

# **Synth Subnet: Synthetic Price Path Data for Predictive Intelligence in LLMs and Agents**

Mode Network, Sam Hyatt, James Ross, Amber Shi

<https://www.mode.network/>

**Date:** February 2025

Synth (Subnet 50) redefines how synthetic financial data is modelled and utilized for predictive intelligence. Powered by Bittensor's decentralized intelligence network, Synth replaces the traditional, opaque data models with an open, verifiable, and continuously evolving synthetic data layer where the models compete, refine, and enhance monte-carlo simulations of price forecasts to unprecedented levels.

Historical financial data is constrained by biases, gaps, and a lack of exposure to extreme market conditions exposing models to overfitting and inaccuracies. Synth's approach solves these issues by building an open competition to reward the most accurate synthetic price distribution data. This data is purpose-built for agents and LLMs to reason, predict, and act with greater precision. As a self-learning, market-driven intelligence engine, Synth's data is set to fuel an economy, where AI can operate more effectively in uncertain and rapidly evolving new financial environments, such as crypto.

The whitepaper is presented as follows:

- Section 1. Introduction
- Section 2. Subnet Architecture
- Section 3. Design Choices
- Section 4. Results and Benchmarks
- Section 5. Roadmap and Future Developments
- Section 6. DTAO
- Section 7. Conclusion
- Section 8. Acknowledgements
- Section 9. About Synth Subnet
- Section 10. References

## 1. Introduction

### 1.1 Context and Purpose

Financial agents and LLMs are still in their infancy, with financial agents first emerging in Blockchains such as Mode Network less than a year ago. While evidence suggests high end-user demand and fledgling yet feasible technology, a crucial question remains: What is required to continuously enhance the performance of financial agents, enabling end users to access them at scale, while preserving crypto's open and permissionless blockchain ethos?

The answer lies in three fundamental principles: open-source, decentralized intelligence, and a synthetic data layer. These three elements combined lay a solid foundation for Agents to continue to thrive in finance. Open-source and decentralized intelligence are two key characteristics of Bittensor and its subnets. At the highest level, Bittensor is a framework for distributed teams to work in parallel to marry the computational power, with the computational problem of the creation of machine intelligence, in the form of digital commodities.<sup>1</sup> Bittensor has a network structure where machine intelligence production is permissionless and is rewarded based on the value of such intelligence. In Synth's case, it is based on the accuracy of BTC price distribution prediction (phase 1). The model outputs (price path predictions) are open source and can be used to refine future models.

By merging open-source datasets with incentive systems and a blockchain-based immutable ledger, Bittensor achieves scalable coordination of computational power. This enables the development of AI-based commodities in a way that prioritizes transparency and fairness.

### 1.2 Why Synth Subnet is Needed

Synth is designed to generate the best in class open source synthetic data to power AI Agent development, with its initial use case in DeFi Agents. In this section, we are going to discuss the two key components of Synth: synthetic data and Agents in DeFi versus traditional finance (banking, CEX trading)

#### Synthetic Data: The Foundation of Agent Intelligence

The effectiveness of AI agents relies heavily on the quality and diversity of data they consume. Machine learning and reinforcement learning models require context-rich datasets to deliver optimal outcomes. Synth addresses this need by providing high-fidelity synthetic data that replicates real-world market dynamics while overcoming the limitations of

real-world financial datasets.

Real-world data often lacks the breadth and depth required for robust AI training<sup>2</sup>. Key limitations include insufficient coverage of market extremes, noise, and incomplete labeling. Synth's synthetic data overcomes these challenges by:

- Encompassing a full range of market conditions, including volatility spikes and black swan events.
- Enabling large-scale, cost-effective data generation for advanced machine learning models.
- Mitigating overfitting and underfitting by providing diverse and balanced datasets.

By 2025, Gartner predicts synthetic data will reduce reliance on real-world datasets by 70%, becoming the primary input for AI models<sup>3</sup>. By 2030, synthetic data is expected to completely overshadow real data in AI models<sup>4</sup>. Synth is at the forefront of this shift, creating a foundation for data-driven innovation in DeFAI.

