Can we imitate stock price behavior to reinforcement learn option price?

Xin Jin

Bank of Montreal

Abstract—This paper presents a framework of imitating the price behavior of the underlying stock for reinforcement learning option price. We use accessible features of the equities pricing data to construct a non-deterministic Markov decision process for modeling stock price behavior driven by principal investor’s decision making. However, low signal-to-noise ratio and instability that appear immanent in equity markets pose challenges to determine the state transition (price change) after executing an action (principal investor’s decision) as well as decide an action based on current state (spot price). In order to conquer these challenges, we resort to a Bayesian deep neural network for computing the predictive distribution of the state transition led by an action. Additionally, instead of exploring a state-action relationship to formulate a policy, we seek for an episode based visible-hidden state-action relationship to probabilistically imitate principal investor’s successive decision making. Our algorithm then maps imitative principal investor’s decisions to simulated stock price paths by a Bayesian deep neural network. Eventually the optimal option price is reinforcement learned through maximizing the cumulative risk-adjusted return of a dynamically hedged portfolio over simulated price paths of the underlying.

Index Terms—Behavioral finance, option pricing, dynamic hedging, Bayesian deep neural network, visible-hidden Markov network, reinforcement learning

I. INTRODUCTION

Modern quantitative finance employs the key idea of replicating portfolio to bind financial derivative pricing and hedging to cross-sectional variations of tradable assets’ returns. Therefore, effective modeling the underlying return has become in demand and pivotal. Traditional models attempt to accurately mirror the asset price movements as an analytical stochastic process or a combination of multiple random processes, such as the geometric Brownian motion in the Black–Scholes model [1], the compound Poisson process in the Merton jump diffusion model [2], and the mean reverting square-root process combining two Wiener processes in Heston model [3]. However, framing the evolution of asset prices into stochastic processes is either too basic to cope with complexity or it relies on numerous undetermined parameters to be calibrated frequently.

Deep Neural Networks (DNNs) [6], [7], [8], [9] have emerged as a powerful computational toolkit in view of the following excellent features. The Universal Approximation Theorem [5] exhibits that a DNN can approximate any highly nonlinear multivariate function arbitrarily well. DNNs are data-driven function approximators without presumed mathematical form, thereby relieving heavy parameter calibration.

Moreover, DNNs are capable of approximating random functions with the help of Bayesian approaches [10], [11], [20].

Limitations of traditional models motivate us to develop an innovational way that characterizes asset price movements by behavioral analysis with the aid of DNNs. This paper focuses on equity options and accordingly underlying assets refer only to stocks. We choose Markov Decision Process (MDP) [4] to associate the investor’s motive/behavior with stock price changes. At each time step, MDP is in a state (spot price of the underlying stock in our scenario) and the decision maker (Principal Investor (PI)) chooses an action (PI’s decision) based on the current state. PI is defined as the collection of investors who have a leading influence on stock market and PI’s decision indicates the quantity of the capital that PI chooses to invest in or sell out a particular underlying stock. We approximately quantize PI’s decision with the accessible features of the equities pricing data. Our MDP is thought to be non-deterministic as it randomly moves into a new state with a probability that is influenced by the chosen action.

Bayesian Deep Neural Network (B-DNN) [10], [11], [12], [13] introduces Bayesian statistics to DNN weights inference. Against a point estimate of DNN weights, B-DNN learns a probability distribution over the network weights. We utilize B-DNN to detect the probability distribution over price change induced by PI’s decision through a two-step process (1. Bayesian inference, 2. regression analysis). In the first step, we infer the posterior distribution over B-DNN weights. Whereafter, we perform a regression analysis to determine the relationship between price change and PI’s decision. This is done by modeling the distribution over price change as a posterior predictive distribution which depends on the posterior distribution obtained in the first step. Owing to the merit of DNN, the second step does away with a closed-form expression for non-linear multi-dimensional feature mapping. Instead, the regression analysis is built on a linear combination of basis vectors with few parameters to be pre-set. We implement Bayesian inference and regression analysis with real stock market data to evaluate the model’s performance. The numeric results reveal that PI’s decision and the stock price change are approximately positive correlated. This proves that our approach can practically pinpoint the impact. We apply B-DNN in a cutting-edge way to behavioral finance research. Some recent excellent applications of machine learning in finance focus on using neural networks to perfect traditional quantitative models based on characterizing random processes, such as calibrating multivariate Lévy processes [14] and calibrating Heston Model [15].

We apply leader-follower type Keynesian beauty contest

X. Jin is with Bank of Montreal, Canada. (e-mail: Xin.Jin@bmo.com; felixxinjin@gmail.com).

[25], [26] to explain the approximately positive correlation between the stock price change and PI's decision: a single retail investor easily believe that the instantaneous observation of PI's decision is the most investors' concurrent decision to follow. Leader-follower type Keynesian beauty contest discloses the behavior pattern of a single retail investor but not PI's strategy to make decision that is also the last remaining piece of our MDP. In contrast to conventional MDP that associates a single type of factor (i.e., state) to the action selection by a policy function, we believe that the action selection depends on two types of factors (i.e., visible state and hidden state) and build a visible-hidden Markov network to model the dependency of PI's decision making on visible factor and hidden factor. An algorithm is proposed to update the visible-hidden Markov network for best possibly fitting the visible state sequence and the observation sequence. Through the instrumentality of a trained visible-hidden Markov network for imitating PI's behavior from observations, we simulate the paths of PI's successive decision making and further transform them to stock price paths through Bayesian inference and regression analysis. Imitation by means of our visible-hidden Markov network is proved experimentally a robust approach to effectively process the input signal with low Signal-to-Noise Ratio (SNR), so it is particularly suitable for the financial fields. Imitation learning has become a hot topic in robotics research. However, some imitation learning algorithms in robotics, such as maximum entropy based Inverse Optimal Control (IOC) [16], [17], [18] and imitation learning from state observation without action information [19], require demonstration samples with high SNR and are not easy to be applied to financial scenarios.

The final step in reaching our ultimate goal is designing an optimization criterion for pricing the option and a dynamic hedging algorithm that can efficaciously take advantage of the cross-sectional information yielded by simulated price paths of the underlying stock. To this end, a self-financing portfolio is built to dynamically replicate the option price. The self-financing restriction raises a risk of exhausting cash for rebalancing the portfolio during the life of the option owing to a volatile price of the underlying stock. With the intention of minimizing this risk, the option price is comprised of a risk-free price and a risk-adjustment. The self-financing restriction also guarantees that the portfolio's value change between adjacent time steps comes from the portfolio's return. Thus, we are capable of exploiting reinforcement learning to maximize the cumulative risk-adjusted return for recursively learning the hedging position and pricing the option.

II. EFFECT OF THE PRINCIPAL INVESTOR'S DECISION

As described in the introduction section, we are motivated to model the evolution of the underlying stock price by virtue of DNN and behavioral analysis, rather than analytically characterize stochastic processes in traditional way. With this end in view, a non-deterministic MDP is designed to associate stock price evolution with PI's successive decision making. This section is on a quest for the impact of PI's decision on stock price change.

