

Price Elasticity of Gas Demand on L1 and L2: Evidence from Ethereum and Arbitrum

Pranay Anchuri   

Offchain, Princeton, United States

Akaki Mamageishvili   

Offchain, Zurich, Switzerland

Abstract

We estimate the causal price elasticity of gas demand on Ethereum mainnet (L1) and Arbitrum One (L2), a quantity necessary for calibrating fee mechanism simulations, evaluating resource pricing reforms, and explaining observed usage patterns. A two-way fixed effects panel regression instrumented by each wallet’s own lagged base fee removes the congestion-driven endogeneity that causes naive regressions to substantially underestimate demand sensitivity. On Ethereum mainnet (full year 2025), the pooled IV elasticity is -0.006^{***} , near-inelastic: a 10% fee increase reduces total gas demand by approximately 0.06%. On Arbitrum One (October 2025–April 2026), the pooled IV elasticity is -0.036^{**} . Both chains are inelastic in the aggregate, with L2 measurably more responsive than L1. A per-resource decomposition of L2 demand reveals elasticities ranging from modestly elastic computation (-0.027^*) to -0.27^{***} for refunds, with storage growth (-0.15^{***}) and calldata (-0.06^*) in between. Behavioral clustering identifies always-on protocol wallets as near-inelastic and high-volume operators as substantially more responsive, with cluster-level elasticities up to roughly $6\times$ the pooled estimate. These results establish an empirical foundation for downstream simulations and for evaluating fee mechanism designs.

2012 ACM Subject Classification Applied computing → Economics

Keywords and phrases gas pricing, demand elasticity, instrumental variables, panel data, fixed effects, Arbitrum, Ethereum, transaction fee mechanism

Digital Object Identifier 10.4230/LIPIcs.AFT.2026.35

Related Version *Full Version:* <https://arxiv.org/pdf/2606.13555>

1 Introduction

On Ethereum, the base fee adjusts in each block in response to block utilization, targeting 50% block capacity on average. Despite the mechanism’s theoretical properties being well-studied [26, 25, 7, 18], the empirical demand response to fee changes remains poorly quantified. Knowing how much users reduce their onchain resource consumption when fees rise is necessary for multiple reasons. First, it allows calibrating simulations of the fee adjustment rule under different parameters. Second, evaluating potential throughput gains from multi-resource fee mechanisms requires demand elasticity estimates for each resource. Third, it helps explain current usage patterns under different fee regimes. Fourth, it helps find the effects of repricing certain opcodes¹.

Estimating demand elasticity naively is complicated by omitted variable bias. Higher network demand drives the base fee up (through the EIP-1559) and gas usage up (through more users transacting), producing a spurious positive correlation between fees and gas even

¹ The Ethereum ecosystem is considering increasing the gas target/limit while also increasing the gas cost of storage creation; Ethereum improvement proposal numbered 8037 (<https://eips.ethereum.org/EIPS/eip-8037>) is specifically proposed for this. Monad (also an EVM chain) increased the gas cost of storage creation at initialization time relative to Ethereum’s default: <https://category-labs.github.io/category-research/monad-initial-spec-proposal.pdf>



© Pranay Anchuri and Akaki Mamageishvili;
licensed under Creative Commons License CC-BY 4.0


8th Conference on Advances in Financial Technologies (AFT 2026).

Editors: Aggelos Kiayias and Maria Kyrpoulou; Article No. 35; pp. 35:1–35:23

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Raw pooled OLS of log gas usage on log base fee, Arbitrum L2 panel ($N = 16,939,781$ wallet-hour observations). The positive slope ($\hat{\beta} = +0.094$) contradicts the expected negative demand response, illustrating the upward bias induced by congestion-driven endogeneity. The two dense vertical bands at $\ln p \approx 16.1$ and ≈ 16.8 correspond to the Arbitrum base fee floor levels of 0.01 gwei and 0.02 gwei respectively (the floor was raised from 0.01 to 0.02 gwei mid-sample), where the base fee is anchored when demand falls below the gas target.

when isolated to wallet level. Figure 1 illustrates this directly: a pooled OLS regression of log gas usage on log base fee across 16,939,781 Arbitrum wallet-hour observations yields $\hat{\beta} = +0.094$, a positive coefficient that contradicts the expected negative demand response. This is in contrast to a typical demand elasticity estimation scenario where the price change is mostly exogenous, e.g., due to a supply drop. In the context of Ethereum and its rollups, the target supply is fixed and the price (gas fees) follows the demand exactly, as described above. Two-way fixed effects (FE: wallet and time) absorb level differences but still leave within-cell congestion variation uncontrolled. We address this residual endogeneity by using each wallet’s own lagged base fee as an instrument: a fee level that is predetermined relative to current demand decisions and strongly correlated with the current fee through the autocorrelation built into the EIP-1559 fee adjustment rule.

Gas is priced as a single unit on both chains, but the underlying operations vary substantially in their resource footprint: a storage-creating transaction consumes a different mix of compute, I/O, and state resources than a pure computation. Evaluating proposals for multi-dimensional resource pricing requires knowing which dimensions actually respond to fee signals. The same decomposition informs the demand impact of opcode repricing, where increasing the cost of one resource category may or may not shift usage depending on the price sensitivity of that dimension.

The pooled elasticity is an average across all wallet types on the chain: oracle updaters that submit at fixed intervals, DeFi operators that optimize gas costs, and MEV searchers that time activity to congestion. These populations differ fundamentally in their relationship to fee changes, and a single estimate characterizes none of them.

Contributions

1. **Causal elasticity estimates for L1 and L2.** Using each wallet's own lagged base fee as an instrument, we obtain Instrumental Variables (IV) estimates of gas demand elasticity for Ethereum mainnet (-0.006) and Arbitrum One (-0.036).
2. **Per-resource decomposition for L2.** We estimate separate elasticities for seven Arbitrum resource dimensions, revealing that state-creating and state-clearing operations are far more price-sensitive than computation or reads, which can be justified by their time-insensitivity.
3. **Wallet-type heterogeneity.** K-means clustering on behavioral features identifies distinct user types with qualitatively different demand responses, from near-inelastic always-on protocol wallets to substantially elastic high-volume operators.

The pooled elasticity estimates confirm near-inelastic demand on both chains: a 10% fee increase reduces total gas demand by approximately 0.06% on Ethereum and 0.36% on Arbitrum. However, per-resource decomposition reveals that this aggregate inelasticity masks substantial resource-level heterogeneity: refunds (-0.27) and storage growth (-0.15) are the most price-sensitive dimensions, while computation (-0.027) is near-inelastic. Wallet clustering further identifies congestion-timing wallets for which the lagged-fee instrument is invalid, and restricts causal interpretation to the remaining segments.

The remainder of the paper is organized as follows. Section 2 reviews the EIP-1559 mechanism and L1/L2 fee regime differences. Section 3 describes the Ethereum and Arbitrum datasets. Section 4 develops the econometric framework. Section 5 presents the wallet clustering procedure. Sections 6 and 7 report the main elasticity results for L1 and L2 aggregate demand; Section 8 provides the per-resource decomposition. Section 9 discusses applications of elasticity analysis. Section 10 discusses implications and limitations.

2 Background

2.1 The EIP-1559 Fee Mechanism

Permissionless blockchains such as Ethereum need to control usage of their resources, to avoid long-lasting denial of service attacks and keep low hardware or bandwidth requirements for participating nodes in the network. For this reason, the Ethereum chain introduced Ethereum improvement proposal (EIP), numbered 1559, which adjusts the fee for resource consumption. There are different types of resources transactions consume, which are converted into gas units. Ethereum aims to operate on a target gas consumption. To achieve this goal, a simple rule is implemented: when the consumption of gas in the last block is larger than the target, the gas fee (in ETH) goes up; when the consumption of gas in the last block is lower than the target, the gas fee goes down; only when the gas usage is exactly equal to the target the gas fee stays the same. Gas fee changes should cause a short-term demand shift for the resources. Block sizes are usually upper bounded in terms of gas, commonly referred to as a block limit. The limit is set to two times the target. Recently, the target has moved up by a factor of 2 on Ethereum, from 15M to 30M, in an effort to scale the chain, with further increases expected.