#### The Case for AI Agents in DeFi

Financial markets have evolved significantly, with algorithmic trading transforming traditional finance by leveraging large datasets and advanced models to optimize efficiency<sup>5</sup>. A similar transformation is underway in DeFi, but with two important distinctions. The case for Agents in DeFi versus traditional finance is grounded in these two unique properties of DeFi:

**Data Accessibility:** DeFi's decentralized and permissionless nature democratizes access to advanced financial strategies, offering a sharp contrast to the opacity of traditional financial markets, where high-frequency trading firms monopolize advanced tools and strategies. In DeFi, all the transaction, wallet, asset holding, and strategies can be easily extracted from an open ledger, enabled by the underlying blockchain. Users and innovators have the opportunity to experiment with AI agents and shape a more inclusive financial ecosystem.

**Complex Interactions:** DeFi workflows often require sophisticated strategies that are ideal for automation by agents. These strategies, which can seamlessly integrate into a single agentic workflow to optimize yield, include:

<sup>2</sup>

<https://www.technologyreview.com/2020/11/18/1012234/training-machine-learning-broken-real-world-health-nlp-computer-vision/>

<sup>3</sup>

<https://www.gartner.com/en/articles/you-ll-be-breaking-up-with-bad-customers-and-9-other-predictions-for-2022-and-beyond>

<sup>4</sup>

<https://aibusiness.com/data/gartner-can-synthetic-data-drive-the-future-of-ai->

<https://www.ft.com/content/62ea61e8-60f1-4aa1-9107-564ef4106dc1>

<sup>1</sup> <https://bittensor.com/whitepaper>

- Efficient bridging and swapping across chains.
- Looping strategies to utilize leverage.
- Optimizing LP rewards while mitigating impermanent loss.
- Maximizing yield in farming pools and lending protocols.

In short, these agents derive useful data from the ledger, automate complex trading strategies, optimize liquidity management, and address interconnected financial operations seamlessly, positioning DeFi as the perfect ecosystem for their adoption. Unlike traditional financial markets, DeFi's data accessibility and complexity enable a broader range of participants to leverage agent-driven innovation.

---

## 2. Subnet Architecture

### 2.1 High-Level Overview

At a high level, participants in the Synth subnet perform take parts in one or more of the following roles:

**Miners** generate probabilistic forecasts of cryptocurrency price movements, starting with BTC. Specifically, they must produce multiple simulated price paths for a given asset, starting from the current time and extending over specified time increments and horizons.

**Validators**, after the time horizon has passed, judge the accuracy of each miner's predicted paths compared to how the price moved in reality. The validator evaluates the miners' probabilistic forecasts using the Continuous Ranked Probability Score (CRPS). The CRPS is a proper scoring rule that measures the accuracy of probabilistic forecasts for continuous variables, considering both the calibration and sharpness of the predicted distribution. The lower the CRPS, the better the forecasted distribution predicted the observed value.

### 2.2 Key Roles

**Miners:** Initially, miners will generate 100 simulated paths for Bitcoin's future price at 5-minute intervals over the next 24 hours. Instead of predicting single-point future prices, the Synth Subnet emphasizes accurately quantifying uncertainty. Miners are expected to represent the probability distribution of future prices through their simulations, incorporating realistic price dynamics such as volatility clustering and skewed, fat-tailed distributions of price changes. Over time, this task will expand to include forecasts for multiple assets, necessitating accurate modeling of correlations between asset prices.

#### Initial Parameters:

- **Start Time** ( $t_0$ ): 1 minute after the request.
- **Asset:** Bitcoin (BTC).
- **Time Increment** ( $\Delta t$ ): 5 minutes.
- **Time Horizon** ( $T$ ): 24 hours.
- **Number of Simulations** ( $N_{sim}$ ): 100.