We define PI as the aggregation of investors who are considered as the dominating force in the equity market. PI could be composed of large institutional investors or a sizable group of retail investors acting in concert. PI's decision is concerned with the amount of money flowing into or out of a stock within a time slot. We approximately quantize PI's decision as a signed number $d \in \Omega_d$ in terms of the opening price $o \in \Omega_o$, the highest price h , the lowest price l , the closing price $c \in \Omega_c$, and the volume v :

$$d = \left(v h \frac{h - o}{h - l} \right) \text{sign}(c - o). \quad (1)$$

This representation allows the difference between the opening and closing prices to indicate the direction of the money flow. Moreover, the difference between the opening and highest prices becomes a primary determinant of the money flow volume. Through simple math deduction, this synthetic variable is proved to have a monetary unit, and we consider it as USD in this paper.

The sample spaces Ω_d and Ω_o (or Ω_c) constitute respectively the MDP action space and state space. The price change

$$g = \frac{c}{o} - 1 \quad (2)$$

functions as an indicator of the MDP state transition. We presume that the price change is incompletely impacted by PI's decision and partly random. This implies that the price change follows a probability distribution which is dependent on PI's decision. We intend to use a data-driven regression method to determine this probability distribution and verify our presumption. For this aim, we resort to B-DNN [10], [11], [12], [13] which applies Bayesian statistics to DNN weights inference. By further merging B-DNN with Gaussian process [20], [21], the price change distribution is simplified as a gaussian distribution.

Given a dataset $\mathcal{D} := \{(d_n, g_n)\}_{n=0}^{T-1}$ which contains pairs of quantized PI's decision and the price change sampled from T consecutive time slots, we apply Stochastic Gradient Descent (SGD) coupled with the dataset $\mathcal{D}$ to train a DNN. The DNN takes an input data d_n and returns a predicted value $f_{\theta}(d_n) \in \mathbb{R}$ to compare with the observed value c_n for producing the empirical loss, where $\theta = [\theta_1, \dots, \theta_M]^T$ signifies the DNN weights. The loss is defined in terms of the loss function $l(g_n, f_{\theta}(d_n))$ and the $L2$ regularizer $r(\theta)$ with a regularization parameter λ :

$$\begin{aligned} \mathcal{L}(\theta, \mathcal{D}) &= L(\theta, \mathcal{D}) + r(\theta) \\ &= \sum_{n=0}^{T-1} l(g_n, f_{\theta}(d_n)) + \frac{\lambda}{2} \theta^T \theta. \end{aligned} \quad (3)$$

In contrast to fit a point estimate of DNN weights, B-DNN learns a distribution over the weights. We take advantage of this merit of B-DNN to infer the posterior distribution over the weights (Bayesian inference) and therefrom derive a posterior predictive distribution which measures the distribution over price change induced by PI's decision (Regression analysis).

In order to implement Bayesian inference on the weights, the loss is straightforwardly embedded in the posterior distribution over the weights via the exponential family form

$P(\theta | \mathcal{D}) := \frac{e^{-\mathcal{L}(\theta, \mathcal{D})}}{\int e^{-\mathcal{L}(\theta, \mathcal{D})} d\theta}$. It implies both the likelihood and the prior also belong to the exponential family with the expression as $P(\mathcal{D} | \theta) := \frac{e^{-L(\theta, \mathcal{D})}}{\int e^{-L(\theta, \mathcal{D})} d\theta}$ and $P(\theta) := \frac{e^{-r(\theta, \mathcal{D})}}{\int e^{-r(\theta, \mathcal{D})} d\theta}$. In such a manner, the Maximum A Posteriori (MAP) estimation $\hat{\theta} = \arg \max_{\theta} P(\theta | \mathcal{D})$ is equivalent to minimize mean square error. However, the immediate MAP estimation is often computationally impractical.

Laplace approximation [22], [23], [24] provides a computationally feasible Bayesian inference on weights. The main idea is to approximate likelihood times prior with a Gaussian density function in respect that likelihood times prior is a well-behaved uni-modal function in $\mathcal{L}^2$ space. The posterior distribution over the weights is consequently derived to be the below Gaussian distribution (see derivation in Appendix A):

$$P(\theta | \mathcal{D}) \approx \mathcal{N}(\theta | \theta^*, [\nabla_{\theta\theta}^2 L(\theta^*, \mathcal{D}) + \lambda \mathbf{I}_M]^{-1}), \quad (4)$$

where $\theta^* = [\theta_1^*, \dots, \theta_M^*]^T$ is the mode of $P(\mathcal{D} | \theta) P(\theta)$.

After obtaining the posterior over the weights, we then seek for the posterior predictive distribution over price change $P(g^* | d^*, \mathcal{D})$, where $d^* \in \mathbb{R}$ denotes a test sample of quantized PI's decision and g^* signifies the corresponding observation of the price change. The high dimensional projection and the nonlinear relationship between d^* and $f_{\theta^*}(d^*)$ impede to directly find $P(g^* | d^*, \mathcal{D})$ in closed form. In order to overcome this impediment, a new random variable y which has a readily accessible posterior predictive distribution is constructed and can be linearly expressed in terms of g^* . As a result, $P(g^* | d^*, \mathcal{D})$ is inferred from the given posterior predictive distribution $P(y | d^*, \mathcal{D})$. As presented in [20], y is possible to be designed as a linear combination of a Gaussian process and additive gaussian noise. Different from [20], we calculate the parameters of the Gaussian process (i.e., $m(d^*)$ in Equation (8) and $k(d^*, d^*)$ in Equation (9)) based on $P(\theta | \mathcal{D})$ instead of $P(\theta | \mathcal{D}')$ where $\mathcal{D}' := \{(y_n, c_n)\}_{n=0}^{T-1}$ is a transformed dataset and y_n is a sample of the constructed variable. In addition, we consider the residual (ε in Equation (7)) as additive independent identically distributed (i.i.d) Gaussian noise rather than additionally bring in the second-order partial derivatives of the loss to represent noise.

The random variable y is designed as the following combination of the standard linear regression model and the first order derivative of the loss function:

$$y = \phi(d^*)^T \theta + \nabla_{f_{\theta^*}(d^*)} l(g^*, f_{\theta^*}(d^*)), \quad (5)$$

where DNN weights follow a prior distribution $\theta \sim \mathcal{N}(\mathbf{0}, \mathbf{K}_{\theta\theta})$ with the covariance matrix $\mathbf{K}_{\theta\theta} = \lambda^{-1} \mathbf{I}_M$. And $\phi(d^*)$ is a Jacobian matrix whose entries form a set of basis functions to project the input sample into M dimensional feature space:

$$\begin{aligned} \phi(d^*) &= \nabla_{\theta} f_{\theta}(d^*) |_{\theta=\theta^*} \\ &= \begin{bmatrix} \frac{\partial f_{\theta}(d^*)}{\partial \theta_1} |_{\theta=\theta_1^*} \\ \vdots \\ \frac{\partial f_{\theta}(d^*)}{\partial \theta_M} |_{\theta_M=\theta_M^*} \end{bmatrix}, \end{aligned} \quad (6)$$

Figure 1: GME behavior observed in each 5-minute time slot from 09:35 to 16:00 during 6 consecutive trading days. Earlier time slots are with lighter shades.