A key design goal of EIP-1559 is fee predictability: the base fee is algorithmically calculated from the recent block utilization. This allows wallets to anticipate costs and plan resource consumption accordingly. This predictability both induces per-transaction optimization and deters participation during high-fee regimes. These are two distinct behavioral responses, and the former is the primary focus of this work.

2.2 L1 vs. L2 Fee Regimes

Major rollups adopt EIP-1559 or a variant of it to price their native resources and manage congestion. L2 fees are typically lower than L1 fees for two reasons: first, EIP-4844 expanded blob throughput so that L1 data availability costs for rollups are now minimal; second, lower hardware and bandwidth requirements reduce the marginal cost of an L2 block relative to an L1 block. Recent gas target increases on L1 have compressed the gap further, making L1 fees relatively cheap as well. The total fee incurred by an L2 end user consists of a native L2 congestion fee and an L1 data availability component; our analysis focuses on the native L2 base fee.

Recently, Arbitrum updated its pricing regime, introducing a gas ladder². This pricing mechanism allows temporary outbursts of gas usage while keeping long-term targets fixed.

An interesting difference between these two types of chains is that Ethereum does not have a lower bound on gas price; it can go as low as 1 wei (10^{-18} ETH), while Arbitrum does have it. Recently, with the pricing mechanism update, Arbitrum increased it from 0.01 gwei to 0.02 gwei. A potential explanation for this is to deter spam transactions, the activity which is delegated to block builders on L1. Other major rollups enforce base fee floors that are substantially lower: OP Mainnet sets its minimum base fee to 0 wei (no effective floor), Base Mainnet to 0.005 gwei (5,000,000 wei), and Scroll to approximately 0.00012 gwei (set via configuration rather than the protocol). Arbitrum's 0.02 gwei floor is the highest among these rollups, roughly $4\times$ Base's and orders of magnitude above the others.

The Arbitrum Nitro node is instrumented to track gas consumption across multiple resource dimensions (computation, storage reads and writes, storage growth, history growth, calldata, and refunds), attributing each opcode's cost to the corresponding category during execution. While the current Arbitrum fee mechanism charges a single gas price, the per-resource data provide a richer view of which operations are sensitive to fee changes, and form the basis for the per-resource demand decomposition in Section 8.

Note that when the base fee is at the floor for extended periods, fee variation is compressed and the identifying variation exploited by the instrument becomes sparse. This constraint is specific to L2 estimation and does not arise on L1, where fee variation is continuous and unbounded from below.

3 Data

3.1 Ethereum L1

The Ethereum dataset comprises all transactions executed on the Ethereum mainnet during 2025. Wallets are included in the study if they submitted at least 500 transactions during the year, ensuring sufficient within-wallet variation for the fixed effects estimation. After filtering on this threshold, the panel data contains 44,457 wallets observed over 8,751 hourly periods, yielding a total of 35,321,371 wallet-hour observations.

For each wallet-hour, the dependent variable is the wallet's total gas used, the independent variable is the wallet's own average base fee in that hour in gwei, and the control is the number of transactions submitted.

Both L1 and L2 panels are at hourly granularity. On Ethereum, we use daily time fixed effects rather than hourly: the Ethereum base fee adjusts each block (~ 12 seconds), but the dominant fee variation is between days rather than within a day. Daily fixed effects efficiently absorb aggregate day-level demand shocks; the within-day, within-wallet hourly fee variation identifies β .

² Arbitrum gas-ladder AIP, Arbitrum Forum, 2025.

■ **Table 1** L2 per-resource dimension coverage. Coverage is the fraction of total wallet-hours (16,939,781) where the resource consumption is non-zero.

Dimension	Description	Wallets	Wallet-hour obs.	Coverage (%)
computation	In-memory EVM execution	23,119	16,939,781	100.0
historyGrowth	Logs and events (append-only)	20,052	13,702,494	80.9
storageAccessRead	Reading existing state	22,652	16,126,735	95.2
storageAccessWrite	Updating existing state	20,883	13,475,351	79.6
storageGrowth	Increasing persistent state size	18,990	8,767,695	51.8
l2Calldata	Transaction calldata size on the child chain	20,799	15,782,132	93.2
refund	Gas refunded from storage-clearing operations	17,770	8,103,016	47.8

3.2 Arbitrum L2

The Arbitrum L2 dataset covers all Arbitrum One transactions from October 2025 through April 2026 (seven months). We use the same minimum threshold of 500 transactions to select the wallets. The resulting panel contains 23,119 wallets observed over 5,087 hourly periods, yielding a total of 16,939,781 wallet-hour observations.

In addition to aggregate gas usage, the Arbitrum dataset records per-transaction gas consumption across seven resource dimensions: *computation*, *historyGrowth*, *storageAccessRead*, *storageAccessWrite*, *storageGrowth*, *l2Calldata*, and *refund*. Table 1 describes each dimension. The dataset was constructed by instrumenting the Nitro execution engine to attribute each opcode’s gas cost to one of these resource categories during transaction execution; the same opcode may map to different categories depending on execution context (for example, an *SSTORE* counts toward *storageGrowth* on first write but *storageAccessWrite* on updates). Computation is non-zero in every wallet-hour by construction; refund appears in fewer than half, reflecting the discretionary nature of state-clearing operations.

3.3 Panel Data Construction

Panels are constructed by aggregating raw transaction data to the (wallet, hour) level for both chains. For each wallet-hour, all transactions submitted by that wallet in the hour are aggregated to obtain total gas used and transaction count, and the base fee is averaged across those transactions. Time fixed effects are daily for Ethereum (365 day dummies over the hourly panel) and hourly for Arbitrum. We exclude system and precompile addresses from both panels.

The instrument, the wallet’s own lagged base fee, is computed by retrieving the most recent period during which the wallet was active and the average base fee during that period.

4 Econometric Framework

In this section, we develop the estimation strategy for the price elasticity of gas demand. We begin by showing why a naive regression of gas usage on the base fee produces a biased estimate, formalize the bias using the omitted-variable formula, and then build the solution in two steps.

- Two-way fixed effects to absorb persistent wallet heterogeneity and time-dependent shocks.
- Instrumental variables to remove the residual endogeneity that fixed effects cannot eliminate.

4.1 The Identification Problem

The parameter of primary interest in this work is β , the price elasticity of gas demand: the percentage change in a wallet's gas consumption following a one-percent increase in the base fee. To isolate the source of bias, consider first a model with a single regressor,

$$\log g_i = \beta \log p_i + u_i, \quad (1)$$

where g_i is gas used, p_i is the base fee, and $u_i \sim N(0, \sigma^2)$. We use a log-log specification throughout.

The ordinary least squares (OLS) estimator satisfies

$$\hat{\beta}^{\text{OLS}} = \frac{\widehat{\text{Cov}}(\log p, \log g)}{\widehat{\text{Var}}(\log p)}, \quad (2)$$

and converges in probability to

$$\text{plim } \hat{\beta}^{\text{OLS}} = \beta + \frac{\text{Cov}(\log p, u)}{\text{Var}(\log p)}. \quad (3)$$

OLS is consistent if and only if $\text{Cov}(\log p, u) = 0$.

In our setting, this condition fails because of (latent) aggregate demand pressure, or congestion. Denote by D_t the aggregate level of demand in the time period t . Under the EIP-1559 adjustment rule, the base fee for each block is adjusted by a factor that is proportional to the excess utilization of the preceding block. Therefore, congestion raises the fee: $\text{Cov}(\log p, D_t) > 0$. At the same time, a high-congestion period reflects users submitting more transactions, which raises the usage: $\text{Cov}(\log g, D_t) > 0$. The term D_t is part of the error u in Equation (1) and these two conditions combine to give $\text{Cov}(\log p, u) > 0$. Therefore, the bias term in Equation (3) is positive. OLS overestimates the true elasticity β , biasing it toward zero or even reversing its sign. This is a structural bias and cannot be eliminated by increasing the sample size.

4.2 Two-Way Fixed Effects Model

The congestion described in Section 4.1 operates on two easily separable levels. First, wallets differ persistently: a high-frequency trading bot and a casual user differ in their gas consumption across all time periods for reasons unrelated to the base fee. Second, the network-level congestion shocks are common across all wallets within a given period. For example, a fee spike during an NFT drop affects all wallets simultaneously. Wallet fixed effects absorb the first source and the time fixed effects absorb the second.