Miners must return ( $N_{sim}$ ) price paths by the start time, with each path containing price predictions at intervals determined by:

$$t_i = t_0 + i \times \Delta t, \text{ for } i = 0, 1, 2, \dots, N$$

Where  $N = \frac{T}{\Delta t}$  is the total number of increments.

**Validators:** as shown below

### 2.2.1 Application of CRPS to Ensemble Forecasts

In our setup, miners produce ensemble forecasts by generating a finite number of simulated price paths rather than providing an explicit continuous distribution. The CRPS can be calculated directly from these ensemble forecasts using an empirical formula suitable for finite samples.

For a single observation  $x$  and an ensemble forecast consisting of  $N$  members  $y_1, y_2, \dots, y_N$ , the CRPS is calculated as:

$$CRPS = \frac{1}{N} \sum_{n=1}^N |y_n - x| - \frac{1}{2N^2} \sum_{n=1}^N \sum_{m=1}^N |y_n - y_m|$$

Where:

The first term  $\frac{1}{N} \sum_{n=1}^N |y_n - x|$  measures the average absolute difference between the ensemble members and the observation  $x$ .

The second term  $\frac{1}{2N^2} \sum_{n=1}^N \sum_{m=1}^N |y_n - y_m|$  accounts for the spread within the ensemble, ensuring the score reflects the ensemble's uncertainty.

This formulation allows us to assess the miners' forecasts directly from their simulated paths without the need to construct an explicit probability distribution.

### 2.2.2 Application to Multiple Time Increments

To comprehensively assess the miners' forecasts, the CRPS is applied to sets of price changes over different time increments. These increments include short-term and long-term intervals (in the case of the initial checking prompt parameters, these will be 5 minutes, 30 minutes, 3 hours, 24 hours).

For each time increment:

**Predicted Price Changes:** The miners' ensemble forecasts are used to compute predicted price changes over the specified intervals

**Observed Price Changes:** The real asset prices are used to calculate the observed price changes over the same intervals. We recommend the validators collect and store the prices by sending requests to the Pyth oracle at each time increment, to be used at the end of the time horizon.

**CRPS Calculation:** The CRPS is calculated for each increment by comparing the ensemble of predicted price changes to the observed price change.

The final score for a miner for a single checking prompt is the sum of these CRPS values over all the time increments.

## 2.3 Calculation of Leaderboard Score

### 2.3.1 Normalization Using Softmax Function

After calculating the sum of the CRPS values, the validator normalizes these scores across all miners who submitted correctly formatted forecasts prior to the start time. The normalized score  $S_i$  for miner  $i$  is calculated as:

$$S_i = \frac{e^{-\beta \cdot CRPS_i}}{\sum_j e^{-\beta \cdot CRPS_j}}$$

where:

- $CRPS_i$  is the sum of  $CRPS$  values for miner  $i$  on that day
- $\beta = \frac{1}{1000}$  is the scaling factor
- The negative sign ensures better forecasts (lower  $CRPS$ ) receive higher scores

Any miners who did not submit a correct prediction are allocated a normalised score of 0 for that prompt.

### 2.3.2 Exponentially Decaying Time-Weighted Average (Leaderboard Score)

The validator is required to store the historic request scores for each miner. After each new request is scored, the validator recalculates the ‘leaderboard score’ for each miner, using an exponentially decaying time-weighted average over their past per request scores, up to a threshold of 30 days in the past. This approach emphasizes recent performance while still accounting for historical scores. The leaderboard score for miner  $i$  at time  $t$  is calculated as:

$$L_i(t) = \frac{\sum_j w_j S_{i,j}}{\sum_j w_j}$$

where:

- $S_{i,j}$  is the normalized score of miner  $i$  at request  $j$ .
- $w_j = e^{-\lambda(t-t_j)}$  is the weight assigned to the score  $S_{i,j}$ .
- $t_j$  is the time of request  $j$ .
- $\lambda = \frac{\ln 2}{h}$  is the decay constant, with half-life  $h = 10$  days.
- The sum runs over all requests  $j$  such that  $t - t_j \leq T$ , where  $T = 30$  days is the threshold time.