As the linear regression model is built on the feature space in Equation (6), it successfully leaves the difficulty of processing non-linear high-dimensional mapping back to DNN.

For a quadratic loss function $l(g^*, f_{\theta^*}(d^*)) = \frac{1}{2}(g^* - f_{\theta^*}(d^*))^2$, the constructed random variable y can be simplified in terms of g^* from Equation (5):

$$\begin{aligned} y &= \phi(d^*)^T \theta + (f_{\theta^*}(d^*) - g^*) \\ &= \varphi(d^*) + \varepsilon, \end{aligned} \quad (7)$$

where $f_{\theta^*}(d^*) - g^*$ is the residual that can be modeled as additive independent identically distributed Gaussian noise $\varepsilon \sim \mathcal{N}(0, \sigma^2)$. Gaussian process is an effective tool for describing a distribution over functions [21]. From the function-space view, the process $\varphi(d^*)$ is specified by a Gaussian process $\varphi(d^*) \sim \mathcal{GP}(m(d^*), k(d^*, d^*))$, where we calculate the mean function $m(d^*)$ and the covariance function $k(d^*, d^*)$ based on the posterior distribution over the weights $P(\theta | \mathcal{D})$ obtained in Equation (4):

$$\begin{aligned} m(d^*) &= \phi(d^*)^T \mathbb{E}[\theta | \mathcal{D}] \\ &= \phi(d^*)^T \theta^*, \end{aligned} \quad (8)$$

$$\begin{aligned} k(d^*, d^*) &= \phi(d^*)^T \mathbb{E}[\theta\theta^T | \mathcal{D}] \phi(d^*) \\ &= \phi(d^*)^T [\nabla_{\theta\theta}^2 L(\theta^*, \mathcal{D}) + \lambda \mathbf{I}_M]^{-1} \phi(d^*). \end{aligned} \quad (9)$$

Equation (7) exhibits that the random process y is a linear combination of a Gaussian process and additive Gaussian noise. Therefore, the posterior predictive distribution of y can be easily deduced from the statistical description of the Gaussian process and Gaussian noise as:

$$P(y | d^*, \mathcal{D}) = \mathcal{N}(\phi(d^*)^T \theta^*, k(d^*, d^*) + \sigma^2). \quad (10)$$

Since we always prefer to perform a prediction based on the DNN weights at the mode of the posterior θ^* rather than any possible DNN weights, only the determinate realization $\phi(d^*)^T \theta^*$ should be extracted from the Gaussian process instead of the entire Gaussian process $\phi(d^*)^T \theta$ for calculating g^* . Thuswise g^* is a linear combination of a random variable and a constant:

$$\begin{aligned} g^* &= y - \phi(d^*)^T \theta^* + f_{\theta^*}(d^*) \\ &= y + C. \end{aligned} \quad (11)$$

Consequently, the posterior predictive distribution over price change is derived from Equation (10) and (11):

$$P(g^* | d^*, \mathcal{D}) = \mathcal{N}(f_{\theta^*}(d^*), k(d^*, d^*) + \sigma^2). \quad (12)$$

The observation of the price change g^* shares the same posterior predictive variance with y owing to the fact that the predicted value $\phi(d^*)^T \theta^*$ is a determinate realization produced at the mode of the posterior θ^* and the deduction of the predicted value from y is incapable of eliminating the immanent uncertainty in the inference on weights and the residual. When we consider an input vector of test samples $\mathbf{d}^* = [d_1^*, \dots, d_N^*]^T$, Equation (12) turns into the below form:

$$P(g^* | \mathbf{d}^*, \mathcal{D}) = \mathcal{N}(f_{\theta^*}(\mathbf{d}^*), \text{diag}(k(\mathbf{d}^*, \mathbf{d}^*) + \sigma^2 \mathbf{I}_N)). \quad (13)$$

During Jan. 26, 2021-Feb. 02, 2021, GameStop Corporation (GME) stock price swiftly rose and then fell with the unusually high price and volatility because a short squeeze was triggered by organized retail investors. We would like to explore the behavioral finance behind this event through our model. The intraday data for the 6 consecutive trading days is exercised to infer B-DNN coefficients and be referred to produce testing samples for validating regression analysis. Due to the unreachability, the pre/post-market data is not included even though some huge price fluctuations occurred in the pre/post-market sessions. We structure a dataset $\mathcal{D}$ as 77 samples in each of 6 episodes to accommodate the equity pricing data in the

Figure 2: Regression with $\sigma = 2.5$ and $\lambda = 0.7$

period from 09:35 to 16:00 for the 6 consecutive trading days. Each sample is a five-dimensional vector which enumerates open-high-low-close prices and volume observed in a 5-minute time slot. We discard the pricing data in the first time slot (i.e., 09:30 to 09:35) because the open price is subject to the pre-market trading and gaps up or down from the previous day's close. Moreover, the volume in the first time slot is extremely high and influenced by the pre-market news. Based on Equation (1), (2) and dataset $\mathcal{D}$, we yield the training dataset $\mathcal{D}_i := \{(d_n, g_n)\}_{n=0}^{T_i-1}$ which contains the pairs of quantized PI's decision and price change sampled from T_i consecutive time slots in the i th episode as illustrated in Figure 1. The samples are plotted in chronological order. The lighter shade indicates that the sample is taken in an earlier time slot.

We first train B-DNN respectively on an episode basis with the training dataset $\mathcal{D}_{i,i=1\dots 6}$ to determine the posterior distribution over the weights and then work out the posterior predictive distribution over price change led by testing samples of quantized PI's decision. The testing samples are produced to be evenly spaced over the interval between the largest and smallest training sample of quantized PI's decision in each episode. There are two fine-tuned hyperparameters σ and λ for optimizing our model's performance. Since σ has a significant influence on posterior predictive variance as shown in Equation (12), the optimal value of σ can be found out by the "68-95-99.7 rule". We carry out Bayesian inference and regression analysis with two sets of parameters ($\sigma = 2.5$, $\lambda = 0.7$, and $\sigma = 1$, $\lambda = 0.1$) and exhibits the results in turn by Figure 2 and Figure 3. It is observed that $\sigma = 2.5$ offers a fitting posterior predictive variance to satisfy in principle the

Figure 3: Regression with $\sigma = 1$, $\lambda = 0.1$

"68–95–99.7 rule", while $\sigma = 1$ leads a too small posterior predictive variance to match the actual data distribution. The experimental evaluation also demonstrate that the overfitting is effectively prevented by λ . As shown in Figure 2, the overfitting is fully suppressed by a relatively large regularization parameter $\lambda = 0.7$. On the other hand, overfitting emerges when the regularization parameter decreases to a smaller value, for instance, $\lambda = 0.1$ as illustrated in Figure 3. Furthermore, we can observe from the numeric results that the stock price change has an approximately positive correlation with quantized PI's decision.

