still being included at the margin across different demand states and shed light on miner behavior in processing transactions with relatively low gas prices.

The estimates in Table A7, for the pre-London period using the 5th percentile block gas price as the dependent variable have the same signs and are similar to our baseline estimates when using the block median gas price in Table 1, with slightly larger estimated coefficients. In Table A8, we perform the same robustness exercise for the post-London gas price model. Our baseline results from Table 2 remain robust, except for the estimate of the block legacy share, which is negatively associated with the 5th percentile gas price in columns (1) and (3). This could be due to users of legacy transactions underestimating the minimum transaction fee required, however, this result is not robust, as seen in the first-differenced specifications (2) and (4).

5 Conclusions

The Ethereum blockchain recently underwent two significant code updates (the 'London' and 'Merge' upgrades), which represent a pivotal shift in the platform's economic model of transaction fee determination and operational dynamics. We analyze the economic implications of these platform upgrades by focusing on the Ethereum gas price, that is, transaction fee per unit of computational complexity. We use a large dataset of block and transaction level data directly downloaded from the Ethereum blockchain and spanning two years and nearly 900 million transactions. We consider both demand-driven factors such as the utilization/congestion rate and the transaction type composition, as well as supply factors, such as the block gas limit and block creation rate.

We document the impact of the London upgrade on the gas price determination mechanism through the introduction of a base fee plus priority fee model which significantly reduced block congestion and stabilized the transaction fees paid by users. These effects were further enhanced by the Merge upgrade, a shift from proof-of-work to a proof-of-stake consensus mechanism, which drastically reduced the platform's energy consumption and enhanced its transaction processing efficiency.

Our empirical analysis reveals key economic determinants influencing the Ethereum gas price. Prior to the August 2021 London upgrade, when the gas price was fully determined by users' bids/offers, we document very high block utilization rates (nearly 100%) and elevated gas prices. These high utilization rates and prices were mitigated after the London upgrade through the algorithmic determination of the base fee which rises when utilization is high and falls when utilization is low, and an increase in the block gas limit. In the post-London period we show that the variability of gas prices substantially decreased and there was an overall reduction in the gas price, affirming the upgrades' effectiveness in addressing the pre-London platform congestion and inefficiencies.

Using time-series regression analysis, we find that the block utilization rate is significantly and positively associated with gas prices prior to the London upgrade. After the London upgrade, the first lag of block utilization has a similar effect, attributable to the algorithmic determination of the base fee which is a direct function of the prior-block utilization rate. We also analyze the impact of the within-block transaction type composition on gas prices and find that a larger fraction of contract call transactions is associated with higher gas prices, suggesting that contract call transactions may be more urgent on average. Additionally, we find that 'legacy transactions', that is, transactions that continue using the pre-London gas price bidding format, pay higher gas prices on average in the post-London period. Consistent with economic theory, a higher block gas limit, reflecting increased gas supply, is statistically significantly associated with lower gas prices in both the pre-London and post-London periods.

We contribute to the growing economics literature on blockchain platforms by providing a granular, block-level data analysis of Ethereum's transaction fee dynamics and the underlying supply and demand factors affecting them. The transition toward a more predictable, stable, and efficient transaction fee mechanism, coupled with the drastic reduction in electric energy consumption after the London and Merge upgrades enhances Ethereum's scalability and sustainability and sets a leading example for future blockchain and digital platform innovations.

References

Brown, C., Chiu, J. and Koeppl, T. (2021), 'What drives Bitcoin fees? Using SegWit to assess Bitcoin's long-run sustainability', *Journal of Financial Market Infrastructures* **9**(4).

De Vries, A. (2023), 'Cryptocurrencies on the road to sustainability: Ethereum paving the way for Bitcoin', *Patterns* **4**(1).

Dickey, D. and Fuller, W. (1979), 'Distribution of the estimators for autoregressive time series with a unit root', *Journal of the American Statistical Association* **74**(366), 427–431.

Donmez, A. and Karaivanov, A. (2022), 'Transaction fee economics in the Ethereum blockchain', *Economic Inquiry* **60**(1), 265–292.

Jain, A., Jain, C. and Krystyniak, K. (2023), 'Blockchain transaction fee and Ethereum Merge', *Finance Research Letters* **58C**(104507).

Kapengut, E. and Mizrach, B. (2023), 'An event study of the Ethereum transition to proof-of-stake', *Commodities* **2**(2), 96–110.

Koutmos, D. (2023), 'Network activity and Ethereum gas prices', *Journal of Risk and Financial Management* **16**(10).