### 2.3.3 Allocation of Emissions

At the end of each day, the leaderboard scores are then raised to the power of an exponent  $\alpha$  (e.g.,  $\alpha = 2$ ) to amplify performance differences. The adjusted scores determine each miner's share of the total emissions for that day

Adjusted Scores:

$$AdjScore_{i,t} = (L_{i,t})^\alpha$$

Emission Allocation:

$$P_{i,t} = \frac{AdjScore_{i,t}}{\sum_j AdjScore_{j,t}} \times Total$$

## 2.4 Overall Purpose

The system creates a competitive environment through:

1. Implementing CRPS scoring objectively measures forecast quality across multiple time increments.
  2. Using ensemble forecasts to calculate CRPS from a finite ensemble of simulations.
  3. Applying CRPS to different time increments to evaluate both short-term and long-term predictions.
  4. Normalizing scores ensures fair comparison using softmax function.
  5. Calculating leaderboard scores and allocating emissions rewards consistent performance and encourages competition.
-

### 3. Design Choices

#### 3.1 Key Decisions and Rationale

We aim to minimize restrictions on miners, allowing them maximum flexibility to innovate and explore the model and best approach to generate synthetic data. At the same time, we have implemented several design choices to ensure alignment between Synth's objectives and the outputs generated by its participants. These choices are guided by three key questions:

- How can individual contributions be effectively recognized and rewarded?
- What mechanisms can drive sustained improvement?
- How can strategies that exploit the system be deterred?

#### 3.2 Incentive Structures

##### 3.2.1 Open competition

Synth is committed to onboarding top modelling teams from fields like crypto, finance, and science—including options traders, high-frequency trading groups, AI agent developers, and signal processors—to launch miners.

We aim to foster a community of ideas sharing collaboration between miners wherever possible. We endeavour to facilitate the onboarding of new miners, especially those new to Bittensor, with clear instructions and blog posts on how to build benchmark models.

##### 3.2.2 Reward Distribution

Synth aims to reward innovation and accuracy. The Beta parameter in our SoftMax scoring function allows us to control how much objective differences in miner performance translates to differences in scores. While the Alpha parameter in our weights calculation allows us to control the steepness of the rewards distribution across miners. We can also adapt the drop-off rate in the moving average, to ensure we're sufficiently filtering out luck and noise, while not putting off new miners. Our team of data scientists will adjust these parameters as the competition progresses, to ensure the highest quality miners are heavily rewarded and incentivized, while allowing new miners to quickly climb the leaderboard if they're providing innovation.

We'll aim to have around 60% of the rewards going to the top 4 miners, and most of the remainder going to the next 10, with a token amount going to the other miners. However the exact nature of the distribution depends on the comparative scores.

#### 3.2.3 Ensuring Continuous Improvement

The Synth team will continue to sequentially release miners with improved benchmark models, aiming to demonstrate that the best miners are pushing further ahead of each benchmark over time. Once these benchmarks have been surpassed, they can be made open-source to provide new miners with a stronger starting point, allowing them to focus on more innovative, cutting-edge advancements. The benchmarks will incorporate the latest techniques from the field of financial time series modelling, such as GARCH and latent stochastic volatility models.

The rewards distribution is normalized, making it highly dependent on the number of miners and the quality of their outputs. On other subnets, this creates challenges in tracking objective miner improvements over time. However, our subnet's scoring mechanism incorporates the "Sum of CRPS scores" which is calculated independently for each miner. This allows for tracking objective performance (or comparative performance of the best miner against the benchmarks) as the contest progresses. Synth will release a tool to visualize the progression of these scores.

The increasing complexity of the challenges—from single assets on fixed timeframes to multiple assets on variable timeframes, and eventually to correlated assets—will ensure continuous innovation. Miners will need to invest sustained effort and research to remain at the top of the leaderboard.