III. INVERSE PI'S DECISION FROM VISIBLE AND HIDDEN STATES

The previous section presents a DNN-based approach to determine the state transition (i.e., price change) after executing an action (i.e., principal investor's decision). We still wonder how an action is made and especially how associate an action with the current state (i.e., open price). Answering this query equals the proffer of the last remaining piece of our MDP. This section will find out the underlying law to simulate PI's successive decision making.

When we ground the numerical results on the real equities pricing data in the previous section, it is observed that higher quantized PI's decision approximately leads to higher stock price change. This phenomenon can be theoretically explained by the Keynesian beauty contest [25], [26]. Keynesian beauty contest describes a beauty contest where voters win a prize for voting the most popular faces among all voters. Under such rule of the contest, a voter tends to guess the other voters'

Figure 4: Visible-hidden Markov network

choices and match the choice of the majority based on the own guess. Keynes believed that the similar behavior pattern applies to the investors within the stock market. Most investors attempt to make the same decision that the majority do rather than build a decision on their own evaluation of fundamental profitability. However, it is very hard to guess other investor's concurrent decision for a retail investor. For this reason, the stock market evolves into a leader-follower type Keynesian beauty contest where a single retail investor (i.e., follower) tends to follow the instantaneous observation of PI's decision and thinks it as the most investors' concurrent decision. The leader-follower type Keynesian beauty contest explains well the approximately positive correlation between PI's decision and the stock price change. Nevertheless, PI's decision making is still inadequate to be inferred from the leader-follower type Keynesian beauty contest.

We believe that two types of factors influence PI's decision making: visible and hidden. The visible factor is observable, while the hidden factor is hard to be quantized or detected, for example, the impact of news on investors' anticipation of the investment return and thus on their investment behavior.

The conventional MDP applies a policy function to map state space to action space. However, the policy function is not sufficient for mapping two state spaces to action space. In order to identify the dependency of PI's decision making on visible factor and hidden factor, we build a visible-hidden Markov network as shown in Figure 4. The hidden factor is modeled by a random variable (i.e., hidden state) which changes through time and the visible factor is modeled by another random variable (i.e., visible state). The hidden state transits over time slots, and its transition is a Markov process. In the same time slot, the hidden state, the visible state, and the random variable representing PI's decision (i.e., observation) are pairwise correlated. We design a chain for connecting these random variables to enable the calculation of the dependencies among them. This chain design differentiates the visible-hidden Markov network from hidden Markov models [27] which considers only two coinstantaneous components without the visible state. In each time slot, the hidden state is directly connected to the visible state and then indirectly chained to the observation via the visible state.

We specify the notations used in the visible-hidden Markov network below:

- $\zeta = \{\zeta_0, \zeta_1, \dots, \zeta_{J-1}\}$: hidden state sample space,
- $\nu = \{\nu_0, \nu_1, \dots, \nu_{K-1}\}$: visible state sample space,
- $\omega = \{\omega_0, \omega_1, \dots, \omega_{L-1}\}$: observation sample space,
- $Z = \{z_0, z_1, \dots, z_{J-1}\}$: initial hidden state distribution,

$\mathbf{A} = \{a_{ij}\} \in \mathbb{R}^{J \times J}$: hidden state transition matrix,
 $a_{ij} = P\{\text{hidden state } \zeta_j \text{ at } t+1 \mid \text{hidden state } \zeta_i \text{ at } t\},$
 $\mathbf{B} = \{b_i(k)\} \in \mathbb{R}^{J \times K}$: visible state probability matrix,
 $b_i(k) = P\{\text{visible state } v_k \text{ at } t \mid \text{hidden state } \zeta_i \text{ at } t\},$
 $\mathbf{C} = \{c_k(l)\} \in \mathbb{R}^{K \times L}$: observation probability matrix,
 $c_k(l) = P\{\text{observation } o_l \text{ at } t \mid \text{visible state } v_k \text{ at } t\},$
 $H = \{h_0, h_1, \dots, h_{T-1}\}$: hidden state sequence,
 $S = \{s_0, s_1, \dots, s_{T-1}\}$: visible state sequence,
 $O = \{o_0, o_1, \dots, o_{T-1}\}$: observation sequence,
 $\Gamma = \{\mathbf{A}, \mathbf{B}, \mathbf{C}, Z\}$: Visible-hidden Markov network.

In our scenario, the hidden state indicates bullish or bearish news received just before a time slot, so the hidden state sample space ζ is with a size of 2. The visible state and the observation chained to the same hidden state represent respectively the opening price and quantized PI's decision in the same time slot influenced by the just acquired market information. The visible state sample space ν and the observation sample space ω are finite and discrete compared with the continuous sample space of the opening price and quantized PI's decision. For this reason, we borrow the concept of the quantization in digital signal processing to map values from the sample space of the opening price and quantized PI's decision to ν and ω :

$$\bar{s} = \left\lfloor \frac{s}{s_{max} - s_{min}} \right\rfloor, \quad (14)$$

where $\lfloor \cdot \rfloor$ is the floor function. s stands for a sample value of the opening price or quantized PI's decision. s_{max} and s_{min} mark respectively the upper and lower limits of the sample space of the opening price or quantized PI's decision. Q denotes the size of ν or ω which is K or L . $\bar{s}$ is mapped from s to constitute ν or ω .

The hidden state is initialized from a distribution Z and transits based on a transition probability matrix $\mathbf{A}$. In the same time slot, the visible state is related to the hidden state by a probability matrix $\mathbf{B}$ and the observation is associated with the visible state by another probability matrix $\mathbf{C}$ to form a chain. Our objective is to find $\Gamma = \{\mathbf{A}, \mathbf{B}, \mathbf{C}, Z\}$ which best possible fits the visible state sequence S and the observation sequence O . To this end, an algorithm is proposed to modulate Γ by iteratively increasing $P(S, O \mid \Gamma)$:

Step 1: Initialize Γ and set $P(S, O \mid \Gamma) = 0$.

