A Macro Finance Model for Proof-of-Stake Ethereum

Urban J. Jermann*

Wharton School of the University of Pennsylvania and NBER

September 6, 2023

Abstract

With the switch to proof-of-stake, Ethereum has implemented fundamental changes to its monetary and fee policies. This paper studies the money supply, fees, and prices under this new regime within a dynamic macro finance model. The model's long run properties are characterized analytically. The fraction of staked ETH is shown to depend on a single parameter representing the utility value of using ETH relative to the value of computational effort on the network. The market capitalization in terms of the dollar numeraire is shown to be invariant to the staking reward factor. For a preliminary quantitative analysis, the model is calibrated with the short sample of available price and fee data. The share of ETH staked in the long-run is estimated to be 26%. The implied long-run money supply can plausibly be below the current level. The potential welfare gains from alternative and more activist staking yield policies are likely to be unpredictable and small.

Keywords: Cryptocurrencies, EIP1559, money supply. JEL codes: E42, G12.

1 Introduction

In September 2022, the Ethereum network switched to a proof-of-stake consensus mechanism. At that occasion, the rate at which ETH is issued has been significantly lowered. A year before that, Ethereum had implemented changes to its fee policies resulting in a majority of fee income flowing to holders of ETH. A priori one would expect these changes to have a positive impact on the price of ETH. Given market volatility, such an impact has been hard

*I am grateful for comments from Tom Winberry, Yao Zeng, Will Cong, seminar participants at USC, UT Dallas, University of Wisconsin, Oxford, Warwick, LBS, and for research assistance from Max Yang. The most recent version of this paper is available at http://finance.wharton.upenn.edu/~jermann/. Email: jermann@wharton.upenn.edu.

to detect. This paper studies money supply, fee, and price dynamics under this new regime with tools from macro finance.

In proof-of-work blockchains such as Bitcoin and Ethereum before the recent changes, miners are compensated with newly issued tokens and transaction fee income. This expense is necessary for the security of the network. With proof-of-stake, the network is secured by the collateral that is staked and that needs to be compensated. This offers an opportunity for long-term investors to receive yield payments on their staked ETH. It remains an open question how much ETH will be staked in the long run.

In this paper, I present a dynamic stochastic equilibrium model where ETH is bought to either be used or staked. Using ETH provides utility; staking ETH gives financial rewards. Agents also use the computational resources of the Ethereum network. The prices of ETH and of computational services are taken as given by individual agents. These prices are determined in equilibrium to clear markets. The supply of ETH follows the rules spelled out by the Ethereum protocol. The model is driven by exogenous shocks to the productivity or adoption, shocks to agents' stochastic discount factor (interest rates) and shocks to the demand for fast transactions.

The long-run properties of the model are characterized analytically. At the model's steady state, the share of ETH staked depends on a single parameter capturing the utility value of computational services relative to the usage value of ETH. Intuitively, if using ETH provides a relatively large utility value, more ETH will be held by users in the long run, and this leads to a smaller share that is staked. None of the other parameters matter in the long run because they imply equal proportional changes in the amounts of ETH that are used and staked.

The analysis solves for the steady state money supply and price level as a function of a few model parameters. In the model, the validator reward factor (effectively, the base_reward_factor) is irrelevant in equilibrium. At steady state, the money supply increases in this factor but the price of ETH declines proportionally, so as to leave the market capitalization in the dollar numeraire unaffected. More generally, this factor–if constant–is shown to be irrelevant for model dynamics.

For a preliminary quantitative analysis, model parameters are calibrated with data on ETH prices and fees starting in August 2021 when the new fee policies were implemented. With additional data from Lido's liquid staking token stETH, the long-run share of staked ETH is estimated at 26% with a relatively tight statistical confidence interval. Other large proof-of-stake networks (e.g. Cardano, Solana) have larger staking shares. From the perspective of the model, a small staking share reflects the fact that the utility to users from holding ETH is more valuable than the utility from the computational services of the network. Based on this estimate and the steady state solution, the implied long-run money supply can plausibly be below its current level. As time goes by more data will become available and more accurate estimates can be obtained.

In the calibrated model, the price of ETH is mostly driven by shocks to productivity/adoption; fluctuations in fees and the amounts of ETH staked and used are driven primarily by discount rate shocks. The aggregate money supply in the model is stationary, but the mean reversion after shocks is estimated to be extremely slow. Based on experiments with the quantitative model, the potential welfare gains from alternative and more activist staking yield policies can be unpredictable and are likely to be small. However, such policies can have first-order effects on the volatility of the staking share.

To the best of my knowledge, this is the first dynamic equilibrium model for proof-of-stake Ethereum. The paper also contributes to a growing literature on economic analyses of proof-of-stake blockchains. Cong, He and Tang (2022) study staking as a function of rewards within a dynamic portfolio model and present cross-sectional evidence on a set of cryptocurrencies. Kose, Rivera and Saleh (2021) theoretically analyze equilibrium staking levels when investors have heterogenous trading horizons. Fanti, Kogan and Viswanath (2019) develop a model for valuing tokens in proof-of-stake payment systems based on given processes for fee income and money creation. Saleh (2021) studies conditions under which proof-of-stake generates consensus. Different from these, my model has a number of Ethereum-specific elements, features jointly determined money supply and fees, and offers quantitative perspectives.

The popular site https://ultrasound.money/ presents long-run projections for the Ethereum money supply based on the model from Elowsson (2021). The model represents the steady state of the ETH supply based on assumptions about the staking reward yield and the burn rate. In my model, these two are determined endogenously based on the assumptions on user preferences and technologies combined with market clearing and intertemporal optimization. This model is nested by my model, and I present a comparison in the paper.

Dynamic economic models for cryptocurrencies more generally have been developed and studied for instance in Biais, et al (2019), Cong, Ye and Wang (2022), Jermann (2021), Jermann and Xiang (2022), Li and Mayer (2020), and Mei and Sockin (2022). On some of the recent changes in Ethereum see Kim (2021), Bitmex Research (2021), Migalabs (2022), and Wahrstatter (2022).

Section 2 provides some background on the recent changes in Ethereum. Section 3 presents the model and analytical characterization. The quantitative analysis is in Section 4 and policies are studied in Section 5.

2 Fees and rewards to staking

I start by describing elements of fee policies and rewards to staking, and their relation to the supply of ETH. In addition to providing institutional background, this motivates some of my model design choices.

Under the current regime, supply changes in ETH are determined by two components,

$$\text{Change in ETH supply} = \text{Validator rewards} - \text{Burnt fees}.$$

Validators of blocks that participate in the consensus process receive compensation rewarding them for their costs and efforts. These rewards are paid out with newly issued ETH and add to the money supply. Fees are paid to execute transactions on the blockchain. A majority of these fees end up being burnt. That is, the ETH collected as fee income is taken out of circulation and this reduces the outstanding supply. I will start describing fee policies in more detail and then consider the rewards to staking.

2.1 Fee policies

In response to high and volatile fees, and to allocate block space more efficiently, fee policies have been upgraded under what is known as Ethereum Improvement Proposal (EIP) 1559. This upgrade was introduced through the London hard fork on August 5, 2021. Under this system, fees for executing transactions on the blockchain have two components: a base fee and an priority fee (or tip). The base fee is burnt and the priority fee goes to validators.

The base fee is determined based on network activity by the following formula

$$f_n^b = \left[1 + 0.125\left(\frac{G_{n-1}}{G^*} - 1\right)\right] f_{n-1}^b,$$

(https://ethereum.org/en/developers/docs/gas/). The base fee in block n, f_n^b, equals the base fee from the previous block adjusted in proportion to how much network activity G_{n-1} deviates from the target level G^*. Network activity is measured in term of *gas*, G, the unit measuring the amount of computational effort to execute specific operations on the Ethereum network. For instance, a simple transfer of ETH from one account to another requires $21,000$ units of gas. Based on this formula, the base fee increases when activity is above target and decreases when activity is below target.

The block size target G^* is set to 15 million units of gas, with a maximum block size of 30 million units of gas. If there is a burst in demand for block space that leads to blocks reaching their maximum size, the base fee increases by 12.5%. A new block is produced every 12 seconds. Assuming that blocks are full say ten times in a row would produce a base fee change by a factor of $1.125^{10} = 3.25$. Clearly, this mechanism has to potential to temper bursts in demand relatively quickly. One would expect the level of activity over the medium term to remain close to the target level.