### 3.3 Limitations, Risks, and Exploits

#### 3.3.1 Model/Output Copying

Synth aims to allow validators to make predicted paths available to external AI agents and human users shortly after they're submitted by miners. However, doing so introduces the risk of duplication, as miners may attempt to replicate and resubmit results generated by others.

The Softmax function combined with the moving average in our scoring process has the benefit of dissipating rewards from similar or equivalent miners and amplifying the rewards to unique and innovative miners. This should act as a deterrent to copying miners.

The team is researching methods to further mitigate this risk of copying:

- Recommendations to validators to only publicly release a subset of paths, blended across multiple miners, therefore preventing full copying of the best miner

- Additional checks of equality or long-term correlation between miner predictions, with corresponding deductions in scores

- Innovative ‘watermarking’ algorithms for embedding miners’ ids or keys into their predictions. If a suitable method can be found, this could have much wider users across the whole generative AI industry

- Potential partnering with other subnets to perform ‘proof of inference’, to demonstrate that miner’s predictions did arise from their own generative models

```
bt.logging.warning("You are allowing non-registered entities to send requests to your miner. This is a security risk.")
```

---

### 3.3.2 Relay Mining

In theory, a miner could operate a validator node on Synth and exploit this dual role by relaying predictions received by the validator. Instead of performing the required computation, the miner could simply obtain the output from another miner and submit it as their own in response to the original prompt. However, this direct copying would only impact the scoring of that one nefarious validator, who would then be punished in the Yuma Consensus process.

This validator could attempt to repurpose stolen predictions for use by its miner in response to subsequent prompts from other validators. However, the preventative measures in 3.3.1 still apply. Plus, from ‘Phase 2’ onward, there will be considerable variability in the requirements for each prompt, severely impacting the ability of miners to repurpose copied predictions.

### 3.3.3 Validator Impersonation

On a subnet, it is theoretically possible for anyone to issue prompts to miners, bypassing validation and scoring mechanisms. Such unauthorized prompts can result in miners producing data points without receiving proper validation or compensation, undermining the fairness and efficiency of the subnet. To address this, Synth implements robust authentication protocols to prevent validator impersonation and ensure a secure and equitable reward system.

The following code snippet demonstrates how Synth’s authentication mechanism is enforced within the subnet:

```
If not self.config.blacklist.force_validator_permit:
```

```
bt.logging.warning("You are allowing non-validators to send requests to your miner. This is a security risk.")
```

```
if self.config.blacklist.allow_non_registered:
```

## **4. Results and Benchmarks**

Upon the mainnet launch of the subnet, we will continue to update this section of the whitepaper as more quantifiable results become available.

### **4.1 Early Performance Metrics**

On testnet new miners have already managed to significantly outperform the initial ‘Geometric Brownian Motion’ benchmark model. Initially this benchmark miner was earning the majority of the rewards but has since dropped outside the top 5 in the leaderboard, courtesy of new superior miners joining on testnet.

The Synth team will continue to run this GBM miner and look forward to seeing its share of the rewards continue to decay on mainnet as more advanced miners join the competition.

### **4.2 Comparisons with Existing Systems**

The competitive environment, coupled with collaboration within the miner community, is expected to drive significant advancements over any existing system.

The Synth team will continue to monitor cutting edge research, and attempt to contact any authors of new relevant papers, and encourage them to launch miners on our subnet.

### **4.3 Case Studies or Success Stories**

Synth is collaborating closely with AI agent teams to leverage its synthetic data for training AI agents capable of:

- Optimizing liquidity positions.
- Mitigating collateral liquidation risks.
- Managing portfolios with improved risk-adjusted returns.

These teams are excited to have access to synthetic data that encapsulates the complexity and nuances of real-world price data. They find this resource significantly more advantageous than attempting to generate synthetic data in-house or relying on the limitations of historical datasets.

Additionally, Synth plans to evaluate the collective intelligence of its subnet by developing an automated options trading agent. Performance metrics for this agent, as well as opportunities to participate in its associated vault, will be shared as they become available.