Liu, Y., Lu, Y., Nayak, K., Zhang, F., Zhang, L. and Zhao, Y. (2022), 'Empirical analysis of EIP-1559: Transaction fees, waiting times, and consensus security', *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security* pp. 2099–2113.

Nakamoto, S. (2008), 'Bitcoin: A peer-to-peer electronic cash system', available at: https://bitcoin.org/bitcoin.pdf .

Newey, W. and West, K. (1987), 'A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix', *Econometrica* **55**(3), 703–708.

Pierro, G. and Rocha, H. (2019), 'The influence factors on Ethereum transaction fees', *IEEE/ACM Second International Workshop on Emerging Trends in Software Engineering for Blockchain* pp. 24–31.

Appendix A – Additional tables and figures

Table A1 Summary statistics

	Full sample, January 1, 2021 – December 31, 2022					
Variables	Obs.	Mean	Std	25^{th} pctile	50^{th} pctile	75^{th} pctile
Median gas price, Gwei	4,657,807	70.6	115.8	19	47	99
Base fee, Gwei	3,277,292	55.8	105.3	14.6	37	77.3
Block utilization	4,657,807	0.66	0.34	0.35	0.74	0.99
Contract call share	4,657,807	0.65	0.13	0.59	0.67	0.74
Legacy share	4,657,807	0.48	0.37	0.16	0.31	1.00
ETH price, 10^3 USD	4,657,807	2.36	0.99	1.51	2.13	3.11
Gas limit, 10^6 gas/block	4,657,807	25.1	7.47	15	30	30
	Pre-London period, January 1, 2021 – August 4, 2021					
	Obs.	Mean	Std	25^{th} pctile	50^{th} pctile	75^{th} pctile
Median gas price, Gwei	1,377,279	96.6	90.4	30	79.2	137.5
Base fee, Gwei	n.a.	-	-	-	-	-
Block utilization	1,377,279	0.99	0.03	0.99	0.99	0.99
Contract call share	1,377,279	0.63	0.15	0.56	0.66	0.74
Legacy share	1,377,279	1.00	0.00	1.00	1.00	1.00
ETH price, 10^3 USD	1,377,279	2.08	0.06	1.68	2.00	2.41
Gas limit, 10^6 gas/block	1,377,279	13.7	1.24	12.4	12.5	14.9
	London period, August 6, 2021 – September 14, 2022					
	Obs.	Mean	Std	25^{th} pctile	50^{th} pctile	75^{th} pctile
Median gas price, Gwei	2,503,394	72.8	138.1	29	55	95.5
Median priority fee, Gwei	2,503,394	4.6	63.6	1.5	2	2.5
Base fee, Gwei	2,503,394	68.2	117.4	26.7	51.9	91.2
Block utilization	2,503,394	0.52	0.33	0.22	0.48	0.88
Contract call share	2,503,394	0.64	0.13	0.58	0.66	0.73
Legacy share	2,503,394	0.28	0.21	0.13	0.23	0.39
ETH price, 10^3 USD	2,503,394	2.84	0.98	1.89	2.99	3.50
Gas limit, 10^6 gas/block	2,503,394	30	0.04	30	30	30
	Merge period, September 16, 2022 – December 31, 2022					
	Obs.	Mean	Std	25^{th} pctile	50^{th} pctile	75^{th} pctile
Median gas price, Gwei	764,267	17	18.5	11.6	14	18.3
Median priority fee, Gwei	764,267	1.7	12.5	1.5	1.5	1.5
Base fee, Gwei	764,267	15.3	13.3	10.1	12.4	16.6
Block utilization	764,267	0.50	0.24	0.34	0.46	0.64
Contract call share	764,267	0.70	0.10	0.66	0.72	0.77
Legacy share	764,267	0.19	0.10	0.12	0.17	0.23
ETH price, 10^3 USD	764,267	1.31	0.12	1.21	1.28	1.33
Gas limit, 10^6 gas/block	764,267	30	0.001	30	30	30

Note: Summary statistics aggregated to the block level. The base fee was introduced on August 5, 2021; before that, the entire gas price was bid by users.