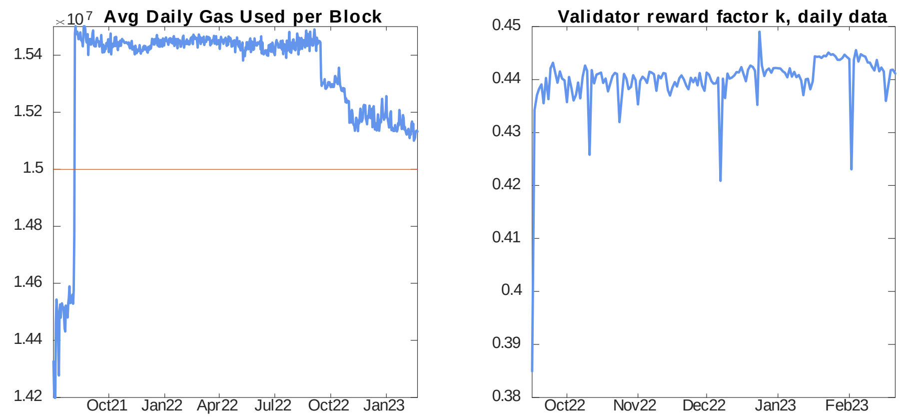

Figure 1.

As shown in Figure 1, left panel, based on data from etherscan.io, the average daily amount of gas used per block has been close to the target level. Starting with the implementation of EIP 1559 August 5, 2021, the average has been about 3% above target and most recently it is running about 1.5% above target. In my simulations with a static demand subject to shocks, the nonlinearities in the formula can generate such deviations from target depending on the volatility and persistence of the shocks. In my model presented below, a period is set to be one month. In line with the data, I will assume that there is a fixed supply of block space available, and in the model the gas price adjusts to clear the market.

The priority fee is an optional payment through which users can speed up their transactions. As discussed in more detail below, priority is essential for a class of transactions grouped under the term MEV (maximal extractable value). Such transactions can seek to frontrun others, and this can contribute to higher priority fees.

2.2 Staking ETH

To be a validator of transactions, ETH needs to be staked and serves as collateral to insure that validators behave honestly and responsively. Staking is compensated with income from three sources: rewards for activities associated with validating transactions, transaction priority fees, and MEV (maximal extractable value).

Validating transactions involves a number of tasks that are compensated according to the rules of the Ethereum protocol. The most significant is the proposal reward which goes to a validator that is randomly selected to propose a block. The average reward received by a validator is proportional to the staked amount M_i^S divided by the square root of the total amount staked by all validators

$$k\frac{M_i^S}{\sqrt{\sum_j M_j^S}}, \tag{1}$$

with constant k. The actual amount received by an individual validator depends on the tasks they are selected for; for instance, there is only one proposer per block. The exact amount received also depends on the performance of the execution. These rewards are paid with newly created ETH. According to Buterin (2022), scaling by the square-root of the aggregate amount strikes a compromise between a constant per unit reward kM_i^S and a constant total reward $kM_i^S/\sum_j M_j^S$, and avoids the worst consequences of each one. The former would make the total amount staked potentially very uncertain and could compromise the security of the network if only a small share of ETH is staked. It could also result in a very volatile total issuance. The later could incentivize manipulation of the consensus process to discourage participation so that manipulators could collect a larger share of a constant amount of ETH issued (Buterin (2018)). This formula makes the aggregate amount paid $k\sqrt{\sum_j M_j^S}$ increasing in the total amount staked but at a decreasing rate.

Rewards are available for validating each block. With blocks created every 12 seconds, there are 7200 blocks created per day. If the performance of the validators does not fluctuate too much, the formula should also hold approximately for reported daily staking amounts for some constant. To verify that and to guide the specification of my model, I use daily data from etherscan.io to generate such daily values for k as

$$k\left(t\right)=\frac{\left[M\left(t\right)-M\left(t-1\right)\right]+\text{FeesBurnt}(t)}{\sqrt{M^S\left(t\right)}}$$

with $M\left(t\right)$ the aggregate supply at date t. As shown in Figure 1, right panel, the implied daily k has experienced only small fluctuations around a stable mean. The coefficient of variation is below 1%, at .007. As shown below, in the model, the coefficient k impacts the levels of the equilibrium money supply and the price of ETH. However, equilibrium staking shares and the dynamic behavior of the logs of money and prices are invariant to k.

Staking ETH has been enabled since the end of 2020. Withdrawing staked ETH has been possible only since April 2023. As shown in Figure 2, in May 2023, about 15% of the supply of ETH is staked. The amount of staked ETH has increased monotonically until unstaking was enabled. After a short period where withdrawals and new deposits were balanced, the staking share continues to increase.[1]

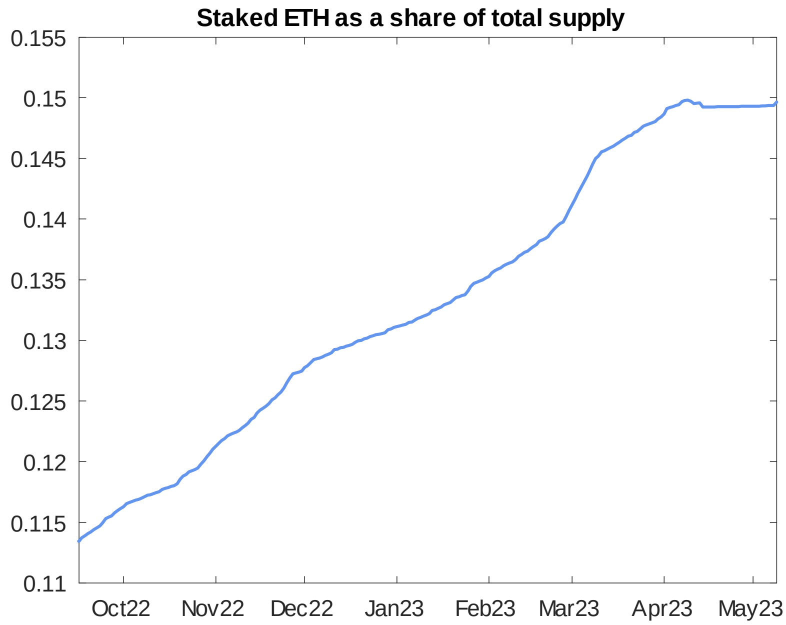

Figure 2.

3 Model

Agents in the model buy ETH to either use it or stake it. Using ETH gives them utility; staking ETH gives them financial rewards. Agents also use gas and derive utility from it. The prices of ETH and gas are taken as given by individual agents. These prices are determined in equilibrium to clear markets. The supply of ETH follows the rules spelled out in EIP1559. The model is driven by exogenous shocks to the productivity/adoption, shocks to agents' stochastic discount factor (interest rates) and shocks to the priority fee driven by the demand for speed.

Time is discrete and indexed by t. Agents maximize

$$\max_{C_{t+j},G_{t+j},M^U_{t+1+j},M^S_{t+1+j}} \sum_{j=0}^{\infty} E_t \beta^j \Lambda_{t+j} \left[C_{t+j} + v\left(A_{t+j}, p_{t+j} M^U_{t+j}, G_{t+j}\right)\right]$$

with the discount factor $\beta\Lambda_{t+j}$ depending on a constant β and stochastic component Λ. C is consumption. The utility function $v(..)$ is increasing in real balances, pM^U, (with p the

[1] The Ethereum protocol restricts daily deposits and withdrawals to the same amounts, 1796 validators each for 32 ETH. This explains the flat portion in April/May 2023.

price in USD per ETH, and M^U ETH holdings), G the quantity of gas, and productivity A which can also be thought of as measuring adoption. This is a standard money-in-utility specification augmented with gas as a second good. It can also be thought of as a production function requiring two inputs, money and computational effort, similarly to a standard production function requiring capital and labor as inputs. Money holdings M_{t+1}^U and M_{t+1}^S are purchased in period t and provide services and rewards in period $t+1$ (they are time$-t$ measurable). This is the standard timing convention used in macroeconomic models.

Productivity/adoption follows an exogenous process

$$\begin{aligned} A_t &= \gamma A_{t-1} \exp(z_t), \\ z_t &= \rho_z z_{t-1} + \sigma_z \varepsilon_{z,t}. \end{aligned} \tag{2}$$

with trend growth rate γ, persistence parameter $|\rho_z| < 1$, $\varepsilon_{z,t}$ an iid random variable with mean zero and standard deviation of 1, and volatility parameter σ_z. The shock z_t has permanent effects on the level of A_t. This specification is intended to capture the significant uncertainty about the future uses and adoption of Ethereum.