Demand equation

The demand equation for wallet i in time period t is

$$\log g_{it} = \alpha_i + \gamma_t + \beta \log p_{it} + \phi \log n_{it} + u_{it}, \quad (4)$$

where p_{it} is wallet i 's average base fee in period t , n_{it} is the number of transactions submitted (a control variable), and u_{it} is the residual. The term α_i absorbs all the time-invariant heterogeneity between wallets. These include bot versus human, typical gas usage per transaction, etc. The term γ_t absorbs all shocks at the period-level experienced by all wallets. These include network congestion level, time of day, day of week, and macro demand events. We also include the transaction count n_{it} because a wallet with more transactions mechanically consumes more gas. Controlling for n_{it} isolates the response in usage to the fee.

The control for n_{it} reflects a deliberate choice about the margin being estimated. The total gas g_{it} combines two distinct decisions made by the wallets: how many transactions to submit in a period t and how much gas to spend in each transaction. The extensive margin, transaction count, captures the participation decisions, whereas the intensive margin, gas used in a transaction, captures things such as reducing gas usage or substituting toward cheaper resources. By controlling for n_{it} , we partial out the extensive margin and identify β as the intensive margin elasticity. The latter is directly relevant in evaluating the efficiency of different fee mechanisms, since fee mechanisms should not merely deter activity but enable wallets to optimize their resource usage. Because n_{it} is fee-responsive, conditioning on it isolates the per-transaction (intensive-margin) response rather than the total effect, which is the margin most relevant to mechanism efficiency.³ We condition on n_{it} rather than instrument for it because our estimand is the intensive-margin elasticity: instrumenting for n_{it} would recover the total elasticity, a different quantity that conflates fee-driven participation changes with per-transaction resource adjustments.

After absorbing the terms α_i and γ_t , the elasticity parameter β is identified from within-wallet and within-period variation. More specifically, β measures how a wallet i 's gas usage changes when the fee deviates from both the wallet's own time-average and the period's cross-wallet average.

Two-way demeaning

Estimating β from Equation (4) requires either one-hot encoding the fixed effects or demeaning the outcome and the regressors. Let $\bar{y}_{i\cdot}$ denote the wallet time-mean, $\bar{y}_{\cdot t}$ the period cross-wallet mean, and $\bar{y}_{\cdot\cdot}$ the overall mean. The demeaned outcome is

$$\tilde{y}_{it} = y_{it} - \bar{y}_{i\cdot} - \bar{y}_{\cdot t} + \bar{y}_{\cdot\cdot}, \quad (5)$$


and the same transformation is applied to the regressors $\log p_{it}$ and $\log n_{it}$.

By the Frisch–Waugh–Lovell theorem [14, 20], OLS on the demeaned model is identical to including dummy variables for α_i and γ_t . We demean all variables by subtracting wallet and period means in alternation until convergence. This process terminates in fewer than 30 iterations in all panel datasets used in this study; the post-estimation fixed-effects recovery described in the full version of the paper applies the same algorithm to the full residuals and converges within 500 iterations.

Time granularity

Both analyses use hourly panels; the granularity of γ_t differs. On Ethereum mainnet, the base fee adjusts by at most $\pm 12.5\%$ per block (approximately 12 seconds), but the dominant variation is between days rather than within a day. We therefore use daily time fixed effects: γ_t absorbs day-level aggregate demand shocks, and β is identified from within-day, within-wallet hourly fee deviations. On Arbitrum, hourly fixed effects are used because the L2 base fee adjusts on an hourly schedule.

³ Omitting $\log n_{it}$ from both stages reverses the sign of the estimate (+0.052 on Ethereum, +1.52 on Arbitrum): high-fee periods attract wallets that submit more and heavier transactions, so total gas co-moves positively with fees when transaction count is uncontrolled. Conditioning on n_{it} isolates the per-transaction response.



■ **Figure 2** Causal structure of the gas demand model. Congestion u_{it} drives both the base fee p_{it} (base fee update rule) and gas usage g_{it} (through aggregate demand), inducing the endogeneity that biases OLS. The instrument z_{it} (the wallet’s lagged fee) satisfies relevance (z_{it} predicts p_{it}) and the exclusion restriction (z_{it} has no direct path to g_{it} , marked $\times$). Instrumental variables estimation uses only the variation in p_{it} attributable to z_{it} , which is uncorrelated with u_{it} .

Residual endogeneity

The above two-way demeaning substantially attenuates the bias. The FE-OLS estimate of β is -0.007 on Ethereum and -0.016 on Arbitrum, both negative and consistent with downward-sloping demand. Within-period demand shocks remain in $\tilde{u}_{it}$ even after demeaning. For example, a wallet that transacts more in period t because of within-period congestion that also raised the fee in that period, contributes a positive correlation between $\tilde{p}_{it}$ and $\tilde{u}_{it}$. The two-way FE removes between-period congestion but not within-period endogeneity. This motivates the instrumental variables approach below.

4.3 Instrumental Variables

Figure 2 illustrates the causal structure of the model. Congestion u_{it} drives both the base fee p_{it} and gas usage g_{it} through separate channels, creating the endogeneity that results in an inconsistent estimate of the elasticity β . An instrumental variable z_{it} that shifts the fee but has no direct effect on gas demand breaks this linkage.

For each wallet i in period t , the instrument is the wallet’s own average base fee in its immediately preceding active period:

$$z_{it} = \log p_{i,t-1}, \quad (6)$$

where $t - 1$ denotes the most recent period before t in which wallet i submitted at least one transaction.

For z_{it} to be a valid instrument, the properties of *Relevance* and *Exclusion* must be established.

Relevance

The instrument must be correlated with the current fee: $\text{Cov}(z_{it}, p_{it}) \neq 0$. The base fee exhibits strong autocorrelation because EIP-1559 adjusts it continuously in response to prior block utilization, so the fee in one period is a strong predictor of the fee in the next. The first-stage partial F -statistic confirms this correlation empirically. Results reported in Sections 6 and 7 far exceed the conventional threshold of 10 for reliable inference.

■ **Table 2** Fraction of active hours in above-median-fee periods, by cluster (L2).

Cluster	Label	n wallets	Median	P25	P75
C0	Hi-vol burst	2,608	0.367	0.268	0.474
C1	Long-running	3,279	0.481	0.328	0.596
C2	Low-intensity	5,818	0.270	0.239	0.323
C3	Fee-avoiding	4,253	0.270	0.224	0.329
C4	Variable gas	697	0.302	0.248	0.363
C5	Cong-active [†]	6,466	0.754	0.400	0.879

[†]Excluded from IV estimation (Section 4.3).

Exclusion Restriction

The validity of the instrument also relies on the assumption that it is uncorrelated with the error u_{it} . The lagged base fee is *predetermined*, as it is a function of gas utilization in periods preceding t . Therefore, it is established before a wallet makes a marginal transaction decision at time t . Moreover, the two-way FE controls for both the time-invariant wallet characteristics and time-specific demand shocks. Under these conditions, the lagged fee affects the current usage only through its influence on the current base fee.

However, the exclusion restriction property is violated by wallets that strategically time their activity with congestion cycles. If a wallet's activity is sparse then the lagged fee z_{it} represents the peak of a prior congestion. Now, if the congestion is serially correlated then the instrument z_{it} becomes an informative signal of the current error term u_{it} . This leads to a correlation between the instrument and the latent demand shocks.

To mitigate the bias introduced by this channel, we identify such congestion sensitive wallets based on the proportion of wallet's activity during periods where the base fee exceeds the median hourly base fee. Wallets with an anomalously high fraction are the ones that are disproportionately active during spikes. For these wallets, the lagged fee is invalid as an instrument. We identify this as cluster C5 (congestion-active, Section 5): the median C5 wallet conducts 75.4% of its active hours in above-median-fee periods (IQR: 40.0–87.9%), compared with 27.0–48.1% for the five retained clusters (Table 2). We therefore restrict causal interpretation to the five clusters whose timing patterns are consistent with the exclusion restriction.

4.4 Calculating β using Two-Stage Least Squares (2SLS)

Building on the discussion in this section, we now describe how we estimate the elasticity β . Prior to estimation, all variables are two-way demeaned as described in Section 4.2.