Table A2 Pre-London, Gas Price – Robustness (using daily gas supply)
(January 1, 2021 to August 4, 2021)

	log median gas price, Wei	
	(1)	(2)
	levels	first differences
block utilization	4.341***	3.383***
	[0.127]	[0.139]
log daily gas supply	-11.74***	n.a.
	[0.064]	–
block contract call share	1.319***	1.837***
	[0.037]	[0.041]
log ETH price in USD	1.351***	n.a.
	[0.018]	–
adj. R-squared	0.183	0.032
sample size	1,377,279	1,377,278

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. The dependent variable 'log median gas price' is the natural logarithm of the median gas price in Wei (10^{-18} ETH) of all transactions within a block; 'block utilization' is the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log daily gas supply' is the natural logarithm of the gas limit of all blocks recorded on a given day; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions in that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD. To account for the observations with zero median gas price (4,535 observations or 0.3% of the sample) we add 1 Wei to all block median gas prices before taking logs.

Table A3 Post-London, Gas Price – Robustness (using daily gas supply)
(August 6, 2021 – December 31, 2022)

	log median gas price, Wei			
	London		London + Merge	
	(1) levels	(2) first differences	(3) levels	(4) first differences
L.block utilization	0.111***	0.131***	0.114***	0.129***
	[0.001]	[0.0001]	[0.001]	[0.0001]
log daily gas supply	-4.597***	n.a.	-4.302***	n.a.
	[0.086]	–	[0.089]	–
block contract call share	0.179***	-0.025***	0.153***	-0.025***
	[0.009]	[0.0006]	[0.009]	[0.0007]
log ETH price in USD	1.589***	n.a.	1.543***	n.a.
	[0.007]	–	[0.007]	–
block legacy share	0.100***	0.163***	0.112***	0.149***
	[0.008]	[0.0008]	[0.007]	[0.0008]
adj. R-squared	0.443	0.291	0.571	0.280
sample size	2,503,393	2,503,392	3,267,660	3,267,659

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In columns (1) and (3) the dependent variable 'log median gas price' is the natural logarithm of the median gas price in Wei (10^{-18} ETH) of all transactions within a block; 'L.block utilization' is the first lag of block utilization, defined as the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log daily gas supply' is the natural logarithm of the total gas supplied on the day block t is recorded; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions within that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD; 'block legacy share' is the number of transactions in a block using the pre-London legacy (bid gas price) format divided by the total number of transactions in that block. In columns (2) and (4), the dependent and independent variables are the first differences of the respective block-level variables from columns (1) and (3). To control for the algorithmic increase in the block creation rate in the Merge period, in columns (3) and (4) we also include a dummy variable which equals 0 before Sep. 15, 2022 and 1 from Sep. 15 onward. Columns (1) and (2) use data from the period Aug. 6, 2021 to Sep. 14, 2022; columns (3) and (4) use data from the period Aug. 6, 2021 to Dec. 31, 2022.

Table A4 Post-London, Priority Fee – Robustness (using daily gas supply)
(August 6, 2021 to December 31, 2022)

	log median priority fee, Wei			
	London		London + Merge	
	(1)	(2)	(3)	(4)
block utilization	0.092***	-0.009***	0.064***	-0.038***
	[0.003]	[0.004]	[0.003]	[0.003]
log daily gas supply	-0.215***	n.a.	-0.346***	n.a.
	[0.046]	–	[0.044]	–
block contract call share	2.119***	2.604***	1.837***	2.309***
	[0.037]	[0.044]	[0.032]	[0.040]
log ETH price in USD	-0.281***	n.a.	-0.209***	n.a.
	[0.007]	–	[0.006]	–
block legacy share	2.642***	3.866***	2.440***	3.468***
	[0.011]	[0.024]	[0.011]	[0.024]
log block base fee	0.175***	1.470***	0.157***	1.222***
	[0.003]	[0.025]	[0.003]	[0.023]
adj. R-squared	0.170	0.165	0.165	0.145
sample size	2,503,394	2,503,393	3,267,661	3,267,660

P-values: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In columns (1) and (3) the dependent variable 'log median priority fee' is the natural logarithm of the median priority fee in Wei (10^{-18} ETH) among all transactions within a block (imputed for legacy transactions as the difference between the gas price and the block base fee); 'block utilization' is the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log daily gas supply' is the natural logarithm of the total gas supplied on the day block t is recorded; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions in that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD; 'block legacy share' is the number of legacy transactions in a block divided by the total number of transactions in that block. 'log block base fee' is the natural logarithm of the base fee on a given block. In columns (2) and (4), the dependent and independent variables are the first differences of the respective block-level variables in columns (1) and (3). To control for the algorithmic increase in the block creation rate in the Merge period, in columns (3) and (4) we include a dummy variable which equals 0 before Sep. 15, 2022 and 1 for Sep. 15 onward. Columns (1) and (2) use data from Aug. 6^{th}, 2021 to Sep. 14^{th}, 2022 while columns (3) and (4) use data from Aug. 6^{th}, 2021 to Dec. 31^{st}, 2022. To account for the observations with zero median priority fee (9,289 observations or 0.3% of the sample) we add 1 Wei to all median priority fees before taking logs.