In the model, gas is one of the productive inputs for the utility of Ethereum. Based on EIP1559, I assume that the supply of gas is fixed at the target level $G^*$and the fee price per gas f_t clears the gas market. The fee has two components, the base fee and the priority fee. I assume the share of the priority fee in the total fee is represented by $0 < \phi_t < 1$, a function of an exogenous shock $z_{\phi,t}$. Specifically

$$\begin{aligned} \ln \phi_t &= \phi_0 + z_{\phi,t}, \\ z_{\phi,t} &= \rho_\phi z_{\phi,t-1} + \sigma_\phi \varepsilon_{\phi,t}, \end{aligned} \tag{3}$$

with parameters $(\phi_0, \rho_\phi, \sigma_\phi)$.[2] The base fee amount $(1-\phi_t) f_t G_t$ is burnt. The priority fee is available as a staking reward to validators.

In the model, I focus on the financial rewards to validators and abstract from the costs and efforts required. Validation activities are handled to a significant extent by staking pools or exchanges in exchange for a portion of the rewards. For instance, as of the end of 2022 market leader Lido charges 10%.

As discussed in Section 2.2, staking is compensated with income from three sources: rewards for activities associated with validating transactions, transaction priority fees, and

[2] A richer specification where $\phi(.)$ also depends on the gas price f_t is empirically supported with $\partial\phi(.)/\partial f_t < 0$. But it essentially makes no difference quantitatively. Therefore this is omitted.

MEV (maximal extractable value). Based on equation (1), validators receive on average

$$\frac{k}{\sqrt{\bar{M}_t^S}}$$

per unit of staked ETH, with $\bar{M}^S = \int M_j^S dj$ the aggregate amount staked based on a measure one of agents. As further discussed below, it is assumed that the idiosyncratic uncertainty is fully diversified and the validator receives the average payment.

The priority fee income, given by $\phi_t f_t \bar{G}_t$, with $\bar{G}$ the aggregate amount of gas used, is available to the block proposer. The probability that an individual staker is selected as the block proposer equals the individual's staked amount divided by the total amount staked. In practise, a large share of staking is done through pools. In addition, the idiosyncratic uncertainty averages out over time. In the model, it is assumed that the individual uncertainty is fully diversified, so that the priority fee income per unit of staked ETH is

$$\frac{\phi_t f_t \bar{G}_t}{\bar{M}_t^S}.$$

In practise, proposers are often not directly receiving the priority fees, instead the fee recipient of a block is a blockbuilder which in turn makes a direct transfer payment to the proposer that is typically very close (within 2-3%) of the fee. Blockbuilders put together blocks with transactions that benefit from being early in line, possibly to frontrun regular transactions. This is known as MEV, maximal extractable value. Block proposers are incentivized by blockbuilders either through fees or direct payments. With blockbuilding a relatively competitive activity, a majority of the MEV should go to the proposers. Based on this, MEV payments to proposers are not separately modelled but are part of the model's equilibrium priority fees.

An agent's budget constraint summarizes these elements

$$C_t + p_t M_{t+1}^S + p_t M_{t+1}^U = p_t M_t^S \left(1 + \frac{k}{\sqrt{\bar{M}_t^S}} + \frac{\phi_t f_t \bar{G}_t}{\bar{M}_t^S}\right) + p_t M_t^U - p_t f_t G_t + Y_t.$$

Note, the individual agent takes the aggregate $\bar{M}^S$ as given, and in equilibrium $M^S = \bar{M}^S$. Specifically, a measure one of identical individual agents aggregate to a representative agent. Equivalently, the aggregate amount of gas $\bar{G}$ is taken as given by individual agents. For conciseness, going forward, the notation will not explicitly acknowledge this distinction.

Market clearing requires that the aggregate money supply equals the aggregate amounts staked and used

$$M_{t+1} = M_{t+1}^S + M_{t+1}^U.$$

Based on EIP1559, the aggregate money supply in the model evolves as

$$M_{t+1} = M_t + k\sqrt{M_t^S} - (1-\phi_t)\, f_t G_t$$

with $k\sqrt{M^S}$ the aggregate staking reward and $(1-\phi_t)\, f_t G_t$ the burnt base fees. Finally, equilibrium for gas requires the agents demand equals the fixed supply $G_t = G^*$.

The first-order conditions for optimal behavior for G_t, M_{t+1}^U, and M_{t+1}^S are given by

$$\frac{\partial v\left(A_t, p_t M_t^U, G_t\right)}{\partial G_t} = p_t f_t \tag{4}$$

$$p_t = E_t \beta \frac{\Lambda_{t+1}}{\Lambda_t}\left(\frac{\partial v\left(A_{t+1}, p_{t+1} M_{t+1}^U, G_{t+1}\right)}{\partial\left(p_{t+1} M_{t+1}^U\right)} + 1\right) p_{t+1} \tag{5}$$

$$p_t = E_t \beta \frac{\Lambda_{t+1}}{\Lambda_t}\left(\frac{k}{\sqrt{M_{t+1}^S}} + \frac{\phi_{t+1} f_{t+1} G_{t+1}}{M_{t+1}^S} + 1\right) p_{t+1}, \tag{6}$$

The optimality condition for gas, equation (4), characterizes the static choice equalizing the marginal utility benefit to the marginal cost, $p_t f_t$, expressed in the numeraire. Optimal choices for M_{t+1}^U and M_{t+1}^S embed the intertemporal trade-off of buying ETH at the current price p_t in exchange of the expected discounted payoff next period.

For additional analytical analysis and to study the model quantitatively, the utility function needs a functional form. Utility is assumed to be represented by

$$\frac{A}{1-\sigma}\left[\left(\frac{pM^U}{A}\right)^{1-\theta} G^\theta\right]^{1-\sigma}$$

with parameters $0 < \theta < 1$ and $\sigma > 0$. As shown below, this function has homogeneity properties that produce a steady state with constant money holdings and the price of ETH growing at the rate γ. To solve the model, the price of ETH can be scaled by the productivity level. The scaled price is itself stationary, and in the steady it is constant. As an equivalent representation, the productivity term could be moved entirely outside the parentheses as $A^{1-(1-\theta)(1-\sigma)}$.

The marginal utilities are

$$\frac{\partial v\left(A, p_t M_t^U, G_t\right)}{\partial G_t} = \theta \frac{A\left[\left(\frac{pM^U}{A}\right)^{1-\theta} G^\theta\right]^{1-\sigma}}{G}$$

$$\frac{\partial v\left(A, pM^U, G_t\right)}{\partial\left(pM^U\right)} = (1-\theta)\,\frac{A\left[\left(\frac{pM^U}{A}\right)^{1-\theta} G^\theta\right]^{1-\sigma}}{pM^U}.$$

An equilibrium is given by exogenous processes for (A_t, Λ_t, ϕ_t), with price and fee processes (p_t, f_t) such that the processes for $\left(M_t^S, M_t^U, G_t\right)$ satisfy the three first-order conditions, the money supply equation, and market clearing for money and gas.

Note, for $\sigma < 1$, there is a zero-price equilibrium with $p_t = 0$ and arbitrary money holdings and fees. We focus on equilibriums with non-zero prices.

3.1 Steady state

I solve for a steady state where the price of ETH grows at the rate of growth of the exogenous productivity/adoption γ, so that the price scaled by productivity $\hat{p}_t \equiv p_t/A_t$ is constant, and the other variables are constant. This characterizes the long-run properties of the dynamic system.

The steady state is the model solution when the variances of all shocks are set to zero, that is z_t, $z_{\phi,t}$ and the shocks driving the discount factor Λ_t. A system of four equations, the three first-order conditions and the money supply equation, can be solved analytically for the four variables M^S, M^U, $\hat{p}$ and f. After some algebra, see the appendix, the solution is summarized in the following proposition.