In the first stage, the demeaned fee is regressed on the demeaned instrument and the demeaned transaction count (control):

$$\tilde{p}_{it} = \pi \tilde{p}_{i,t-1} + \delta \tilde{n}_{it} + v_{it}. \quad (7)$$

Here $\tilde{p}_{i,t-1}$ is shorthand for the two-way demeaned instrument $\tilde{z}_{it}$, where $z_{it} = \log p_{i,t-1}$. The fitted values $\hat{p}_{it}$ capture only the variation in the fee that is driven by the instrument.

In the second stage, the fee is replaced with the fitted values from the first stage.

$$\tilde{g}_{it} = \beta \hat{p}_{it} + \phi \tilde{n}_{it} + \varepsilon_{it}. \quad (8)$$

35:10 Price Elasticity of Gas Demand on L1 and L2

With a single endogenous regressor (base fee) and a single instrument, the estimate from 2SLS reduces to

$$\hat{\beta}^{IV} = \frac{\widehat{\text{Cov}}(z, \tilde{g})}{\widehat{\text{Cov}}(z, \tilde{p})}. \quad (9)$$

Interpretation

The 2SLS estimator in Equation (9) identifies the local average treatment effect (LATE). It is the average demand response for wallets whose fee in period t is affected by the lagged fee. This is the subpopulation of wallets that respond to prior price levels when deciding the current gas usage. Wallets that execute on a fixed schedule regardless of the base fee are not shifted by the instrument and consequently have negligible contribution to the estimate.

Preview of main results

Applying this framework to the full panels, the pooled IV elasticity is -0.006^{***} ($\text{SE} = 0.0005$) on Ethereum and -0.036^{**} ($\text{SE} = 0.014$) on Arbitrum. Both estimates are near-inelastic and are detailed in Sections 6 and 7. Section 5 and Section 8 decompose these pooled estimates by behavioral cluster and resource dimension, respectively.

5 Wallet Clustering

The IV β estimate from Section 4 is a precisely identified quantity, but it is a weighted average of demand responses across all wallet types in the data. In general-purpose blockchain networks like Ethereum and Arbitrum, wallets span a wide range of behavioral regimes: oracle updaters, arbitrageurs, DeFi operators, and retail users. These wallets differ not only in their average gas consumption but also in how they respond to fee changes. An oracle contract that must submit a transaction at regular intervals to maintain the on-chain state faces no meaningful price sensitivity. At the same time, a high-volume arbitrageur may substitute toward lower fee periods. Pooling all wallet types into a single regression Equation (4) produces an average elasticity estimate that does not characterize any particular wallet class.

Clustering the wallets into behaviorally homogeneous groups prior to IV estimation recovers within-cluster elasticities that are interpretable as the demand response of a coherent user type. It also serves a diagnostic role in our analysis: wallets that activate disproportionately during congestion spikes are localized in a single cluster identifiable by their high-fee fraction, as discussed in Section 4.3.

5.1 Clustering Procedure

Wallets are partitioned using k -means clustering based on a set of behavioral features computed from each wallet's transaction history over the entire sample period. Before clustering, we also standardize the features to zero mean and unit variance. We select the number of clusters k with the elbow method: the within-cluster sum of squared distances is computed for $k = 2, \dots, 12$, and k is chosen at the point of diminishing returns.

Elbow plots for all three analyses are shown in the full version of the paper. All three exhibit a clear kink at $k = 6$.

■ **Table 3** Clustering features. The top five are used in all analyses; the bottom four are added for the L2 per-resource clustering only. All features are standardised to zero mean and unit variance before clustering.

Feature	What it captures
<i>Behavioural features (all analyses)</i>	
<code>log(tx_frequency)</code>	Transaction rate per active hour; separates high-throughput operators from occasional users.
<code>log(gas/tx)</code>	Mean gas per transaction; separates simple transfers from complex contract interactions.
<code>log(active_hours)</code>	Total hours with at least one transaction; separates persistent from episodic wallets.
<code>gas_cov</code>	Coefficient of variation (σ/μ) of gas per transaction; low values indicate a fixed operation type, high values indicate volatile mixed activity.
<code>high_fee_fraction</code>	Fraction of active hours above the median hourly fee; identifies wallets that transact preferentially during congestion. Also used as the exclusion restriction diagnostic (Section 4).
<i>Resource composition features (L2 per-resource only)</i>	
<code>frac_computation</code>	Share of gas from EVM opcode execution.
<code>frac_storageAccessRead</code>	Share from reading existing storage slots.
<code>frac_l2Calldata</code>	Share from L2 calldata.
<code>frac_storageGrowth</code>	Share from creating new storage slots.

5.2 Features

Table 3 lists all features used for clustering. The five behavioral features are used in both the total-gas and per-resource analyses. For the resource level analysis, we add four resource composition features. Each of these features measures the fraction of a wallet’s total gas attributable to a given resource dimension averaged over the entire time period.

The remaining dimensions (`storageAccessWrite`, `historyGrowth`, `refund`) are excluded from the clustering features as their variation is already captured by the included features.

5.3 Cluster Profiles

5.3.1 L1 clusters


Table 4 reports the six L1 clusters, ordered by size. Figure 3 shows a t-SNE projection of the five-dimensional feature space, confirming that the resulting clusters occupy geometrically distinct regions.

The clusters reflect qualitatively different wallet types. C3 is the most persistently active, with a median of 3309 active hours (38% of all hours in the year). C1 has the highest frequency (25 transactions per active hour) but operates in short bursts (53 median active hours) and generally avoids high-fee periods ($\text{hff} = 0.34$). These wallets are likely to be opportunistic and sensitive to the base fee. C0 is highly active during high-fee periods suggesting price insensitivity or time-sensitive applications. C4 is the smallest cluster but exhibits the highest gas coefficient of variation, indicating wallets that alternate between simple and complex transactions. C5 is similar to C1 with respect to high-fee fraction but uses $3.5\times$ more gas per transaction.

35:12 Price Elasticity of Gas Demand on L1 and L2

■ **Table 4** L1 (Ethereum) wallet clusters ($k = 6$, full year 2025). Medians reported for each feature. tx_freq: transactions per active hour; gas/tx: average gas per transaction (thousands); act. hrs: active hours in sample; gas CoV: coefficient of variation of gas per transaction; hff: high-fee fraction.

Cl.	Wallets (%)	tx_freq	gas/tx (K)	act. hrs	gas CoV	hff
C0	15,135 (34.0%)	2.52	130	330	0.92	0.69
C1	3,609 (8.1%)	25.1	50	53	0.96	0.34
C2	11,625 (26.1%)	1.54	35	543	0.68	0.50
C3	4,123 (9.3%)	2.93	114	3,309	0.82	0.50
C4	1,280 (2.9%)	3.04	106	368	3.33	0.53
C5	8,685 (19.5%)	2.86	171	303	1.01	0.34



■ **Figure 3** t-SNE projection of L1 wallet feature vectors ($k = 6$, five behavioral features). Each color represents one cluster. Clusters occupy well-separated regions, confirming meaningful behavioral separation.

5.3.2 L2 clusters (total gas)

Table 5 reports the six L2 clusters. Figure 4 shows the corresponding t-SNE projection.


C0 is the most active cluster by transaction rate (59 transactions per active hour) and with the highest mean hourly gas consumption (16.9M gas per hour) but with a very short median activity. C1 combines modest transaction activity with a very high number of active hours. C2 is the biggest cluster and has both low transaction frequency and gas usage. C3 has the lowest high-fee fraction, indicating wallets that are sensitive to the base fee. C4 has very high gas coefficient of variation. C5 is the cluster with highest high-fee fraction. The wallets transact predominantly in the most expensive gas fee regimes. As discussed in Section 4.3, the exclusion property is implausible for this cluster.

L2 clusters (per-resource)

The per-resource analysis uses nine features (five behavioral plus four resource composition fractions). Six clusters emerge at the same $k = 6$ elbow. C1 has high transaction frequency, high-fee fraction, and transactions dominated by computation and storage reads. C2 is the

■ **Table 5** L2 (Arbitrum) total-gas wallet clusters ($k = 6$, October 2025–April 2026). Medians reported. gas/hr: mean gas used per active hour (millions). Other columns as in Table 4.