Table A5 Pre-London, Gas Price – Robustness (replace log(0) with 0)
(January 1, 2021 to August 4, 2021)

	log median gas price, Wei	
	(1)	(2)
	levels	first differences
block utilization	4.384***	3.395***
	[0.127]	[0.139]
log block gas limit	-11.10***	-27.44***
	[0.060]	[2.497]
block contract call share	1.347***	1.835***
	[0.037]	[0.041]
log ETH price in USD	1.344***	n.a.
	[0.018]	–
adj. R-squared	0.187	0.032
sample size	1,377,279	1,377,278

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In column (1) the dependent variable 'log median gas price' is the natural logarithm of the median gas price in Wei (10^{-18} ETH) of all transactions within a block; 'block utilization' is the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log block gas limit' is the natural logarithm of the block gas limit; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions in that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD. In column (2) the dependent and independent variables are the first differences of the respective block-level variables from column (1). To account for the observations with zero median gas price (4,535 observations or 0.3% of the sample) we replace the natural logarithm of zero with zero.

Table A6 Post-London Results, Priority Fee – Robustness (replace log(0) with 0)
(August 6, 2021 to December 31, 2022)

	log median priority fee, Wei			
	London		London + Merge	
	(1)	(2)	(3)	(4)
	levels	first differences	levels	first differences
block utilization	0.093***	-0.002	0.066***	-0.032***
	[0.003]	[0.004]	[0.003]	[0.003]
log block gas limit	-20.84***	-33.93***	-22.85***	-34.72***
	[1.031]	[1.758]	[1.026]	[1.795]
block contract call share	2.121***	2.604***	1.839***	2.309***
	[0.038]	[0.044]	[0.033]	[0.040]
log ETH price in USD	-0.282***	n.a.	-0.216***	n.a.
	[0.006]	–	[0.005]	–
block legacy share	2.632***	3.868***	2.429***	3.469***
	[0.012]	[0.024]	[0.011]	[0.024]
log block base fee	0.177***	1.495***	0.159***	1.244***
	[0.003]	[0.025]	[0.002]	[0.023]
adj. R-squared	0.170	0.165	0.165	0.145
sample size	2,503,394	2,503,393	3,267,661	3,267,660

P-values: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In columns (1) and (3) the dependent variable 'log median priority fee' is the natural logarithm of the median priority fee in Wei (10^{-18} ETH) among all transactions within a block (imputed for legacy transactions as the difference between the gas price and the block base fee); 'block utilization' is the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log block gas limit' is the natural logarithm of the block gas limit; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions in that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD; 'block legacy share' is the number of legacy transactions in a block divided by the total number of transactions in that block. 'log block base fee' is the natural logarithm of the base fee on a given block. In columns (2) and (4), the dependent and independent variables are the first differences of the respective block-level variables in columns (1) and (3). To control for the algorithmic increase in the block creation rate in the Merge period, in columns (3) and (4) we include a dummy variable which equals 0 before Sep. 15, 2022 and 1 for Sep. 15 onward. Columns (1) and (2) use data from Aug. 6, 2021 to Sep. 14, 2022 while columns (3) and (4) use data from Aug. 6, 2021 to Dec. 31, 2022. To account for the observations with zero median priority fee (9,289 observations or 0.3% of the sample) we replace the natural logarithm of zero with zero.

Table A7 Pre-London Results, Gas Price – Robustness (5th percentile gas price)
(January 1, 2021 to August 4, 2021)

	log pct5 gas price, Wei	
	(1)	(2)
	levels	first differences
block utilization	5.612***	4.525***
	[0.171]	[0.185]
log block gas limit	-11.77***	-34.32***
	[0.067]	[3.861]
block contract call share	3.051***	3.742***
	[0.038]	[0.042]
log ETH price in USD	1.665***	n.a.
	[0.020]	–
adj. R-squared	0.139	0.059
sample size	1,377,279	1,377,278

P-values: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Notes: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In column (1) the dependent variable 'log pct5 gas price' is the natural logarithm of the 5th percentile gas price in Wei (10^{-18} ETH) of all transactions within a block; 'block utilization' is the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log block gas limit' is the natural logarithm of the block gas limit; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions in that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD. In column (2) the dependent and independent variables are the first differences of the respective block-level variables from column (1). To account for the observations with zero 5th percentile gas price (7,958 observations or 0.6% of the sample) we add 1 Wei to all block 5th percentile gas prices before taking logs.