---



## 5. Roadmap and Future Developments

### 5.1 Phase 1 Synth MVP

During testnet and initial mainnet launch, Miners will construct models that generate 100 predicted paths of BTC price data for the next 24 hours, with 5-minute increments.

#### Features:

- Live prediction visualizer: Real-time graphical representation of predicted BTC price distributions.
- Historic synthetic data API: Easy access to past synthetic price data for analysis and model training.

#### Uses for Agents:

- Access large-scale, accurate synthetic BTC price data for offline training.
- Request BTC price paths for the next 24 hours to guide automated transaction decisions.

### 5.2 Phase 2 Synth Bespoke

During Phase 2, Miners will create versatile models capable of generating price paths for multiple crypto assets over variable timeframes.

#### Features:

- Bespoke prediction requests API: Generate synthetic data tailored to specific needs.
- Natural language interface: Powered by a custom LLM for real-time interaction and personalized predictions.

#### Uses for Agents:

- Generate synthetic data customized to specific tasks, facilitating offline training for diverse scenarios.
- Query live predictions of price dynamics across multiple tokens to inform automated trading strategies, retaining sole ownership of predictions.

### 5.3 Phase 3 Synth Portfolio

Miners will develop models to predict correlated price paths, incorporating the prices of multiple tokens per time step.

#### Features:

- Multidimensional path viewer: Visualize the evolution of multiple token prices over time.
- Correlation matrices and heat maps: Analyze relationships and interdependencies between assets.

#### Uses for Agents:

- Train portfolio management AI agents with synthetic data that accurately represents cross-token correlations and volatility interactions.
- Access hyper-realistic synthetic data ideal for optimizing diversification and risk management strategies.

### 5.4 Phase 4 Synth Universe

Miners will expand their models to generate synthetic data for industries beyond DeFi, including commodities, weather, traffic, sports, lifespan, and relationship outcomes.

#### Features:

- Versatile LLM interface: Interact with Monte Carlo data across multiple domains.
- Broad user applications: Build sophisticated tools leveraging Synth path data for various use cases.

#### Uses for Agents:

- Enhance LLM comprehension of complex, multi-domain dynamics.
  - Enable AI agents to strategize effectively in uncertain environments and with incomplete information.
  - Transform decision-making in fields like health, travel, and logistics with actionable predictive insights.
-

## 6. The Move to DTAO

The move to dTAO presents an opportunity to demonstrate the true value of the models being built on the Synth subnet and the data they produce.

embedding Synth-generated insights directly into onchain financial mechanisms, Synth alpha token enables a sustainable, incentive-aligned ecosystem where the value of alpha tokens is directly linked to the predictive power of Synth's models.

---

### 6.1 Commercial Strategy: Monetizing Synthetic Intelligence

The features and integrations we are building on top of Synth will allow us to develop new revenue models and drive value to alpha token holders. The intention for all of the features and use cases on our roadmap will be to provide a basic free version, alongside a pro tier that will require an amount of staked Synth-dTAO to access.

For substantial usage, such as acquiring vast amounts of historical data or high-frequency live predictions for AI agent decision-making, payments in Synth-DTAO alpha token will be required. Through our partnership with the Mode AI Builders community, we are actively developing new features on top of Synth, with potential profit-sharing mechanisms for the DTAO these features generate.

Our data science team will introduce a VIP tier for the Synth mining community, accessible through a Synth-DTAO alpha token staking requirement. This VIP program will include:

- Exclusive insights into leading miner performance and strategy analysis
- In-depth research posts on new benchmarks and modeling techniques

### 6.2 Utility of the Alpha Token: Expanding DeFi Applications

The utility of Synth-DTAO extends beyond access to premium data services—it is being actively integrated into DeFi applications, ensuring its role as a core economic unit in AI-driven finance. Upcoming integrations include:

- Bridging and EVM Representation – Expanding Synth's presence across multiple chains, enabling cross-chain token utility and interoperability.
- Lending and Borrowing Markets – Using Synth-generated probabilistic models to power risk-adjusted lending and collateral management.
- Automated Market Maker (AMM) Liquidity Pools – Establishing Synth-dTAO trading pairs with key DeFi assets to enhance liquidity and market participation.
- Dollar-Cost Averaging (DCA) Capabilities – Providing agent-driven execution of automated trade strategies, leveraging Synth's predictive intelligence.