Cl.	Wallets (%)	tx_freq	gas/hr (M)	act. hrs	gas CoV	hff
C0	2,608 (11.3%)	58.74	16.87	24	1.02	0.53
C1	3,279 (14.2%)	8.43	2.14	2,858	1.25	0.51
C2	5,818 (25.2%)	2.31	0.17	532	0.97	0.54
C3	4,253 (18.4%)	5.24	1.11	166	0.90	0.34
C4	697 (3.0%)	3.42	0.51	898	4.42	0.48
C5	6,466 (28.0%)	16.81	0.90	117	1.48	0.90



■ **Figure 4** t-SNE projection of L2 total-gas wallet feature vectors ($k = 6$, five behavioral features). Clusters are well-separated. C5 (high-fee fraction 0.90, upper right) forms a tight distinct region.

high-volume cluster. C3 (35.0%) is the largest cluster, with general DeFi activity. C4 is computation-dominant. C5 wallets have significant storage-growth fractions. Full cluster profiles are reported in the full version of the paper.

Adversarial-activity validation for both L2 analyses is reported in the full version of the paper.

6 Results: Ethereum L1

6.1 First Stage

The instrument is strongly relevant on L1. The overall first-stage coefficient is $\hat{\pi} = 0.480$. Although the partial F -statistics are inflated by the large sample size, the first-stage coefficient confirms a genuinely strong first stage, reflecting the strong autocorrelation in the daily L1 base fee. The first-stage partial R^2 is 0.252, confirming that the lagged base fee explains a substantial share of current-fee variation even after partialling out the fixed effects.

■ **Table 6** Ethereum L1 IV elasticity estimates by cluster. Standard errors clustered at the wallet level. ** $p < 0.01$, *** $p < 0.001$; n.s. not significant at 5%.

Group	Wallets (%)	$\hat{\beta}_{OLS}$	$\hat{\beta}_{IV}$	SE	Partial F
Pooled	44,449 (100%)	-0.0074	-0.0059***	0.0005	11,873,625
C0 (hi-fee act.)	15,135 (34.0%)	-0.0110	-0.0236***	0.0015	616,714
C1 (bursty)	3,609 (8.1%)	+0.0011	+0.0071 n.s.	0.0051	111,703
C2 (low-intens.)	11,625 (26.1%)	-0.0015	-0.0016**	0.0006	1,147,259
C3 (always-on)	4,123 (9.3%)	-0.0022	-0.0010 n.s.	0.0007	15,170,218
C4 (variable)	1,280 (2.9%)	-0.0116	-0.0066 n.s.	0.0036	298,077
C5 (fee-avoid.)	8,685 (19.5%)	-0.0308	-0.0212***	0.0022	592,378

6.2 IV Estimates

Table 6 reports 2SLS estimates for the pooled sample and each of the six clusters in Section 5.3.1. The pooled elasticity is -0.0059^{***} ($SE = 0.0005$): a 10% increase in the base fee reduces total L1 gas demand by approximately 0.06%. L1 demand is near-inelastic. The aggregate near-inelasticity reflects the large share of low-intensity (C2, 26.1%) and always-on (C3, 9.3%) wallets whose consumption does not vary with the fee. Given the size of the panel, statistical significance is easily attained; the economically relevant takeaway is that aggregate L1 gas demand is, for practical purposes, insensitive to fee changes. The two most elastic clusters are C0 (34.0%), with $\hat{\beta} = -0.0236^{***}$ ($SE = 0.0015$), and C5 (19.5%), with $\hat{\beta} = -0.0212^{***}$ ($SE = 0.0022$). Wallets in C3 (9.3%) have the highest first-stage F ($\hat{\pi} = 0.689$, $F = 15,170,218$) yet the elasticity is near-zero and statistically insignificant (-0.0010 , $p = 0.17$). These are wallets that do not respond to fee signals.

On L1, the OLS and IV estimates are close in magnitude: daily fixed effects already absorb most of the between-day congestion that drives the bias, so instrumenting moves the pooled estimate only slightly and the sign of the adjustment varies across clusters. The correction is larger on L2 (Section 7), where hourly variation leaves more residual endogeneity.

The negative slope is visible after removing wallet and day fixed effects, confirming a downward-sloping demand relation in the partialled-out data. See full version of the paper.

7 Results: Arbitrum L2 Aggregate Demand

7.1 First Stage

The Arbitrum total-gas panel uses hourly time fixed effects. The pooled first-stage coefficient is $\hat{\pi} = 0.052$, significantly lower than on Ethereum. This reflects the frequency with which Arbitrum's base fee rests at its lower bound. Because the partial F -statistic scales with sample size, its large value reflects the size of the panel rather than instrument strength; the first-stage partial R^2 is only 0.007, confirming that the lagged fee explains a negligible share of current-fee variation on L2 after absorbing the fixed effects. The L2 estimates are therefore identified from limited fee movement.

The pooled partial F -statistic is 118,405, well above the commonly used threshold. Per-cluster F -statistics range from 1,361 (C0, high-volume but sparsely active per hour) to 113,875 (C1, long-running), confirming the relevance property of the instrument.

■ **Table 7** Arbitrum L2 aggregate IV elasticity estimates by cluster. Standard errors clustered at the wallet level. ** $p < 0.01$, *** $p < 0.001$; n.s. not significant at 5%; n.i. not identified. † excluded from causal interpretation.

Group	Wallets (%)	$\hat{\beta}_{OLS}$	$\hat{\beta}_{IV}$	SE	Partial F
Pooled	23,119 (100%)	-0.0161	-0.0359**	0.0141	118,405
C0 (hi-vol.)†	2,606 (11.3%)	-0.0440	-1.163 n.i.	0.1795	1,361
C1 (long-run.)	3,279 (14.2%)	-0.0285	-0.0782**	0.0268	113,875
C2 (low-intens.)	5,818 (25.2%)	-0.0131	-0.0057 n.s.	0.0374	4,662
C3 (fee-avoid.)	4,253 (18.4%)	-0.0311	+0.0792 n.s.	0.0574	2,709
C4 (variable)	697 (3.0%)	-0.0306	-0.1272 n.s.	0.1687	2,749
C5 (cong.-act.)†	6,466 (28.0%)	-0.0102	+0.0314 n.s.	0.0249	7,163

7.2 IV Estimates

Table 7 reports IV estimates for the Arbitrum aggregate demand model. The pooled elasticity is -0.0359^{**} ($SE = 0.0141$), six times larger in magnitude than the Ethereum estimate. A 10% fee increase on Arbitrum reduces total gas demand by an estimated 0.36%, less elastic in absolute terms but measurably more responsive than Ethereum.

The most credible cluster-level estimate is C1 (long-running high-gas wallets, 14.2%), with $\hat{\beta} = -0.0782^{**}$ ($SE = 0.0268$, $F = 113,875$). The median number of active hours is highest in this cluster, giving the instrument strong identifying power. The IV estimate for C2 (25.2%) is near-zero and statistically insignificant (-0.0057 , $p = 0.88$). Wallets in this cluster are occasional, low-frequency users that do not adjust gas usage in response to fee variation.

We exclude C5 and C0 from causal interpretation. C5 (28.0%, median high-fee fraction = 0.90) yields a positive estimate ($+0.0314$, $p = 0.21$), consistent with exclusion restriction failure: these wallets time activity to congestion periods, so the lagged fee captures position during the previous congestion period. C0 (11.3%) yields $\hat{\beta} = -1.163$, an anomalously large magnitude that is driven by very low active hours and is not identified.

The pooled estimate includes the C0 and C5 wallets whose instruments are potentially invalid (Section 4.3). Since C5 contributes a near-zero positive estimate and C0 an anomalously large negative one, the pooled -0.036 is an average over heterogeneous reliability; the valid-cluster estimates in Table 7 provide the within-cluster causal figures.

A natural alternative identification strategy would exploit the $0.01 \rightarrow 0.02$ gwei base fee floor increase mid-sample as a quasi-experimental discontinuity. This change, however, binds only for transactions priced at or near the floor – primarily spam and dust activity – while the 500-transaction wallets in our sample operate well above the floor throughout. The floor increase therefore does not induce a sharp discontinuity in the fee faced by our study population and cannot serve as an instrument for this sample.

The two-way demeaned binscatter confirms a negative demand slope after removing wallet and hour fixed effects and log transaction count (see full version).