Step 2: Calculate the conditional joint probability function $\alpha_t(i)$, $\beta_t(i)$, $\gamma_t(i, j)$ and the conditional probability function $\gamma_t(i)$:

$$\begin{aligned} \alpha_{t=0}(i) &= P(s_0, o_0, h_0 = \zeta_i \mid \Gamma) \\ &= z_i b_i(s_0) c_{s_0}(o_0), \end{aligned} \quad (15)$$

$$\begin{aligned} \alpha_{t>0}(i) &= P(s_0, s_1, \dots, s_t, o_0, o_1, \dots, o_t, h_t = \zeta_i \mid \Gamma) \\ &= \sum_{j=0}^{K-1} \alpha_{t-1}(j) a_{ji} b_i(s_t) c_{s_t}(o_t), \end{aligned} \quad (16)$$

$$\beta_{t=T-1}(i) = P(s_{T-1}, o_{T-1} \mid h_{T-1} = \zeta_i, \Gamma) = 1, \quad (17)$$

Figure 5: Simulation of PI's successive decision making

Figure 6: Simulation of GME price movements

$$\begin{aligned} & \beta_{t < T-1}(i) \\ &= P(s_{t+1}, s_{t+2}, \dots, s_{T-1}, o_{t+1}, o_{t+2}, \dots, o_{T-1} \mid h_t = \zeta_i, \Gamma) \\ &= \sum_{j=0}^{K-1} a_{ij} b_j(s_{t+1}) c_{s_{t+1}}(o_{t+1}) \beta_{t+1}(j), \end{aligned} \quad (18)$$

$$\begin{aligned} \gamma_{t=0 \dots T-2}(i, j) &= P(h_t = \zeta_i, h_{t+1} = \zeta_j \mid S, O, \Gamma) \\ &= \frac{\alpha_t(i) a_{ij} b_j(s_{t+1}) c_{s_{t+1}}(o_{t+1}) \beta_{t+1}(j)}{P(S, O \mid \Gamma)} \\ &= \frac{\alpha_t(i) a_{ij} b_j(s_{t+1}) c_{s_{t+1}}(o_{t+1}) \beta_{t+1}(j)}{\sum_{i=0}^{J-1} \alpha_{t=T-1}(i)}, \end{aligned} \quad (19)$$

$$\begin{aligned} \gamma_{t=T-1}(i) &= P(h_t = \zeta_i \mid S, O, \Gamma) \\ &= \frac{\alpha_{t=0}(i)}{\sum_{i=0}^{J-1} \alpha_{t=0}(i)}, \end{aligned} \quad (20)$$

$$\begin{aligned} \gamma_{t=0 \dots T-2}(i) &= P(h_t = \zeta_i \mid S, O, \Gamma) \\ &= \sum_{j=0}^{J-1} \gamma_t(i, j). \end{aligned} \quad (21)$$

Step 3: Re-estimate Γ :

$$z_i = \gamma_{t=0}(i), \quad (22)$$

$$a_{ij} = \frac{\sum_{t=0}^{T-2} \gamma_t(i, j)}{\sum_{t=0}^{T-2} \gamma_t(i)}, \quad (23)$$

$$b_i(k) = \frac{\sum_{t \in \{0,1,\dots,T-1\}, s_t=\nu_k} \gamma_t(i)}{\sum_{t=0}^{T-1} \gamma_t(i)}, \quad (24)$$

$$c_k(l) = \frac{\sum_{t \in \{0,1,\dots,T-1\}, s_t=\nu_k, o_t=\omega_l} \gamma_t(i)}{\sum_{t=0}^{T-1} \gamma_t(i)}. \quad (25)$$

Step 4: If $P(S, O | \Gamma) = \sum_{i=0}^{J-1} \alpha_{t=T-1}(i)$ climbs, go to Step 2, otherwise end the loop.

The parameters of a visible-hidden Markov network is initialized before entering the iterative process and then updated in each iteration until $P(S, O | \Gamma)$ doesn't grow anymore. $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{Z}$ are row stochastic, moreover, their rows respectively belong to a standard $J-1$, $J-1$, $K-1$, and $L-1$ simplex which are the support of the Dirichlet distribution of order $J-1$, $J-1$, $K-1$, and $L-1$. Therefore, we can employ the Dirichlet distribution to randomize $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{Z}$ to keep away from a local maximum which prevents the evolution of the model.

Figure 1c and 1d clearly show that PI's decision making strategy varies from one episode to another: PI preferred to give a decision with large absolute value at the early stage on Jan. 28, 2021, on the other hand, tended to spread large trades over time on Jan. 29, 2021. In light of this signs, our determination of a visible-hidden Markov network is episode-based. We measure Γ in each episode to generate time series of the probable evolution of PI's decision making and further transform them to stock price paths through Bayesian inference and regression analysis.

We imitate PI's successive decision making based on observations from the training dataset $\mathcal{D}_{i,i=1\dots 6}$. The hyperparameters are given the following values: $J = 2$, and $K = L = 30$. We parametrize Dirichlet distribution with $\alpha_i = 1000$, so that the initial $\mathbf{A}$, $\mathbf{B}$, $\mathbf{C}$, and $\mathbf{Z}$ are row stochastic and have near uniform values, where α_i denotes any component of the parameter vector α . Figure 5 exhibits three simulated paths of PI's successive decision making along with the real evolution of PI's decision. In view of article space, we only show the experimental results in two representative episodes for comparison in figures. According to Figure 5a and 5b, we find that the simulated paths on Jan. 29, 2021 are more volatile than those on Jan. 28, 2021. This correctly reflects the discrepancy of the volatility of real PI's decision on these two days and thus well proves the effectiveness of our algorithm. The B-DNN is trained by the training dataset $\mathcal{D}_{i,i=1\dots 6}$ when we perform Bayesian inference in the section II. We continue to carry out regression analysis grounded on this trained B-DNN for each evolutionary path of PI's decision. For the sake of simplicity, the price change is simulated as its posterior predictive mean. Figure 6 illustrates our simulated price movements of the underlying stock together with the real price path.

IV. REINFORCEMENT LEARN OPTION PRICE GROUND UPON IMITATIVE PI'S CHRONOLOGICAL DECISION MAKING

Our algorithm introduced in the section III probabilistically imitates PI's chronological decision making. With the help

of B-DNN inference-regression presented in the section II, likely evolutionary paths of PI's decision are transformed into probable price paths of the underlying stock. These simulated price paths of the underlying stock yield the cross-sectional information for optimizing a dynamically hedged portfolio. This section expounds how to reinforcement learn option price using the cross-sectional information stemmed from imitative PI's chronological decision making.

Suppose that we sell a European option with the terminal payoff $H(S_T)$, where S_T denotes the underlying stock price at expiry date T . For the sake of the suitability of our algorithm for both calls and puts, switching the below expression of the terminal payoff will turn pricing/hedging a call option into a put option:

$$H(S_T) = \begin{cases} \max(S_T - K, 0), & \text{if calls} \\ \max(K - S_T, 0), & \text{if puts.} \end{cases} \quad (26)$$

With the proceeds from the sale of the option, we build a portfolio Π_t which consists of a_t units of the stock with price S_t and a deposit in the risk-free bank account B_t . With the aim of dynamically hedging the option, we rebalance this portfolio to keep its value replicating the value of the option at time $t \leq T$:

$$\Pi_t = a_t S_t + B_t. \quad (27)$$

At maturity we close out the position as we no longer need hedge the option. Consequently, the terminal payoff is the portfolio value at maturity:

$$\Pi_T = H(S_T) = B_T \quad (28)$$

We ignore transaction costs or other frictions in this paper due to the fact that low rates or fixed commissions are offered to the customers with high volume of trades executed (e.g., High-Frequency Trading (HFT) firms). We force our dynamically hedged portfolio Π_t to be self-financing in such a manner that there is no external infusion or withdrawal of cash over the lifetime of the option, and the spread of portfolio values at different times comes from the continuously compounded risk-free interest rate r and the change in underlying stock price ΔS_t :

$$\Pi_{t+1} - e^r \Pi_t = a_t \Delta S_t + \kappa f_{\theta^*}(F) a_t S_t, \quad (29)$$

where $\Delta S_t = S_{t+1} - e^r S_t$ and $F = (a_{t+1} - a_t) \Theta S_t$.