Ultimately, Synth-DTAO will become the payment layer for knowledge and real-world understanding among AI agents. By

## 7. Conclusion

This whitepaper demonstrates that Synth can develop synthetic probabilistic models on Bittensor that surpass centralized approaches. By combining an open competitive environment with blockchain-based incentives, Synth creates a transparent and scalable framework for advancing AI agent model training and achieving mass adoption in crypto.

With its phased roadmap, Synth aims to expand from BTC price predictions to multi-asset and multi-domain applications. Its use of CRPS scoring and dynamic incentives ensures fairness and continuous innovation, empowering miners and validators to produce high-fidelity synthetic data.

Synth has already shown potential in optimizing liquidity, mitigating risks, and enhancing portfolio management for AI-driven agents. By democratizing access to high-quality synthetic data, Synth accelerates the development of advanced AI agents while fostering collaboration and innovation.

As Synth evolves, it is poised to become a cornerstone of the agentic economy, bridging the gap between traditional systems and the next generation of autonomous, AI-driven solutions in decentralized finance.

### 7.1 Call to Action for Collaboration or Adoption

We invite developers, miners, validators, and innovators to join Synth in advancing decentralized intelligence. Whether contributing to model development, ensuring data integrity, or leveraging high-fidelity synthetic data, your participation can drive innovation in DeFi and AI. Together, we can shape the future of AI-driven solutions and the agentic economy.

---

## 8. Acknowledgements

The authors of this paper extend their gratitude to the Bittensor community for their support in creating this subnet, assisting with incentive design, and providing valuable insights on technical implementation.

We also thank our advisors from General Tao Ventures and Routable 21 for their prompt guidance and strategic direction when needed.

Lastly, we acknowledge and appreciate the contributions of the Synth subnet team: Mode, James Joss, Sam Hyatt, Stanislav Parkhomenko, and Amber Shi.

---

## 9. About Synth Subnet

### 9.1 Mission and Vision

Synth redefines how markets approach uncertainty, replacing vague estimations with quantified probabilities. By bridging the gap between data scarcity and agent intelligence, Synth empowers DeFi participants and AI developers to make informed, data-backed decisions with confidence.

Synth's ultimate goal is to become the foundational data layer for the agentic economy. Unlike traditional price prediction systems that offer single-point forecasts, Synth captures full price distributions, delivering nuanced insights for AI agents. The vision of Synth is to start with single asset price paths, move to multiple asset price paths, then integrate correlated asset paths and finally expand to generate Monte Carlo simulated data for other industries. E.g. commodities, weather, traffic, sports, lifespan, relationship outcomes.

### 9.2 Organizational Background

Synth is a flagship project developed by Mode. Since its launch in January 2024, Mode has rapidly emerged as a leading EVM

Layer 2 chain, achieving \$500 million in Total Value Locked (TVL) within its first six months and supporting over 70 deployed DeFi protocols. Mode has also launched an AI Agent App Store, featuring a range of user-facing DeFi agents designed to enhance accessibility and usability. Through its accelerator program, Mode has provided support to over 20 early-stage DeFi and AI agent teams, fostering innovation and growth in the agentic and DeFi economy. These initiatives reflect Mode's strategic focus on advancing the agentic economy on-chain and its commitment to establishing robust infrastructure for AI-driven financial systems. The mainnet launch of Synth in January 2025 marks a pivotal step in realizing this vision, providing a comprehensive data layer for AI agents and contributing to the broader integration of DeFi and AI technologies.

---

## 10. References

<https://bittensor.com/whitepaper>

<https://taostats.io/>

<https://www.mode.network/>

<https://github.com/mode-network/synth-subnet>