Proposition 1 *At a steady state where all shocks are set to* 0

$$\begin{aligned}
M^S &= \left(\frac{\beta\gamma}{1-\beta\gamma}\right)^2\left(\frac{k}{1-\phi}\right)^2 \\
M^U &= \left(\frac{\beta\gamma}{1-\beta\gamma}\right)^2\left(\frac{k}{1-\phi}\right)^2\left(\frac{1-\theta}{\theta}\right) \\
f &= \frac{\beta\gamma}{(1-\beta\gamma)}\left(\frac{k}{1-\phi}\right)^2\frac{1}{G} \\
\hat{p} &= G^{\varrho\theta(1-\sigma)}\left(\frac{\beta\gamma}{1-\beta\gamma}\right)^{\varrho-2}\theta\left(1-\theta\right)^{\varrho-1}\left(\frac{1-\phi}{k}\right)^2
\end{aligned}$$

with $\varrho \equiv 1/\left[1-(1-\theta)(1-\sigma)\right]$ *and*

$$\frac{M^S}{M^U + M^S} = \theta.$$

Based on the utility function defined above

$$\theta = \frac{G_t \frac{\partial v(.)}{\partial G_t}}{G_t \frac{\partial v(.)}{\partial G_t} + p_t M_t^U \frac{\partial v(.)}{\partial p M_t^U}},$$

that is, the usage value share of gas. Specifically, $G\frac{\partial v(.)}{\partial G}$ is the usage (utility) value of gas and $pM^U \frac{\partial v(.)}{\partial p M^U}$ the usage (utility) value of monetary balances. As the proposition shows, equilibrium usage M^U decreases with θ, staking M^S does not depend on θ. The insight from the model is that as gas generates relatively more usage value than ETH holdings, θ is larger, and less ETH is held for usage which increases the share of staked ETH. None of the other parameters, k, ϕ, β and γ, matter because they all produce equal proportional changes in M^S and M^U. For instance, a higher k increases staking but increases usage by the same factor. Not surprisingly, a higher k leads to more staking as the rewards are higher. Less directly, with more staking income, the fee income also needs to be higher at a steady state with constant money supply. The higher fee income requires an increase in usage in equilibrium. Note, this is a comparative statics comparison.

The aggregate supply of ETH at steady state is given by

$$M = M^S + M^U = \frac{1}{\theta}\left(\frac{\beta\gamma}{1-\beta\gamma}\right)^2 \left(\frac{k}{1-\phi}\right)^2 \tag{7}$$

which is increasing in k as this is the seigniorage component of the staking reward. The money supply is declining in $(1-\phi)$ as this is the share of the total fee that is burnt. The steady state money supply is declining in the steady state staking share, θ. Intuitively, a high staking share is consistent with a relatively low amount of ETH being used, implying a relatively low ETH supply overall.

Based on this expression, we can compute a back-of-the-envelop estimate of the implied long-run money supply. The validator reward factor k can be determined based on the specifications of the Ethereum protocol. Based on my analysis in Section 2, the daily value for this factor is about 0.43, and the annual counterpart is simply multiplied by 365.25, giving a value of 160.6 (based on the unrounded values). The share of fee income burnt $(1-\phi)$ has been on average about 0.8 since August 2021. β is the discount factor, with rate of return $1/\beta - 1$, and γ the long run growth rate of the network's productivity/adoption level. The ratio $\left(\frac{\beta\gamma}{1-\beta\gamma}\right)$

is the payout valuation multiplier. For the S&P500, since 1950, the earnings multiplier has been on average about 20 (http://www.econ.yale.edu/~shiller/data.htm). There are many differences between these two, but this could be considered as an initial guess. As shown later in the paper, the estimated model is roughly consistent with this value, and the utility parameter θ, equivalently, the long run staking share, is estimated at 0.26. Taken together, this implies

$$M = \left(\frac{1}{0.26}\right)(20)^2\left(\frac{160.6}{0.8}\right)^2 = 62 \text{ million},$$

which is substantially below the current supply of about 120 million. Clearly, this estimate depends crucially on the parameter values about which there is considerable uncertainty. However, the example illustrates that one can reasonably expect the supply of ETH to continue declining as it has been since September 2022.[3] In the appendix, I compare Equation (7) to Elowsson (2021)'s model for the steady state money supply.

The total market capitalization of ETH in terms of the numeraire (assumed to be the dollar) is proportional to $\hat{p}M$. Because the price of ETH is growing at γ in steady state, the market capitalization grows at the same rate in steady state. As implied by the proposition, the market capitalization does not depend on k nor $(1-\phi)$, the staking reward factor and the share of priority fees in total fees, respectively. That is because the ratio $(k/(1-\phi))^2$ scales up the quantities measured in units of ETH, M^S, M^U and f, but reduces the price in terms of the numeraire by the same proportion. In other words, these two coefficients have no real effects in the steady state. As shown below, the value of k is also irrelevant for equilibrium dynamics. Note, the price of ETH in the model is only determined up to a constant that determines the level of A.

3.2 Stochastic dynamics

The stochastic dynamic system can be reduced to four equations for four unknowns M^S, M^U, f, and $\hat{p}$. The dynamic system has two endogenous state variables $\left(M_t^U, M_t^S\right)$. The aggregate money supply M_t is not a sufficient state variable. It can be solved for based on its two components.

The dynamic system comprises the three first-order conditions (with the price level stationarized) and the money supply equation:

$$\theta\left[\left(\hat{p}_t M_t^U\right)^{1-\theta} G^\theta\right]^{1-\sigma} = \hat{p}_t f_t G \tag{8}$$

[3] As shown later in the paper, the unconditional expected value of the money supply in the estimated stochastic model which explicitly accounts for nonlinearities is not very different from this back-of-the-envelop estimate.

$$\hat{p}_t = E_t \beta \frac{\Lambda_{t+1}}{\Lambda_t} \gamma \exp(z_{t+1}) \left[(1-\theta) \left(\hat{p}_{t+1} M_{t+1}^U \right)^{(1-\theta)(1-\sigma)-1} G^{\theta(1-\sigma)} + 1 \right] \hat{p}_{t+1} \tag{9}$$

$$\hat{p}_t = E_t \beta \frac{\Lambda_{t+1}}{\Lambda_t} \gamma \exp(z_{t+1}) \left[\frac{k}{\sqrt{M_{t+1}^S}} + \frac{\phi_{t+1} f_{t+1} G}{M_{t+1}^S} + 1 \right] \hat{p}_{t+1} \tag{10}$$

$$M_{t+1}^S + M_{t+1}^U = M_t^S + M_t^U + k\sqrt{M_t^S} - (1-\phi_t) f_t G. \tag{11}$$

The aggregate money supply is obtained by summing the two components

$$M_t = M^S + M_t^U.$$

The dynamic system exhibits invariance to k as shown in the following proposition. The proof is straightforward; it is in the appendix for completeness.

Proposition 2 *For an economy with a given k and equilibrium values $\left(M_t^S, M_t^U, f_t, \hat{p}_t\right)$, and for any positive constant ξ, there is an economy with $k' = k\xi$ for which the equilibrium values are the scaled variables $\left(M_t^S \xi^2, M_t^U \xi^2, f_t \xi^2, \hat{p}_t/\xi^2\right)$.*

The model is driven by the exogenous shock processes for z_t and ϕ_t defined in equations (2) and (3). In addition, the variable component of the discount factor is specified as

$$\begin{aligned} \frac{\Lambda_{t+1}}{\Lambda_t} &= \exp\left(-r_{z,t} + \rho_{\Lambda r} \sigma_r \varepsilon_{r,t+1} - \rho_{\Lambda z} \sigma_z \varepsilon_{z,t+1}\right) \\ r_{z,t} &= \rho_r r_{z,t-1} + \sigma_r \varepsilon_{r,t} \end{aligned}$$

where $r_{z,t}$ represents the variable component of the log of the one-period interest rate and the parameters $\rho_{\Lambda r}$ and $\rho_{\Lambda z}$ index risk prices for interest rate risk and for productivity risk. The one-period interest rate or short rate r_t is defined through $1/\exp(r_t) = E_t \beta \Lambda_{t+1}/\Lambda_t$.

4 Quantitative analysis

Assigning values to the parameters for a quantitative analysis of the model is challenging for at least two reasons. First, there is only a short sample of relevant historical data available. The current policies on fees have been in place only since the implementation of EIP-1559 in August 2021. Second, unstaking ETH has been possible only since April, 2023. The model is designed without restrictions on staking. I first present an estimate for the utility parameter θ based on two first-order conditions from the model. These hold even when unstaking ETH is not possible. The full model is parameterized using data on ETH prices and fees since August 2021 and longer time series on the USD real term structure. This approach implicitly assumes that prices and fees were only moderately affected by the inability to withdraw staked ETH. As time goes by, this preliminary parameterization can be updated.