8 Results: Arbitrum L2 Per-Resource Demand

8.1 Per-Dimension Pooled Estimates

Table 8 reports IV elasticity estimates for each of the seven Arbitrum resource dimensions using the pooled sample. The estimates span a wide range across resource dimensions. These are intensive-margin elasticities, estimated on wallet-hours with positive consumption of

■ **Table 8** L2 per-resource IV elasticity estimates, pooled sample. Intensive margin (wallet-hours with $\text{gas}_d > 0$). $*p < 0.05$, $***p < 0.001$; n.s. not significant at 5%.

Dimension	Wallets	$\hat{\beta}_{\text{OLS}}$	$\hat{\beta}_{\text{IV}}$	SE	Partial F
computation	23,119	-0.0026	-0.0274*	0.0123	118,405
historyGrowth	20,052	-0.0109	-0.0109 n.s.	0.0304	46,410
storageAccessRead	22,652	+0.0166	+0.0344*	0.0153	118,185
storageAccessWrite	20,883	+0.0040	+0.1186***	0.0274	47,843
storageGrowth	18,990	-0.0388	-0.1510***	0.0446	26,104
l2Calldata	20,799	-0.0125	-0.0575*	0.0228	71,727
refund	17,770	-0.0232	-0.2688***	0.0414	30,666

the given resource; they describe the response of wallets that continue to use a resource, not the economy-wide change in its total consumption. Computation gas, the dominant component of most transactions, is modestly elastic at -0.0274^* ($\text{SE} = 0.0123$). Storage growth and refunds are the most price-sensitive: -0.1510^{***} ($\text{SE} = 0.0446$) and -0.2688^{***} ($\text{SE} = 0.0414$) respectively. L2 calldata shows a significant negative elasticity of -0.0575^* ($\text{SE} = 0.0228$). History growth is statistically insignificant from zero ($p = 0.72$).

8.2 Positive Elasticities

The positive IV estimates for storage access reads ($+0.034^*$) and writes ($+0.119^{***}$) do not represent a genuine demand response. When fees rise, wallets with lower-necessity operations exit the sample. This leaves a selected set of wallets whose operations require storage access. In such periods, the conditional average storage access rises even as total volume falls, producing a spurious positive elasticity. Because this selection operates whenever fees rise, it biases every per-resource estimate in the positive direction. The dimensions that remain significantly negative, storage growth and refunds, are therefore conservative: the true intensive-margin response is at least as large as the estimates reported here. The small or positive estimates for computation, reads, and writes should be read as responses among continuing wallets rather than economy-wide demand elasticities.

The share of active wallet-hours using storage reads is effectively constant across base-fee deciles (0.951–0.956), confirming that read-intensive wallets are the last to exit when fees rise and dominate the remaining sample at high fees. Storage writes show a clearer decline (0.835 at the lowest fee decile to 0.773 at the highest), consistent with write-light wallets exiting while write-intensive operations persist. By contrast, the dimensions with the most negative IV estimates, storage growth and refunds, show no systematic selection pattern across deciles, supporting their interpretation as genuine demand responses rather than selection artifacts.


8.3 Most Elastic Resources: Storage Growth and Refunds

Storage growth ($\hat{\beta} = -0.1510^{***}$) and refunds ($\hat{\beta} = -0.2688^{***}$) are the most price-sensitive resource dimensions in the pooled estimate. Both effects are concentrated in C3 of the per-resource clustering (35.0% of the sample): -0.3709^{***} for storage growth ($\text{SE} = 0.0549$) and -0.4157^{***} for refunds ($\text{SE} = 0.0514$).

The wallets in this cluster have routine operations for which creating new storage slots and storage cleanup are not time sensitive. When fees rise, they defer these operations, which lead to a measurable reduction in usage of these resources.

■ **Table 9** L2 per-resource IV $\hat{\beta}$ for key clusters. * $p < 0.05$, *** $p < 0.001$; n.s. not significant at 5%.

Dimension	C2 high-volume (12.8%)	C3 general DeFi (35.0%)
computation	−0.1590***	−0.0232 n.s.
historyGrowth	−0.6859***	−0.0354 n.s.
storageAccessRead	−0.3189***	−0.0274 n.s.
storageAccessWrite	−0.3207***	+0.0069 n.s.
storageGrowth	+0.1274 n.s.	−0.3709***
l2Calldata	−0.3560***	+0.0107 n.s.
refund	+0.0781 n.s.	−0.4157***



■ **Figure 5** Median gas usage by resource dimension for the pooled Arbitrum L2 model (average regime, fee = 0.05 gwei). Error bars show the inter-quartile range. Blue bars have negative $\hat{\beta}$ (demand response); red bars have positive $\hat{\beta}$.


8.4 Cluster-Level Per-Dimension Results

Table 9 presents cluster-level estimates for the two clusters with the most interpretable IV results. C2 (12.8%) shows large negative elasticities across nearly all dimensions except storage growth and refunds, confirming price elasticity in this segment. C3 (35.0%) is near-zero across most dimensions but strongly negative for storage growth and refunds, consistent with the pattern observed in the pooled sample.

Figure 5 summarizes median gas usage and the pooled IV elasticity across dimensions.

9 Downstream applications

The IV estimates from Section 4.3 include the causal parameters of a demand model that can be used to simulate gas usage by the wallets in various gas fee regimes. In this section, we describe how the full demand model is used to generate synthetic gas demand distributions under counterfactual fee regimes. Such simulations help evaluate the performance of fee mechanism designs beyond the observed range. We present the construction of the demand model here. Figure 6 illustrates the model applied to a fee sweep from 0.005 to 5 gwei



■ **Figure 6** Simulated gas demand under counterfactual base fees, **average** regime. Bold line: median. Inner band: 25th–75th percentile. Outer band: 10th–90th percentile. Dotted vertical line: observed median base fee (0.02 gwei).

for the **average** regime ($\hat{\alpha} \in [P_{40}, P_{60}]$, $\hat{\gamma} \in [P_{40}, P_{60}]$, $n_{\text{txs}} = 3$). The downward-sloping median response reflects the pooled L2 elasticity ($\hat{\beta} = -0.036$) across the full fee range; the interquartile and 10th–90th percentile envelopes shift in parallel, confirming that the elasticity acts multiplicatively on the gas distribution rather than compressing or expanding it. A systematic evaluation of specific mechanism parameters is left to future work.

Recovering Fixed Effects

The 2SLS method described in Section 4.4 estimates only $\hat{\beta}$ and $\hat{\phi}$ of the structural demand equation (4). The wallet and time fixed effects are recovered by computing the residuals of the outcome and applying alternating projections. The residuals are then fitted with a Student- t distribution to accommodate heavy tails that are commonly observed in gas usage. The alternating projections procedure and the sampling algorithm used to generate counterfactual gas demand distributions are detailed in the full version.

10 Discussion

10.1 Why L2 Is More Elastic Than L1

Several factors may contribute to the higher measured elasticity on L2. The two chains attract different wallet populations: L2 users have opted into a lower-cost environment and may be more price-sensitive by selection, whereas L1 wallets include a higher share of protocol-level operations that are insensitive to fee changes. Arbitrum’s base fee floor and gas ladder also create a distinct incentive environment compared to Ethereum. Finally, the difference in time granularity between the two estimates, daily fixed effects on L1 versus hourly on L2, may mechanically affect the measured elasticity in ways that are difficult to disentangle from a genuine behavioral difference.

One particular difference is the relative absence of a large share of reverted transactions on L1, unlike L2s, where the sequencer's mempool is kept private and arbitrageurs try to probabilistically backrun state changes. In times of high asset price volatility, which results in higher on-chain fees, such transactions dominate on-chain gas consumption. On L1, the task of filtering reverted transactions is delegated to the block builders.