Table A8 Post-London Results, Gas Price – Robustness (5th percentile gas price)
(August 6, 2021 – December 31, 2022)

	log pct5 gas price, Wei			
	London		London + Merge	
	(1)	(2)	(3)	(4)
	levels	first differences	levels	first differences
L.block utilization	0.093***	0.118***	0.098***	0.119***
	[0.001]	[0.00009]	[0.001]	[0.0001]
log block gas limit	-8.639***	-4.720***	-7.639***	-4.734***
	[1.003]	[0.065]	[1.024]	[0.065]
block contract call share	0.185***	-0.009***	0.158***	-0.013***
	[0.009]	[0.0004]	[0.009]	[0.0004]
log ETH price in USD	1.454***	n.a.	1.420***	n.a.
	[0.007]	–	[0.007]	–
block legacy share	-0.068***	0.004***	-0.050***	0.001***
	[0.008]	[0.0004]	[0.008]	[0.0004]
adj. R-squared	0.414	0.361	0.556	0.360
sample size	2,503,393	2,503,392	3,267,660	3,267,659

P-values: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Note: OLS regressions with block-level data including a constant. Newey-West (HAC) standard errors robust to heteroskedasticity and serial correlation are reported in the square brackets. In columns (1) and (3) the dependent variable 'log pct5 gas price' is the natural logarithm of the 5th percentile gas price in Wei (10^{-18} ETH) of all transactions within a block; 'L.block utilization' is the first lag of block utilization, defined as the ratio of the total gas used by all transactions in a block to the gas limit of that block; 'log block gas limit' is the natural logarithm of the block gas limit; 'block contract call share' is the number of contract call transactions in a block divided by the total number of transactions within that block; 'log ETH price in USD' is the natural logarithm of the hourly Ether (ETH) price in USD; 'block legacy share' is the number of transactions in a block using the pre-London legacy (bid gas price) format divided by the total number of transactions in that block. In columns (2) and (4), the dependent and independent variables are the first differences of the respective block-level variables from columns (1) and (3). To control for the algorithmic increase in the block creation rate in the Merge period, in columns (3) and (4) we also include a dummy variable which equals 0 before Sep. 15, 2022 and 1 from Sep. 15 onwards. Columns (1) and (2) use data from the period Aug. 6, 2021 to Sep. 14, 2022; columns (3) and (4) use data from the period Aug. 6, 2021 to Dec. 31, 2022.

Figure A1: ETH Price in USD over the sample period

Appendix B

B1. Base fee determination

As of the Aug. 5, 2021 London upgrade, the Ethereum gas price (transaction fee per unit of complexity) consists of a “base fee” and a “priority fee”. The base fee of block t is algorithmically determined by the utilization rate and the base fee in the previous block $t-1$,

$$BaseFee_t = BaseFee_{t-1} \times (1 + \frac{1}{8} \times \frac{GasUsed_{t-1} - GasTarget}{GasTarget}) \tag{3}$$

where t is the block number and *GasTarget* is set equal to 15 million (15M) gas which is half of the post-London block *gas limit*. By construction, the base fee can increase or decrease by a maximum of 12.5% from one block to the next, with the magnitude of the increase or decrease determined by the magnitude of the deviation between observed block utilization and the 50% average utilization target.

B2. Post-London gas price structure

After the London upgrade the Ethereum gas price, g_t equals the sum of an algorithmically determined base fee, b_t and a user-submitted optional priority fee, $\pi_t \geq 0$,

$$g_t = b_t + \pi_t \tag{4}$$

Using equation (3) we have:

$$\begin{aligned}
\log(b_t) &= \log(b_{t-1}) + \log(1 + \frac{1}{8} \times \frac{GasUsed_{t-1} - GasTarget}{GasTarget}) \\
&= \log(b_{t-1}) + \log(1 + \frac{1}{4} \times \frac{GasUsed_{t-1} - 15M}{30M}) + \\
&= \log(b_{t-1}) + \log(\frac{7}{8} + \frac{1}{4} \times u_{t-1})
\end{aligned}$$

That is, the base fee of block t (which constitutes the major part of the gas price) is a deterministic function of the previous block' utilization rate, u_{t-1} and its base fee, b_{t-1}. Going back to the gas price, using (4), we obtain

$$g_t = f(u_{t-1}, b_{t-1}) + \pi_t \tag{5}$$

showing the dependence of the gas price on the preceding block's (lagged) utilization rate.