4.1 Limited information estimate for θ

Combining the first-order conditions for ETH used M_{t+1}^U and gas G_{t+1}, equations (8) and (9), with minimal algebra implies

$$1 = E_t \beta \frac{\Lambda_{t+1}}{\Lambda_t} \frac{p_{t+1}}{p_t} \left[\left(\frac{1-\theta}{\theta} \right) \frac{f_{t+1} G_{t+1}}{M_{t+1}^U} + 1 \right]. \tag{12}$$

This equation holds independently from any frictions on staking. Conveniently, the only unobservable element in addition to the share parameter θ is the discount factor $\beta \frac{\Lambda_{t+1}}{\Lambda_t}$. Introducing alternative return $R_{t,t+1}^j$ from outside the model, this becomes

$$0 = E_t \left\{ \frac{\Lambda_{t+1}}{\Lambda_t} \right\} \left\{ \frac{p_{t+1}}{p_t} \left[\left(\frac{1-\theta}{\theta} \right) \frac{f_{t+1} G_{t+1}}{M_{t+1}^U} + 1 \right] - R_{t,t+1}^j \right\}.$$

A reference return $R_{t,t+1}^j$ is required to be priced by users of ETH and should not generate unobservable utility. Lido's liquid ETH staking token stETH satisfies these criteria. In addition, the fluctuations to this return are primarily driven by ETH prices which reduces estimation error. The return can be written as $R_{t,t+1}^{Lid} = \frac{p_{t+1}^{Lid}}{p_t^{Lid}} (y_{t+1} + 1)$, where y_{t+1} is the staking yield and p_t^{Lid} the price of stETH. Staking income is paid out to holders in additional units of stETH. With this, the moment condition can be written as

$$0 = E_t \left\{ \frac{\Lambda_{t+1}}{\Lambda_t} \right\} \frac{p_{t+1}}{p_t} \left\{ \left[\left(\frac{1-\theta}{\theta} \right) \frac{f_{t+1} G_{t+1}}{M_{t+1}^U} + 1 \right] - \left(\frac{p_{t+1}^{Lid}}{p_t^{Lid}} / \frac{p_{t+1}}{p_t} \right) (y_{t+1} + 1) \right\}.$$

From this moment condition, I can estimate θ with data on returns from Lido in addition to data on prices, fees and money supply. Consistent with the model, I assume risk neutrality. Because the volatility of the return spread is relatively low, this should not have a first-order impact. Estimating an SDF would require assumptions on the pricing factors as well as additional reference returns to determine the risk prices.

Based on weekly data going back to the implementation of EIP 1559, August 5, 2021, $\hat{\theta} = 0.26$ with Newey-West standard errors of 0.0014 based on the delta method. To a first approximation $\hat{\theta}$ is determined by the ratio of the averages of $\frac{f_{t+1} G_{t+1}}{M_{t+1}^U}$ and $\left(\frac{p_{t+1}^{Lid}}{p_t^{Lid}} / \frac{p_{t+1}}{p_t} \right) (y_{t+1} + 1) - 1$, with "weighting" by ETH price changes, p_{t+1}/p_t, playing a secondary role. Indeed the averages in annualized terms equal 0.02 and 0.05, respectively. Solving $\left(1 - \hat{\theta}\right) / \hat{\theta} = 0.05/0.02$ implies $\hat{\theta} = 0.286$ which is close to the estimated 0.26.

4.2 Model parameterization

The majority of the parameters are assigned so that the model matches selected moments of the USD real term structure and of ETH fee and price data. A model period is a month. The model is solved by a third-order approximation implemented with Dynare.

Parameter values are reported in Table 1. The utility share parameter θ is set to 0.26 as just described. The trend growth rate γ is set to 1.0025, implying an annualized rate of 3%. Productivity/adoption is assumed to follow a random walk so that $\rho_z = 0$.

Four parameters specifying the SDF are set to match properties of the real term structure: coefficient β, interest rate volatility and persistence parameters σ_r and ρ_r, and interest rate risk price $\rho_{\Lambda r}$. Target moments are the unconditional mean, standard deviation and persistence of the 10-year real Treasury rate, and the unconditional mean of the one-month rate. This part of the SDF is akin to a one-factor term structure model. For convenient calibration, I set $\beta = \hat{\beta} \exp\left(-\frac{(\rho_{\Lambda z}\sigma_z)^2}{2} - \frac{(\rho_{\Lambda r}\sigma_r)^2}{2}\right)$. This makes the term structure independent of the productivity/adoption risk parameter σ_z, and the one-period interest rate is independent of the shock volatility parameters, σ_r and σ_z.

The remaining parameters are determined based on monthly data from etherscan.io for ETH prices, fee incomes, and priority fee shares going back to the implementation of EIP 1559, August 5, 2021. Daily data is averaged or summed (when appropriate) to a monthly frequency.[4]

The three parameters of the priority fee share shock process $\left(\phi_0, \sigma_\phi, \rho_\phi\right)$ are estimated from the time series of the empirical counterpart. The remaining three parameters $(\sigma, \sigma_z, \rho_{\Lambda r})$ jointly match three moments: the standard deviation of the log ETH price change, the standard deviation of the log fee income change, and the average of the fee income divided by the supply of ETH used (not staked). For each of these three moments, one of the three parameters is dominant for the identification. The productivity shock variance σ_z matters most for the price volatility. The utility curvature σ matters most for fee volatility. The productivity risk price $\rho_{\Lambda r}$ matters most for the fee income divided by the supply of ETH used as this is proportional to the marginal utility of using ETH, see Equation (12).

[4]For May 1st 2022, the implied reported value for the priority fee share ϕ_t is negative, which seems to indicate a problem with the data. I remove this value and replace it with the interpolation from the two adjacent days.

ParamNam	ParamVal
{'the' }	0.26
{'gam' }	1.0025
{'rho_z' }	0
{'bethat' }	0.999
{'sig_r' }	0.00156
{'rho_r' }	0.968
{'rho_Lamr'}	34.3
{'sig' }	9.84
{'sig_z' }	0.1698
{'rho_Lamz'}	0.6702
{'phi0' }	-1.6055
{'sig_phi' }	0.201
{'rho_phi' }	0.775

Table 1. Parameter values

Moments based on the simulated model and the selected data moments are displayed in Table 2. The empirical volatilities of the changes in the price of ETH and fee income are very high, with monthly standard deviations of 0.178 and 0.325, respectively. The model matches these moments with volatile productivity shocks and the curvature of the utility function.

VarName	StdMon	MeanAnn	BenchStd	BenchMean
{'dp' }	0.1733	0.0305	0.178	0
{'df' }	0.3177	0	0.325	0
{'fGMu'}	0.0007	0.0206	0.0011	0.019
{'phi' }	0.0688	0.2112	0.066	0.2
{'r' }	0.0062	0.0121	0	0.012
{'y1' }	0.0046	0.0188	0	0
{'y10' }	0.0016	0.0264	0.0017	0.027
{'rs' }	0.1733	0.0857	0	0
{'M' }	0.1144	49.802	0	26.07

Table 2. Model implications and calibration targets

4.3 Model implications

As shown in Table 2, the unconditional mean of the money supply is 49.8 mn ETH. Nonlinearities amplified by risk are an important determinant of this value. In particular, the expected log return to staking, r_s in Table 2, has an annualized mean of 8.57%. With the mean log risk free rate at 1.21%, the return to staking ETH has a risk premium of over 7% annually.

The steady state formula for the money supply, Equation (7), is a reasonable approximation for the expected value of the money supply in the stochastic model if the price/payout multiplier is appropriately risk-adjusted. In particular, replace $\beta\gamma$ in the steady state equation by $E_t \beta \frac{\Lambda_{t+1}}{\Lambda_t} \frac{p_{t+1}}{p_t}$ from the stochastic model. With this, the ratio of $E(M)$ from the simulated stochastic model and the steady state M with the risk adjustment is 0.987 (see Table 4). In the stochastic model, the productivity risk price is an important determinant of this adjustment. It is identified from the data by the average staking yield which also determines the marginal utility of using ETH.

Figure 3 displays the projected paths of the money supply and of staked ETH as a share in the total supply. The projection takes April 2023 as the starting point. The transition of the money supply from its current level of about 120 million to the long run level is expected to take a very long time. The staking share is expected to sharply increase in the short run from its 15% level, and then to transition to its long-run level over an extended period.