10.2 Implications for Fee Mechanism Design

Ethereum and Arbitrum both use a variant of the same congestion management fee mechanism, namely, converting all resource usage into a single gas unit and pricing the gas according to its own EIP-1559 parameters. Note that not all resources are the same. Namely, state growth is a long-term issue for the chain: if the state grows too much, accessing and writing to state memory become slower, and it may no longer fit on disk. Other resources have more immediate constraints. Storage capacity not allocated at a given time can be allocated later, whereas computation and state-access resources not consumed in a given period cannot be recovered. These distinctions call for at least different ratios between limit and target per resource, instead of a constant ratio imposed by converting all resources to a single gas unit, to better utilize the chain. Fully-fledged multi-dimensional resource pricing schemes have been proposed and analyzed by a number of industry and academic works. Our findings on multi-dimensional resource elasticities can help inform optimal parameter ranges for each resource, given EIP-1559-type mechanisms are applied to each resource independently. As an illustration, the comparatively elastic refund and storage-growth dimensions (pooled -0.27 and -0.15) imply that pricing these resources more aggressively would meaningfully curb their use, whereas repricing computation (near-zero elasticity) would mainly raise revenue without changing usage. A multi-dimensional mechanism could therefore apply tighter independent pricing to the elastic, state-affecting resources, where price has traction.

10.3 Limitations

The IV estimator identifies a local average treatment effect: the demand response of wallets whose current fee is shifted by the lagged fee. Always-on wallets (C3 on L1) and congestion-timing wallets (C5 on L2) are excluded from causal interpretation; the pooled estimates do not characterize these populations. The L2 sample covers seven months (October 2025 through April 2026) and reflects the fee regime in effect during this window; as Arbitrum's base fee floor and gas ladder parameters evolve, the measured elasticities may shift. The 500-transaction minimum threshold ensures sufficient within-wallet variation for fixed-effects estimation but excludes low-frequency users, and the estimated elasticities may not generalize to that population. Finally, the per-resource decomposition is feasible for Arbitrum because the Nitro node tracks per-category gas consumption; no equivalent data exists for Ethereum L1.

The methodology applies to any chain with a congestion-responsive fee adjustment rule (an EIP-1559 variant or equivalent), sufficient wallet-level panel data, and temporal fee variation above the protocol floor. For chains with fixed-price auctions or purely flat fees, the lagged-fee instrument loses relevance and a different identification strategy would be required. Chains where the base fee spends most of its time at the floor (as L2 does during low-congestion periods) should use the partial R^2 alongside the partial F to assess instrument strength, as shown here for Arbitrum.

11 Related Work

This work draws on, and contributes to, four research areas: theoretical and empirical study of transaction fee mechanisms, multi-dimensional resource pricing, L2 fee economics, and the empirical estimation of demand elasticity under congestion pricing.

11.1 Transaction Fee Mechanism Design

Roughgarden’s analysis of EIP-1559 [26, 25] establishes the design objectives of the protocol’s base-fee adjustment rule and characterizes its incentive properties. Ferreira et al. [13] develop a class of dynamic posted-price mechanisms for blockspace, and Chung and Shi [7] provide a unified treatment of impossibility results for off-chain incentive-compatible transaction fee mechanisms. Leonardos et al. [18] give a control-theoretic characterization of base-fee dynamics under stochastic demand. The original proposal [6] motivates the mechanism in terms of wallet-level fee predictability, the property that underpins the present empirical strategy: predictable fees make the lagged base fee a relevant instrument for the current fee.

11.2 Empirical Studies of Gas Fees and Demand

Reijsbergen et al. [24] examine the first month after EIP-1559 deployment and document a substantial reduction in fee volatility with limited change in median fees. Liu et al. [19] extend this horizon and quantify the effect of the mechanism on waiting times and consensus security. Pierro and Rocha [23] provide an earlier empirical decomposition of gas-fee covariates including pending transaction count and exogenous costs. The closest predecessor to the present elasticity estimate is Donmez and Karaivanov [11], who use block-level data from the pre-EIP-1559 first-price regime and a queueing-theoretic specification to identify a strongly nonlinear utilization-price relationship, with sharp fee response above 90% block fullness. Their analysis predates EIP-1559, does not exploit wallet-level variation, and does not separate intensive from extensive margins; the present work addresses each of these in the post-EIP-1559 regime and additionally provides per-resource estimates for an L2. Silva [27] studies multi-resource demand elasticities on Ethereum, complementing our data collection; the focus is on the effects on per-resource demand as the Ethereum gas target increases.

11.3 Multi-Dimensional Resource Pricing

Diamandis et al. [10] introduce multi-dimensional EIP-1559-style pricing for non-fungible blockchain resources and characterize the equilibrium fee schedule. Angeris et al. [1] prove that multi-dimensional fees are essentially optimal in a broad class of demand models. Lavee et al. [17] and Kiayias et al. [16] compare the throughput and convergence properties of multi-dimensional versus single-dimensional pricing; Lavee et al. study stationary throughput and find that the ranking depends on demand functions, while Kiayias et al. take a dynamic approach and argue that a single-dimensional mechanism may reach equilibrium prices faster. Crapis et al. [9] study optimal design of a multi-dimensional pricing mechanism when resource usages are pairwise correlated.

11.4 Rollup Economics and MEV on L1/L2

Monnot [22] and Crapis [8] provide foundational economic analyses of rollup fee markets. Crapis et al. [32] study optimal strategies of L2 data posting on L1 and effects on end users.

Solmaz et al. [28] empirically find that a large percentage of L2 gas usage is due to optimistic MEV, which is not observed on L1 since the block builders can and do filter out reverted transactions. This phenomenon is theoretically studied in Mazorra et al. [21] and Wang et al. [33], where a similar conclusion is made: given fixed gas costs, probabilistic backrunning searchers will spend all expected value in reverted transactions (sometimes classified as spam).

11.5 Wallet Identification and Behavioral Clustering

In an account-based blockchain such as Ethereum, identifying wallet type is a prerequisite for any heterogeneity analysis. Victor [31] introduces the standard address-clustering heuristics for Ethereum. Béres et al. [4] extend this with timing- and graph-based features for finer profiling of addresses. Hu et al. [15] provide recent feature-engineering work specifically aimed at separating financial bots and MEV searchers from human users. The clustering procedure in Section 5 adopts features in this tradition (transaction frequency, gas variance, and high-fee timing) and augments them with per-resource consumption fractions for the multi-dimensional analysis.

11.6 Demand Elasticity under Congestion Pricing

The methodological framework used to estimate elasticity in this work has direct precedence in the electricity pricing literature. Borenstein [5] establishes that even small demand elasticities (around -0.025 , a range comparable to the L1 elasticity estimate) translate to substantial welfare gains under dynamic pricing. Fabra et al. [12] use household-level panel data and a renewable-supply instrument to estimate elasticity under Spain’s real-time pricing tariff, providing the closest methodological analogue to our wallet-level panel analysis. Tiedemann et al. [30] show that standard IV is inconsistent when high-frequency time series exhibit autocorrelation in both price and demand, informing the within-period exclusion concerns discussed in Section 4.3. The foundations of IV analysis are described in [2, 29, 34, 3].

12 Conclusion

We estimate the IV elasticities of gas demand for Ethereum L1 and Arbitrum L2. Gas demand is inelastic on both networks: pooled IV estimates of -0.006 (L1) and -0.036 (L2) imply that a 10% fee increase reduces total demand by at most 0.4%. Within L2, per-resource decomposition reveals that the aggregate inelasticity masks substantial resource-level heterogeneity: refunds and new storage slot creation respond meaningfully to price (-0.27 and -0.15 respectively), while computation is near-inelastic (-0.03).

References

- 1 Guillermo Angeris, Theo Diamandis, and Ciamac Moallemi. Multidimensional Blockchain Fees Are (Essentially) Optimal. In *7th Conference on Advances in Financial Technologies (AFT 2025)*, volume 354 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 24:1–24:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.AFT.2025.24.
- 2 Joshua D. Angrist and Alan B. Krueger. Does Compulsory School Attendance Affect Schooling and Earnings? *The Quarterly Journal of Economics*, 106(4):979–1014, 1991. doi:10.2307/2937954.