The cross-sectional information $\mathcal{F}(t) = \{S_t^1, \dots, S_t^U\}$ collects possible prices at time t from all the simulated price paths, where U is the number of simulated paths and $S_t^u, u=1, \dots, U$ denotes the price at time t sampled from the u th path. Our cross-sectional information is derived from the existing PI's decision, so we design an compensation term $\kappa f_{\theta^*}(F) a_t S_t$ to put in the price change caused by our hedging activity if we are also referred to as PI. $f_{\theta^*}(F)$ maps a large-scale volume of buying or selling the underlying stock to a sharp rise or drop in price via our trained B-DNN. We may sell multiple options covering significant amount of shares of the underlying stock. We use Θ to denote the total number of shares of the underlying stock. The coefficient κ

gives a degree of freedom to scale the compensation term. If we are a single retail investor who sells a mini option, then this compensation term vanishes as κ tends to 0. Monte Carlo sample-based reinforcement learning [28] uses the cross-sectional information which is generated by sampling a pre-set probability distribution. Compared with Monte Carlo sample-based reinforcement learning, our algorithm can better fit into each circumstance as we adopt the cross-sectional information that reflects the existing PI's decision and add an compensation term to compensate the effect of our hedging activity on the price change.

When price an option at any time within its lifetime, the risk of running out cash to continue rebalancing the portfolio due to the fluctuation in the market price of the underlying stock should be taken into account. To this end, the ask price C_t at time $t \leq T$ is designed to be the sum of the risk free price and the risk-adjustment weighted by the risk-aversion coefficient η [28] as expressed in Equation (30). The risk free price is the expected value of the portfolio Π_t . The risk-adjustment accumulates the risk from time t until the expiry date T . The risk at time t is quantized as the expected discounted variance of the portfolio Π_t .

$$C_t(S_t) = \mathbb{E}_t \left[\Pi_t + \eta \sum_{t'=t}^T e^{-r(t'-t)} \text{Var} [\Pi_{t'} | \mathcal{F}_{t'}] | \mathcal{F}_t \right]. \quad (30)$$

Geometric Brownian motion is the most commonly used model of stock price behavior. By applying Itô's lemma to the random variable of stock price which follows Geometric Brownian motion, the stock price at time t is proved to be lognormally distributed [29]:

$$\ln S_t \sim \mathcal{N} \left(\ln S_0 + \left(\mu - \frac{\sigma_s^2}{2} \right) t, \sigma_s^2 t \right), \quad (31)$$

where μ is the expected rate of return per time-slot from the stock, and σ_s is the volatility of the stock price.

Equation (31) implies S_t is not martingale due to the constant drift rate $\mu - \frac{\sigma_s^2}{2}$. When price the option, we prefer the calculation to be based on an independent variable without constant drift. Therefore, we expand the independent variable S_t in Equation (27), (29), and (30) to a function in terms of an independent variable without constant drift $\tilde{S}_t$:

$$S_t = \exp \left(\tilde{S}_t + \left(\mu - \frac{\sigma_s^2}{2} \right) t \right). \quad (32)$$

S_t afterwards will signify the expansion of stock price in terms of $\tilde{S}_t$ in all equations.

The following notations will denote for short the other variables which are dependent on $\tilde{S}_t$ in subsequent derivations:

$$\tilde{a}_t = \tilde{a}_t(\tilde{S}_t) = a_t, \quad (33)$$

$$\pi = \pi(\tilde{S}_t, t) = \tilde{a}_t, \quad (34)$$

$$\Delta S_t = e^{\tilde{S}_{t+1} + \left(\mu - \frac{\sigma_s^2}{2} \right) (t+1)} - e^r e^{\tilde{S}_t + \left(\mu - \frac{\sigma_s^2}{2} \right) t}, \quad (35)$$

$$\Pi_t = \Pi_t(\tilde{S}_t), \quad (36)$$

where π is the policy function which maps $\tilde{S}_t$ and t to the position in the stock at time t .

We are aiming to find a policy to price the option with the satisfaction for the only necessary demand of hedging risk, or other kind of statement, catch the optimal or nearly optimal π that minimizes the ask price $C_t(\tilde{S}_t)$. Our goal can be achieved by using reinforcement learning which translates the aim into learning π that maximizes the expected cumulative reward via maximizing the state-value function and the action-value function [30]. Reinforcement learning is found on a MDP. In our scenario, the stock price $\tilde{S}_t \in \mathcal{S}$ is the state and the position in the stock $\tilde{a}_t \in \mathcal{A}$ is the action. Accordingly, the sample space of the stock price $\mathcal{S}$ and the position $\mathcal{A}$ respectively form the state space and the action space. The state-value function for a policy π at time t is denoted as $V_t^\pi(\tilde{S}_t)$ and it is designed as the negative of the option price. Equation (29) formulates the chronological recursive relationship that expresses the portfolio's value at time t in respect of its value at the subsequent time $t+1$. We can exploit this chronological recursive relationship to find the Bellman equation for the state-value function and the reward:

$$V_t^\pi(\tilde{S}_t) = -C_t(\tilde{S}_t) \quad (37)$$

$$= \mathbb{E}_t \left[-\Pi_t - \eta \sum_{t'=t}^T e^{-r(t'-t)} \text{Var} [\Pi_{t'} | \mathcal{F}_{t'}] | \mathcal{F}_t \right] \quad (38)$$

$$= \mathbb{E}_t \left[-\Pi_t - \eta \text{Var} [\Pi_t] - \eta \sum_{t'=t+1}^T e^{-r(t'-t)} \text{Var} [\Pi_{t'} | \mathcal{F}_{t'}] | \mathcal{F}_t \right] \quad (39)$$

$$= \mathbb{E}_t \left[R_t(\tilde{S}_t, \tilde{a}_t, \tilde{S}_{t+1}) + \gamma V_{t+1}^\pi(\tilde{S}_{t+1}) \right], \quad (40)$$

Equation (38) is obtained by plugging Equation (30) into Equation (37). Equation (40) is the Bellman equation for $V_t^\pi(\tilde{S}_t)$, where the reward $R_t(\tilde{S}_t, \tilde{a}_t, \tilde{S}_{t+1})$ is derived as below by using the chronological recursive relationship given in Equation (29):

$$\begin{aligned} R_t(\tilde{S}_t, \tilde{a}_t, \tilde{S}_{t+1}) &= \gamma(\tilde{a}_t \triangle S_t + \kappa f_{\theta^*}(F) S_t) - \eta \text{Var} [\Pi_t | \mathcal{F}_t] \\ &= \gamma(\tilde{a}_t \triangle S_t + \kappa f_{\theta^*}(F) S_t) \\ &\quad - \eta \gamma^2 \mathbb{E}_t \left[\left(\dot{\Pi}_{t+1} - \left(\tilde{a}_t \triangle \dot{S}_t + \kappa f_{\theta^*}(F) \tilde{a}_t \dot{S}_t \right) \right)^2 \right], \end{aligned} \quad (41)$$

where $\gamma = e^{-r}$ is the discount rate, $\dot{\Pi}_{t+1} = \Pi_{t+1} - \bar{\Pi}_{t+1}$, $\Delta \dot{S}_t = \Delta S_t - \Delta \bar{S}_t$, $\dot{S}_t = S_t - \bar{S}_t$, and $(\cdot)$ denotes the sample mean. Maximizing the reward is identical to maximize the cumulative risk-adjusted return as the last term of the reward stands for the risk at time t and the first term quantizes the

portfolio's return between time t and $t + 1$. When η tends to infinity, the hedge becomes a pure risk hedge.