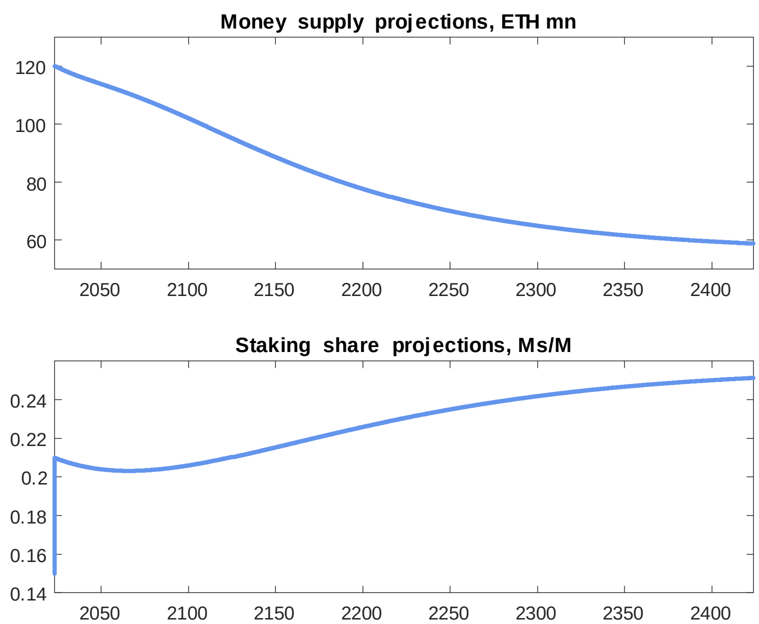

Figure 3: Projections

Figure 4 displays the responses of selected variables to the three shocks. The size of the shocks correspond to one standard deviation in the innovations.

Productivity shock, z. A positive shock in z permanently raises the productivity/adoption level A and the price of ETH by the same log-percent amounts than the shock. Neither money holdings nor the fee income are affected by the shock. This strong separability is due to absence of persistence in z, that is $\rho_z = 0$. Deviations from this case would induce transitory movements in the price as well as a transmission of this shock to the money holdings and fee income. However, the shock persistence is not likely to be very high, so these effects are unlikely to be significant.

Discount rate shock, r_z. The shock corresponds to an increase in the interest rate which drives down the price of ETH. An approximately 20 basis points increase in the interest rate at impact has a large effect on the price, a decline by more than 4%. This is because the interest rate is expected to remain low for several periods. With higher interest rates, staking ETH is less attractive and staked ETH, M^S, declines by a large amount. The marginal product of using ETH increases with the lower price, leading to an increase in ETH used, M^U. The fee income in the last row increases significantly. This is due to an increase in the marginal product of gas brought about by the lower price and the direct effect of a lower price of ETH on the cost of gas in terms of the numeraire; see the first-order condition for gas Equation (8).

The aggregate money supply M only changes by a very smaller amount at impact, however its response to the discount rate shock r_z is extremely persistent. The same applies for the priority fee shock in the third column. For these two shocks, even after 100 periods (months) the money supply is still not reverting back to its long-run level while the other variables have long started their mean reversion. In other words, the forces pushing the aggregate money to its long run level are very weak.

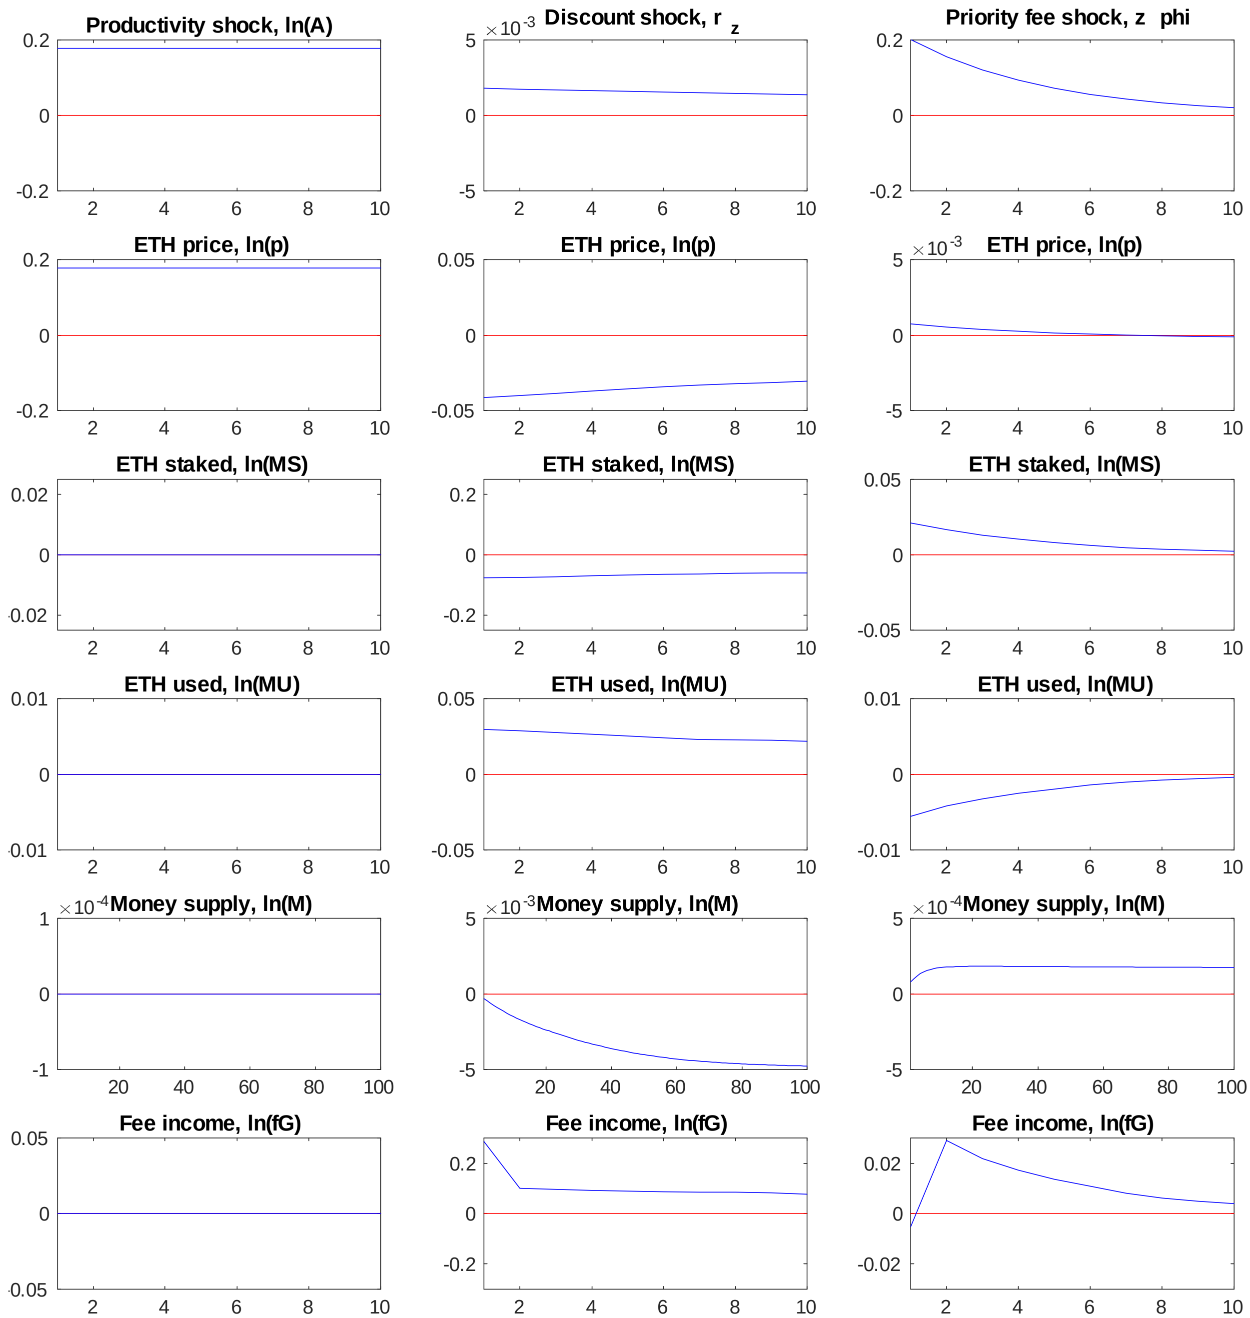

Figure 4. Impulse responses.

Table 3 displays the variance decomposition of selected variables to the three shocks. Consistent with the impulse responses, productivity shocks, in the first column, are the exclusive source of variance for price changes. Discount rate shocks (equivalently, interest rate shocks) in the second column account for the largest variance shares for fees and money holdings. The priority fee shocks primarily drive the priority fee share, but is not a significant driver of any of the endogenous variables in the model.

VarNamevd	VARDECz	VARDECrz	VARDECphiz
{'dp' }	95.86	4.10	0.00
{'df' }	0.00	99.39	1.09
{'phi'}	0.00	0.00	100.00
{'dMs'}	0.00	93.02	4.01
{'dMu'}	0.00	97.37	3.49
{'dM' }	0.00	99.33	1.22

Table 3. Variance decomposition to productivity/adoption shock, z, discount rate shock, rz (r_z), and priority fee share shock, phi (ϕ_z).