- 3 Joshua D. Angrist and Jörn-Steffen Pischke. *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press, 2009.
- 4 Ferenc Béres, István András Seres, András A. Benczúr, and Mikerah Quintyne-Collins. Blockchain Is Watching You: Profiling and Deanonimizing Ethereum Users. In *IEEE International Conference on Decentralized Applications and Infrastructures (DAPPS)*, pages 69–78, 2021. [arXiv:2005.14051](#).
- 5 Severin Borenstein. The Long-Run Efficiency of Real-Time Electricity Pricing. *The Energy Journal*, 26(3):93–116, 2005.
- 6 Vitalik Buterin, Eric Conner, Rick Dudley, Matthew Slipper, Ian Norden, and Abdelhamid Bakhta. EIP-1559: Fee Market Change for ETH 1.0 Chain. Ethereum Improvement Proposal, 2019. URL: <https://eips.ethereum.org/EIPS/eip-1559>.
- 7 Hao Chung and Elaine Shi. Foundations of transaction fee mechanism design. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2023)*, 2023. [arXiv:2111.03151](#).
- 8 Davide Crapis. Rollups are Real: Rollup Economics 2.0. Notion, 2023. URL: <https://daveidecrapis.notion.site/Rollups-are-Real-Rollup-Economics-2-0-2516079f62a745b598133a101ba5a3de>.
- 9 Davide Crapis, Ciamac C. Moallemi, and Shouqiao Wang. Optimal dynamic fees for blockchain resources. In Jeremy Clark and Elaine Shi, editors, *Financial Cryptography and Data Security - 28th International Conference, FC 2024, Willemstad, Curaçao, March 4-8, 2024, Revised Selected Papers, Part I*, Lecture Notes in Computer Science, pages 271–291. Springer, 2024. doi:10.1007/978-3-031-78676-1_15.
- 10 Theo Diamandis, Alex Evans, Tarun Chitra, and Guillermo Angeris. Dynamic Pricing for Non-Fungible Resources: Designing Multidimensional Blockchain Fee Markets. In *5th Conference on Advances in Financial Technologies (AFT 2023)*, volume 282 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 4:1–4:23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.AFT.2023.4.
- 11 Anil Donmez and Alexander Karaivanov. Transaction Fee Economics in the Ethereum Blockchain. *Economic Inquiry*, 60(1):265–292, 2022. doi:10.1111/ecin.13025.
- 12 Natalia Fabra, David Rapson, Mar Reguant, and Jingyuan Wang. Estimating the Elasticity to Real-Time Pricing: Evidence from the Spanish Electricity Market. *AEA Papers and Proceedings*, 111:425–429, 2021. doi:10.1257/pandp.20211007.
- 13 Matheus V.X. Ferreira, Daniel J. Moroz, David C. Parkes, and Mitchell Stern. Dynamic posted-price mechanisms for the blockchain transaction-fee market. In *Proceedings of the 3rd ACM Conference on Advances in Financial Technologies*, 2021. doi:10.1145/3479722.3480993.
- 14 Ragnar Frisch and Frederick V. Waugh. Partial time regressions as compared with individual trends. *Econometrica*, 1(4):387–401, 1933. doi:10.2307/1907330.
- 15 Hanzhi Hu, Tom Beer, and Daniel Perez. Detecting Financial Bots on the Ethereum Blockchain. *CoRR*, abs/2403.19530, 2024. doi:10.48550/arXiv.2403.19530.
- 16 Aggelos Kiayias, Elias Koutsoupias, Giorgos Panagiotakos, and Kyriaki Zioga. One-dimensional vs. multi-dimensional pricing in blockchain protocols. *CoRR*, abs/2506.13271, 2025. doi:10.48550/arXiv.2506.13271.
- 17 Nir Lavee, Noam Nisan, Mallesh M. Pai, and Max Resnick. Does your blockchain need multidimensional transaction fees? *CoRR*, abs/2504.15438, 2025. doi:10.48550/arXiv.2504.15438.
- 18 Stefanos Leonardos, Barnabé Monnot, Danijar Reijsbergen, Efstratios Skoulakis, and Georgios Piliouras. Dynamical Analysis of the EIP-1559 Ethereum Fee Market. *CoRR*, abs/2102.10567, 2021. [arXiv:2102.10567](#).
- 19 Yulin Liu, Yuxuan Lu, Kartik Nayak, Fan Zhang, Luyao Zhang, and Yinhong Zhao. Empirical Analysis of EIP-1559: Transaction Fees, Waiting Times, and Consensus Security. In *Proceedings of the 2022 ACM SIGSAC Conference on Computer and Communications Security*, 2022. doi:10.1145/3548606.3559341.

- 20 Michael C. Lovell. Seasonal adjustment of economic time series and multiple regression analysis. *Journal of the American Statistical Association*, 58(304):993–1010, 1963. doi:10.1080/01621459.1963.10480682.
- 21 Bruno Mazorra, Christoph Schlegel, and Akaki Mamageishvili. Timing games: Probabilistic backrunning and spam. *CoRR*, abs/2602.22032, 2026. doi:10.48550/arXiv.2602.22032.
- 22 Barnabé Monnot. Understanding Rollup Economics from First Principles. Substack, 2022. URL: <https://barnabe.substack.com/p/understanding-rollup-economics-from>.
- 23 Giuseppe Antonio Pierro and Henrique Rocha. The Influence Factors on Ethereum Transaction Fees. In *2nd International Workshop on Emerging Trends in Software Engineering for Blockchain (WETSEB)*, pages 24–31, 2019. doi:10.1109/WETSEB.2019.00010.
- 24 Danijar Reijsbergen, Stefanos Leonardos, Barnabé Monnot, Stratis Skoulakis, and Georgios Piliouras. Transaction Fees on a Honeymoon: Ethereum’s EIP-1559 One Month Later. In *2021 IEEE International Conference on Blockchain (Blockchain)*, 2021. doi:10.1109/Blockchain53845.2021.00049.
- 25 Tim Roughgarden. Transaction Fee Mechanism Design. In *Proceedings of the 22nd ACM Conference on Economics and Computation (EC 2021)*, 2021. doi:10.1145/3465456.3467660.
- 26 Tim Roughgarden. Transaction Fee Mechanism Design for the Ethereum Blockchain: An Economic Analysis of EIP-1559. *Cryptoeconomic Systems*, 2021. arXiv:2012.00854.
- 27 Maria Silva. Empirical Analysis of Price Elasticities for Ethereum State and Burst Resources. Ethereum Research, 2026. URL: <https://ethresear.ch/t/empirical-analysis-of-price-elasticities-for-ethereum-state-and-burst-resources/24166>.
- 28 Ozan Solmaz, Lioba Heimbach, Yann Vonlanthen, and Roger Wattenhofer. Optimistic MEV in ethereum layer 2s: Why blockspace is always in demand. In Zeta Avarikioti and Nicolas Christin, editors, *7th Conference on Advances in Financial Technologies, AFT 2025, Pittsburgh, PA, USA, October 8-10, 2025*, LIPIcs, pages 28:1–28:24. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.AFT.2025.28.
- 29 Douglas Staiger and James H. Stock. Instrumental Variables Regression with Weak Instruments. *Econometrica*, 65(3):557–586, 1997. doi:10.2307/2171753.
- 30 Silvana Tiedemann, Raffaele Sgarlato, and Lion Hirth. Price Elasticity of Electricity Demand: Using Instrumental Variable Regressions to Address Endogeneity and Autocorrelation of High-Frequency Time Series. *Energy Economics*, 2024. arXiv:2306.12863.
- 31 Friedhelm Victor. Address Clustering Heuristics for Ethereum. In *Financial Cryptography and Data Security (FC) 2020*, volume 12059 of *Lecture Notes in Computer Science*, pages 617–633. Springer, 2020. doi:10.1007/978-3-030-51280-4_33.
- 32 Shouqiao Wang, Davide Cripis, and Ciamac C. Moallemi. A framework for combined transaction posting and pricing for layer 2 blockchains. In Christina Garman and Pedro Moreno-Sanchez, editors, *Financial Cryptography and Data Security - 29th International Conference, FC 2025, Miyakojima, Japan, April 14-18, 2025, Revised Selected Papers, Part I*, Lecture Notes in Computer Science, pages 56–72. Springer, 2025. doi:10.1007/978-3-032-07024-1_4.
- 33 Wenhao Wang, Aditya Saraf, Lioba Heimbach, Kushal Babel, and Fan Zhang. Blockspace under pressure: An analysis of spam MEV on high-throughput blockchains. *CoRR*, abs/2604.00234, 2026. doi:10.48550/arXiv.2604.00234.
- 34 Jeffrey M. Wooldridge. *Econometric Analysis of Cross Section and Panel Data*. MIT Press, 2nd edition, 2010.