The action-value function (Q -function) returns the value of taking action a in state s under a policy π at time t . Consequently, Q -function can be interpreted as a state-value function under the condition of $\tilde{S}_t = s$ and $\tilde{a}_t = a$:

$$Q_t^\pi(s, a) = \mathbb{E}_t \left[-\Pi_t \mid \tilde{S}_t = s, \tilde{a}_t = a \right] - \mathbb{E}_t \left[\eta \sum_{t'=t}^T e^{-r(t'-t)} \text{Var} [\Pi_{t'} \mid \mathcal{F}_{t'}] \mid \tilde{S}_t = s, \tilde{a}_t = a \right]. \quad (42)$$

We use $Q_t^{\pi^*}(\tilde{S}_t, \tilde{a}_t^*)$ to stand for the Q -function value which is maximized by the optimal policy π^* which chooses the optimal action $\tilde{a}_t^*$ in state $\tilde{S}_t$ at time t :

$$Q_t^{\pi^*}(\tilde{S}_t, \tilde{a}_t^*) = \max_{\pi} Q_t^\pi(\tilde{S}_t, \tilde{a}_t) = \max_{\tilde{a}_t \in \mathcal{A}} Q_t^\pi(\tilde{S}_t, \tilde{a}_t), \quad (43)$$

where $\pi^* = \pi^*(\tilde{S}_t, t)$ and $\tilde{a}_t^* = \tilde{a}_t^*(\tilde{S}_t)$ are both dependent on $\tilde{S}_t$.

The Bellman optimality equation for the action-value function has a similar structure to the Bellman equation for the state-value function (Equation (40)) and can be expanded by plugging Equation (41) into it:

$$\begin{aligned} Q_t^\pi(\tilde{S}_t, \tilde{a}_t) &= \mathbb{E}_t \left[R_t(\tilde{S}_t, \tilde{a}_t, \tilde{S}_{t+1}) + \gamma Q_{t+1}^{\pi^*}(\tilde{S}_{t+1}, \tilde{a}_{t+1}^*) \right] \\ &= \gamma \mathbb{E}_t \left[\tilde{a}_t \triangle S_t + \kappa f_{\theta^*}(F) \tilde{a}_t^* S_t + Q_{t+1}^{\pi^*}(\tilde{S}_{t+1}, \tilde{a}_{t+1}^*) \right] \\ &\quad - \eta \gamma^2 \mathbb{E}_t \left[\left(\dot{\Pi}_{t+1} - (\tilde{a}_t \triangle \dot{S}_t + \kappa f_{\theta^*}(F) \tilde{a}_t \dot{S}_t) \right)^2 \right]. \end{aligned} \quad (44)$$

The terminal condition gives the optimal Q -function value at expiry in terms of the terminal payoff:

$$Q_T^{\pi^*}(\tilde{S}_T^u, \tilde{a}_T^*) = -H(S_T) - \text{Var}[H(S_T)], \quad (46)$$

where $\tilde{a}_T^* = 0$.

Our purpose is to find $\pi^*(\tilde{S}_t, t)$ or $\tilde{a}_t^*(\tilde{S}_t)$, and $Q_t^{\pi^*}(\tilde{S}_t, \tilde{a}_t^*)$. However, the unknown value of $\tilde{S}_t$ and the incalculable expectation in the Bellman optimality equation prevents us from achieving our goal. For the removal of the barrier, we use the cross-sectional information $\tilde{\mathcal{F}} = \{\tilde{S}_t^u\}_{u=1, \dots, U}$ to take as many likely values of $\tilde{S}_t$ into account as possible, where $\tilde{S}_t^u = \ln S_t^u - \left(\mu - \frac{\sigma_s^2}{2}\right)t$. We believe that our simulated price paths have the same constant drift as those follows a lognormal distribution. Moreover, the incalculable expectation in the Bellman optimality equation is approximated by an empirical mean of U observations.

In our scenario, the state space and the action space are both continuous. We adopt basis functions, for example basis-splines, as an intermediary for mapping two continuous spaces. A state sample $\tilde{S}_t$ is first mapped to a set of basis functions and the optimal action $\tilde{a}_t^*$ is then expressed as a linear combination of basis functions weighted by the optimal coefficients w_{nt}^* :

$$\tilde{a}_t^* = \sum_{n=1}^B w_{nt}^* \psi_n(\tilde{S}_t), \quad (47)$$

where B represents the number of basis functions. Similarly, the optimal Q -function value can be formulated as a linear combination of basis functions weighted by the optimal coefficients φ_{nt}^* :

$$Q_t^{\pi^*}(\tilde{S}_t, \tilde{a}_t^*) = \sum_{n=1}^B \varphi_{nt}^* \psi_n(\tilde{S}_t). \quad (48)$$

Our objective of finding optimal policy and Q -function value is consequently transformed into learning the optimal coefficients w_{nt}^* and φ_{nt}^* . We elaborate the derivation of the following closed-form solution of w_{nt}^* in Appendix B:

$$w_t^* = \mathbf{E}_t^{-1} \mathbf{D}_t, \quad (49)$$

$$\begin{aligned} E_{nm}^t &= \sum_{u=1}^U \left[\psi_m(\tilde{S}_t^u) \triangle \dot{S}_t \Xi \sum_{n=1}^B \psi_n(\tilde{S}_t^u) \right] \\ &\quad + \sum_{u=1}^U \left[\kappa f_{\theta^*}(F) \psi_m(\tilde{S}_t^u) \dot{S}_t \Xi \sum_{n=1}^B \psi_n(\tilde{S}_t^u) \right], \end{aligned} \quad (50)$$

$$D_n^t = \sum_{u=1}^U \left[\left(\frac{\triangle S_t^u}{2\eta\gamma} + \dot{\Pi}_{t+1} \Xi + \frac{\kappa f_{\theta^*}(F) S_t^u}{2\eta\gamma} \right) \sum_{n=1}^B \psi_n(\tilde{S}_t^u) \right], \quad (51)$$

$$\Xi = (\triangle \dot{S}_t + \kappa f_{\theta^*}(F) \dot{S}_t), \quad (52)$$

where w_t^* is a B -dimensional vector with entry w_{nt}^* , $\mathbf{E}_t$ is a matrix of size $B \times B$ with entry E_{nm}^t and $\mathbf{D}_t$ is a B -dimensional vector with entry D_n^t . We can also introduce a regularization parameter ς with a very small value as $\mathbf{E}_t + \varsigma \mathbf{I}_B$. This regularization parameter prevents $\mathbf{E}_t$ from becoming a singular matrix which does not have an inverse.