Priority fee shock, $z_{\phi,t}$. A positive shock increases the fee income for stakers leading to higher staking amounts. This shock has almost no impact on the price of ETH.

The findings on the variance decompositions are broadly consistent with Jermann (2021) where I estimate a basic monetary transaction model for Bitcoin and Ethereum and find that a majority of price fluctuations can be accounted for by shocks to adoption.

MomName	Bench	the=.4	gam=1.5%	sig=5	sig_z=.1	rho_Lamz=.6
"E(M)"	49.802	31.102	32.02	48.291	442.95	118.28
"E(rs)"	0.086	0.086	0.085	0.086	0.049	0.062
"Std(dp)"	0.173	0.174	0.173	0.174	0.109	0.174
"Std(df)"	0.318	0.286	0.297	0.164	0.389	0.355
"E(Ms/M)"	0.275	0.409	0.273	0.271	0.285	0.278
"Std(Ms/M)"	0.072	0.067	0.067	0.064	0.086	0.08
"E(M)/Mss*"	0.987	0.969	1.01	0.977	1.01	0.766

Table 4. Sensitivity analysis

Table 4 documents the sensitivities of model statistics to key parameters. For instance, if the standard deviation of the productivity/adoption shock σ_z were to drop from its current level of 17.8% to 10%, the money supply would increase dramatically. In this case, the reduction in risk would increase the demand for ETH by users and stakers.

5 Optimal monetary policy

This section considers alternative specifications for the staking yield and evaluates their impact on the expected utility created by the network as well as on the risk of the staking share dropping below a given level. First, the dependence on the aggregate amount staked is generalized beyond the square root function. Second, the staking reward coefficient k is made state-contingent.

The staking yield function can be generalized to

$$\frac{k}{\left(M_t^S\right)^x}, \text{ with total issuance } k\left(M_t^S\right)^{1-x},$$

for $x \geq 0$. The current policy for Ethereum corresponds to $x = 1/2$. A smaller x stabilizes the yield and is expected to make the staking share M^S/M more volatile. The limiting case for $x = 0$ with a constant staking yield is problematic because a very low (possibly zero) staking share becomes possible. A parameter value larger than $x = 1/2$ makes the yield more responsive and staking shares are expected to be less volatile. With $x = 1$, to amount of ETH issued per block is constant. According to Buterin (2022), a higher x would strengthen incentives to manipulate the validation process to discourage participation so as to collect a larger share of a relatively stable amount of ETH issued. Discouragement attacks are beyond the scope of my model. The model is well-suited to quantify the risk of a low staking share.

The following proposition characterizes the steady state for arbitrary $x \geq 0$.

Proposition 3 *The steady state money supply equals*

$$M = \frac{1}{\theta}\left(\left\{\frac{\beta\gamma}{1-\beta\gamma}\right\}\left\{\frac{k}{1-\phi}\right\}\right)^{1/x},$$

with the staking share M^S/M and market capitalizations $\hat{p}M^S$ and $\hat{p}M^U$ unaffected by x.

The proof is a straightforward generalization of the proof for Proposition 1.

To evaluate a policy, consider the expected lifetime utility of the network scaled by productivity $\hat{V}_t \equiv V_t/A_t$,

$$\hat{V}_t = \frac{1}{1-\sigma}\left[\left(\hat{p}_t M_t^U\right)^{1-\theta} G^\theta\right]^{1-\sigma} + (\beta\gamma)\, E_t \Lambda_{t+1} \exp\left(z_{t+1}\right) \hat{V}_{t+1}.$$

Given that the steady state market capitalization of unstaked ETH $\hat{p}M^U$ is unaffected by x, steady state lifetime utility is also unaffected by x. Beyond the steady state, the choice of x does impact lifetime utility. Specifically, the period utility is concave in $\hat{p}_t M_t^U$, which implies that fluctuations in market capitalization of unstaked ETH have a welfare cost. I measure utility differences in units of the compensated differential market capitalization Ω defined through $E\left[\hat{V}_t\right] = E\left[\hat{V}_t^{bench}\left(\left\{\hat{p}M^U\right\}(1+\Omega)\right)\right]$. For instance, a value of $\Omega = 0.02$ means that the improvement associated with a policy generating $\hat{V}_t$ relative to the benchmark $\hat{V}_t^{bench}$ is equivalent to a 2% permanent increase in the unstaked market capitalization. While this utility is the natural welfare measure in the model, the security represented by a high value of staked ETH is not valued by it. I can account for that by directly evaluating the risk of a low staking share.

As a second policy tool consider the staking reward coefficient k. Proposition 2 shows neutrality with respect to a constant k. As a consequence, the maximization for a policy maker with respect to k is not well-defined. However, one can consider a policy rule aiming at making deviations from steady state depend on the staking ratio

$$k_t = \bar{k}\exp\left(-\kappa\left(\ln\frac{M_t^S}{M_t} - \ln\theta\right)\right).$$

In this case, the policy can be optimized with respect to the sensitivity parameter κ. Presumably, a positive value for κ would have the potential to stabilize the staking share and affect expected lifetime utility.

Policy	UtilityGain	StkShareBelow15pc	StkShareStd
{'x=1/2' }	0	0.044	0.085
{'x=1' }	-0.0007	0.01	0.073
{'x=1/8' }	0.0012	0.095	0.096
{'x=1/2, kap=10'}	-0.0009	0	0.018

Table 5. Policy experiments.

Results from policy experiments are displayed in Table 5. The first row corresponds to the current specification with $x = 1/2$ with a constant k, that is $\kappa = 0$. Second and third rows correspond to $x = 1$ and $x = 1/8$, respectively. As expected, changing x to 1 is associated with a less volatile staking share. Specifically, the staking share drops below 15% only 1% of the time while for the benchmark this limit is breached 4.4% of the time. Moving the other way, that is $x = 1/8$, produces the opposite effect. More surprisingly, the lifetime utility for $x = 1$ is lower than the benchmark. The share of staked ETH, M^S/M, is more stable, but ETH used in units of the numeraire, pM^U, becomes more volatile, and this reduces utility. Both price and money supply contribute to this (not shown in the table).

In the last row, the sensitivity parameter for the staking share is set to $\kappa = 10$. As shown in the table, the staking share is stabilized but again this stabilization is associated with a loss in utility. One practical challenge is that the designer of the policy would need to know θ. Without that, the policy maker would embark on a futile attempt to stabilize the staking share around a level it cannot.

These quantitative experiments suggest that alternative and more activist 'monetary' policies are likely to have only small welfare effects and that even the sign of these may not be easily predictable. However, such policies can have first-order effects on the volatility of the staking share.

References

Biais, Bruno, Christophe Bisière, Matthieu Bouvard, Catherine Casamatta, 2019, The Blockchain Folk Theorem, The Review of Financial Studies, Volume 32, Issue 5, May 2019, Pages 1662–1715,

Bitmex Research, 2021, Ethereum's Proof of Stake System – Calculating Penalties & Rewards, https://blog.bitmex.com/ethereums-proof-of-stake-system-calculating-penalties-rewards/

Buterin Vitalik, 2022, Serenity Design Rationale, https://notes.ethereum.org/@vbuterin/serenity_design_rationale#Why-proof-of-stake

Buterin Vitalik, 2018, Discouragement Attacks, https://eips.ethereum.org/assets/eip-2982/ef-Discouragement-Attacks.pdf

Cong, Lin William, Zhiheng He and Ke Tang, 2022, Staking, Token Pricing, and Crypto Carry, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4059460

Cong, Lin William, Ye Li and Neng Wang, 2022, Token-based platform finance, Journal of Financial Economics, 144, 972–991.

Elowsson, Anders, 2021, Circulating Supply Equilibrium for Ethereum and Minimum Viable Issuance during the Proof-of-Stake Era, https://ethresear.ch/t/circulating-supply-equilibrium-for-ethereum-and-minimum-viable-issuance-during-the-proof-of-stake-era/10954?u=benjaminion

Fanti, G., L. Kogan, and P. Viswanath. 2019. Economics of Proof-of-Stake Payment Systems, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4320274

Jermann Urban, 2021, Cryptocurrencies and Cagan's model of hyperinflation, Journal of Macroeconomics, vol 69.

Jermann Urban and Haotian Xiang, 2022, Tokenomics: Optimal Monetary and Fee Policies, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4236859

Kim, Christine, 2021 The investment implications of Ethereum Improvement Proposal 1559, https://downloads.coindesk.com/research/EIP-1559-Ethereum-Fee-Market-Upgrade-Explained-1.pdf

Kose John, Thomas Rivera, and Fahad Saleh, 2021, Equilibrium Staking Levels in a Proof-of-Stake Blockchain, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3965599

Li, Ye and Simon Mayer, 2020, Money Creation in Decentralized Finance: A Dynamic Model of Stablecoin and Crypto Shadow Banking, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3757083

Mei, Kevin and Michael Sockin, 2022, The Risk of Adoption on Traded Token Platforms, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4115073

Migalabs, 2022, Validators or value-takers?, https://medium.com/blockchain-biz/validators-or-value-takers-813dd287fee8

Saleh, Fahad, 2021. Blockchain Without Waste: Proof-of-Stake. Review of Financial Studies 34(3):1156-1190.

Wahrstatter Toni, 2022, Practical Guide into Analyzing MEV in the Proof-of-Stake Era, https://medium.com/@toni_w/practical-guide-into-analyzing-mev-in-the-proof-of-stake-era-e2b024509918

6 Appendix

6.1 Proof of Proposition 1: Solution of steady state

For a steady state shock innovations are set to 0. The money supply equation at steady state, that is for a constant level of M, implies

$$f = \frac{k\sqrt{M^S}}{(1-\phi)\,G},$$

note that at steady state $\phi = \exp(\phi_0)$. The first-order condition for M^S imposing price growth of γ per period and after substituting f implies

$$M^S = \left(\frac{\beta\gamma}{(1-\beta\gamma)}\frac{k}{(1-\phi)}\right)^2 \tag{13}$$

so that the gas price equals

$$f = \frac{k\sqrt{M^S}}{(1-\phi)\,G} \rightarrow f = \frac{\beta\gamma}{(1-\beta\gamma)}\left(\frac{k}{1-\phi}\right)^2\frac{1}{G}. \tag{14}$$

The first-order condition for gas with scaled price $\hat{p} \equiv P/A$ is

$$\left[\left(\hat{p}M^U\right)^{1-\theta}G^\theta\right]^{1-\sigma} = \frac{G\hat{p}f}{\theta}.$$

The first-order condition for M^U

$$\left(\frac{1-\beta\gamma}{\beta\gamma}\right)\frac{\hat{p}M^U}{(1-\theta)} = \left[\left(\frac{\hat{p}M^U}{A}\right)^{1-\theta}G^\theta\right]^{1-\sigma}$$

Combining these last two equations, and substituting the solution for f gives

$$M^U = \left(\frac{\beta\gamma}{1-\beta\gamma}\right)^2\left(\frac{k}{1-\phi}\right)^2\left(\frac{1-\theta}{\theta}\right) \tag{15}$$

and use this in the first-order condition for M^U to get

$$\hat{p} = \left(\frac{1-\beta\gamma}{\beta\gamma}\right)^2\left(\frac{\theta}{1-\theta}\right)\left[(1-\theta)\,G^{\theta(1-\sigma)}\frac{\beta\gamma}{(1-\beta\gamma)}\right]^{1/[1-(1-\theta)(1-\sigma)]}\left(\frac{1-\phi}{k}\right)^2, \tag{16}$$

and grouping the exponents produces the expression in the main text. The solutions for M^S, f, M^U and $\hat{p}$ confirm the existence of the conjectured steady state.

6.2 Proof of Proposition 2: Invariance to k

Each equation can be rewritten unchanged for the scaled variables $\left(M_t^S\xi^2, M_t^U\xi^2, f_t\xi^2, \hat{p}_t/\xi^2\right)$ and for $k' = k\xi$.

$$\theta\left[\left(\frac{\hat{p}_t}{\xi^2}\xi^2 M_t^U\right)^{1-\theta} G^\theta\right]^{1-\sigma} = \frac{\hat{p}_t}{\xi^2}\xi^2 f_t G$$

$$\frac{\hat{p}_t}{\xi^2} = E_t\beta\frac{\Lambda_{t+1}}{\Lambda_t}\gamma\exp\left(z_{t+1}\right)\left[(1-\theta)\left(\frac{\hat{p}_{t+1}}{\xi^2}\xi^2 M_{t+1}^U\right)^{(1-\theta)(1-\sigma)-1} G^{\theta(1-\sigma)}+1\right]\frac{\hat{p}_{t+1}}{\xi^2}$$

$$\begin{aligned}
\frac{\hat{p}_t}{\xi^2} &= E_t\beta\frac{\Lambda_{t+1}}{\Lambda_t}\gamma\exp\left(z_{t+1}\right)\left[\frac{k}{\sqrt{M_{t+1}^S}}\frac{\xi}{\xi}+\frac{\phi_{t+1}f_{t+1}\xi^2 G}{M_{t+1}^S\xi^2}+1\right]\frac{\hat{p}_{t+1}}{\xi^2} \\
\frac{\hat{p}_t}{\xi^2} &= E_t\beta\frac{\Lambda_{t+1}}{\Lambda_t}\gamma\exp\left(z_{t+1}\right)\left[\frac{k\xi}{\sqrt{\xi^2 M_{t+1}^S}}+\frac{\phi_{t+1}f_{t+1}\xi^2 G}{M_{t+1}^S\xi^2}+1\right]\frac{\hat{p}_{t+1}}{\xi^2}
\end{aligned}$$

$$\begin{aligned}
\left(M_{t+1}^S\xi^2+M_{t+1}^U\xi^2\right)-\left(M_t^S\xi^2+M_t^U\xi^2\right) &= \xi^2 k\sqrt{M_t^S}-\left(1-\phi_t\right)Gf_t\xi^2 \\
\left(M_{t+1}^S\xi^2+M_{t+1}^U\xi^2\right)-\left(M_t^S\xi^2+M_t^U\xi^2\right) &= \xi k\sqrt{\xi^2 M_t^S}-\left(1-\phi_t\right)Gf_t\xi^2.
\end{aligned}$$

6.3 Comparison to Elowsson's (2021) "Supply equilibrium"

Elowsson (2021) proposed a "supply equilibrium" model, which is a steady state of the money supply equation obtained under some assumptions. This is a popular tool for long run projections, as seen for instance at https://ultrasound.money/. I am comparing here this model to the steady state of my model.

Based on the money supply equation, a steady state with a fixed money supply implies that the issuance for staking rewards, $k\sqrt{M^S}$, equals the burnt fees B

$$k\sqrt{M^S} = B.$$

This model assumes an exogenous staking reward yield defined as

$$y = \frac{k}{\sqrt{M^S}}.$$

With this yield, the amount of ETH staked is determined. The amount burnt is assumed to be given as

$$B = b\left(M - M^S\right),$$

with an exogenous burn rate b.

Under these assumptions, the money supply equals

$$M = \frac{k}{b}\sqrt{M^S} + M^S = \left[\frac{1}{b} + \frac{1}{y}\right]\frac{k^2}{y}.$$

For numerical predictions, values for y and b are needed. For instance, https://ultrasound.money/ for 1/27/2023, reports an issuance reward yield of 4.1% and a current burn rate of 1.9%, combined with my estimate for k, the long run money supply is estimated at[5]

$$M = \left[\frac{1}{0.019} + \frac{1}{0.041}\right]\frac{160.6^2}{0.041} = 48.5 \text{ million}.$$

A challenge for this model is that neither the issuance reward yield y nor the burn rate b can be seen as structural coefficients. Of course, in the long run, there may not be many structural coefficients for which we have good estimates.

For comparison, in my model, the steady state issuance reward yield and burn rate are given by

$$y = \frac{k}{\sqrt{M^S}} = (1-\phi)\left(\frac{1-\beta\gamma}{\beta\gamma}\right)$$

and

$$b = \frac{B}{M - M^S} = (1-\phi)\left(\frac{1-\beta\gamma}{\beta\gamma}\right)\left(\frac{\theta}{1-\theta}\right).$$

Interestingly, neither depends on the staking rewards factor k (effectively the base_reward_factor). The expression for y is intuitive. In equilibrium, the return to staking ETH has to equal the interest rate, $1/\beta - 1$. The return to staking depends not only on y but also on the growth rate of the price of ETH γ and the priority fee income and MEV captured by ϕ. Based on the parameter values used in the example in Subsection 3.1 of the main text, $y = 0.8/20 = 0.04$ and $b = y(.26/.74) = 0.014$. The values for y are close, the implied burn rates b are more distinct. This highlights again the importance of the implied staking share for estimates of the long run money supply.

[5] Note, https://ultrasound.money/ uses $k = 166.4$ based on Elowsson (2021), which is slightly different from mine. With this number, the estimate here would be 7.35% larger.