Once we obtain the optimal action $\tilde{a}_t^*$ from Equation (49) and Equation (47), the optimal Q -function value at time t can be inferred by an expectation in terms of the reward at time t and the optimal Q -function value at time $t + 1$ through the Bellman optimality equation (Equation (44)) at the optimal action $\tilde{a}_t^*$. However, in practice we take the sample from our simulated path instead of the expectation to infer the optimal Q -function value at time t and this will produce error e_u which is the deviation of the observed value from the true value:

$$\begin{aligned} Q_t^{\pi^*}(\tilde{S}_t, \tilde{a}_t^*) &= \mathbb{E}_t \left[R_t(\tilde{S}_t, \tilde{a}_t^*, \tilde{S}_{t+1}) + \gamma Q_{t+1}^{\pi^*}(\tilde{S}_{t+1}, \tilde{a}_{t+1}^*) \right] \\ &= R_t(\tilde{S}_t^u, \tilde{a}_t^*, \tilde{S}_{t+1}^u) + \gamma Q_{t+1}^{\pi^*}(\tilde{S}_{t+1}^u, \tilde{a}_{t+1}^*) + e_u. \end{aligned} \quad (53)$$

The optimal Q -function value at time t can be learned by reaching the criteria of minimizing the sum of squared errors in all simulated paths:

$$\begin{aligned}\varphi_{nt}^* &= \min_{\varphi_{nt}} \sum_{u=1}^U e_u^2 \\ &= \min_{\varphi_{nt}} \sum_{u=1}^U \left(F_u - \sum_{n=1}^B \varphi_{nt} \psi_n(\tilde{S}_t^u) \right)^2, \quad (54)\end{aligned}$$

where $F_u = R_t(\tilde{S}_t^u, \tilde{a}_t^*, \tilde{S}_{t+1}^u) + \gamma Q_{t+1}^*(\tilde{S}_{t+1}^u, \tilde{a}_{t+1}^*)$.

The optimal coefficients φ_{nt}^* is approximated by the least square estimator $\hat{\varphi}_t^*$:

$$\begin{aligned}\varphi_t^* &\approx \hat{\varphi}_t^* \\ &= \mathbf{G}_t^{-1} \mathbf{H}_t, \quad (55)\end{aligned}$$

$$G_{nm}^t = \psi_n(\tilde{S}_t^u) \psi_m(\tilde{S}_t^u), \quad (56)$$

$$H_n^t = \psi_n(\tilde{S}_t^u) \left(R_t(\tilde{S}_t^u, \tilde{a}_t^*, \tilde{S}_{t+1}^u) + \gamma Q_{t+1}^*(\tilde{S}_{t+1}^u, \tilde{a}_{t+1}^*) \right), \quad (57)$$

where φ_t^* is a B -dimensional vector with entry φ_{nt}^* , $\mathbf{G}_t$ is a matrix of size $B \times B$ with entry G_{nm}^t and $\mathbf{H}_t$ is a B -dimensional vector with entry H_n^t .

We are ultimately able to predict the option price $C_t(\tilde{S}_t) = -Q_t^*(\tilde{S}_t, \tilde{a}_t^*)$ and the hedge position $\tilde{a}_t^*$ backward recursively at any time until the expiration date by triggering the following algorithm:

Algorithm of learning option price and hedge position

- 1: Calculate $Q_T^*(\tilde{S}_T^u, \tilde{a}_T^*)$ with Eq. (46) and Π_T with Eq. (28)
 - 2: **for** $t = T - 1$ **to** $t = 0$ **do**
 - 3: Calculate w_t^* with Eq.(49), $\tilde{a}_t^*$ with Eq. (47) and Π_t with Eq. (29)
 - 4: Calculate $R_t(\tilde{S}_t^u, \tilde{a}_t^*, \tilde{S}_{t+1}^u)$ with Eq. (41)
 - 5: Calculate φ_t^* with Eq.(55) and $Q_t^*(\tilde{S}_t, \tilde{a}_t^*)$ with Eq. (48)
 - 6: **end for**
-

V. EXPERIMENTAL EVALUATION

We evaluate our entire algorithm/process on a task of hedging and pricing a European call option based on 100 shares of the underlying GME stock with a period of 6 trading days and a strike price of 100 USD. Suppose that we have learned that the leading investors are the same group of investors who took on the role as PI during Jan. 26, 2021-Feb. 02, 2021. Therefore, our first step is to probabilistically imitate their previous successive decision making from the hidden motives behind them through our visible-hidden Markov network as we presented in the section III. The second step is to transform likely evolutionary paths of PI's decision into probable price paths of the underlying stock by our B-DNN inference-regression process introduced in the section II. For the above two steps, we train a B-DNN and a visible-hidden Markov network on an episode basis with the training data set $\mathcal{D}_{i,i=1\dots 6}$. The simulated price paths along with the real price paths used as the training data samples are compared in Figure

Figure 7: Imitative GME price movements

Figure 8: Reinforcement learn option price and hedging position ground upon imitative PI's chronological decision making

7a. We remove the constant drift from all the prices paths and plot the real state $\tilde{S}_t$ and the simulated states $\tilde{S}_t^u$ in Figure 7b. The drift rate is calculated with $\mu = 0.05$ and $\sigma_s = 0.2926$. The value of σ_s is derived from averaging the daily VIX index from Jan. 26, 2021 to Feb. 02, 2021. The total number of the simulated paths is 1000 and only 10 paths are uniformly sampled to be shown here. The training samples are sampled in each 5-minute time slot. The simulated paths are also with a time resolution of 5 minutes. The price jumps that appear in Figure 7a and 7b are owing to the pricing data in the first time slot (i.e., 09:30 to 09:35). The open price and the volume in this time slot deviate far from the average owing to their high dependence on the pre-market trading and news. For this reason, the training dataset does not include the pricing data in the first 5-min time slot, and we interpolate the real value in the figures. The Visible-hidden Markov network is abbreviated to VHMN in the figures for the benefit of space. We can observe that the visible-hidden Markov network produces imitative price paths that have highly matched behavior pattern with the

real price path. As a consequence, the visible-hidden Markov network is a powerful tool to cope with the intrinsic instability in financial data by engendering imitative variations for each specific scenario.

Our final phase of this task is to perform the algorithm of learning option price and hedge position. We rebalance the portfolio every 5 minutes during 6 trading days. The risk-free interest rate $r = 1.059\%$ is taken the average of daily US 10-year bond yield from Jan. 26, 2021 to Feb. 02, 2021. We set $\kappa = 0$ as we only hedge a mini option based on 100 shares. We also give $\frac{1}{\eta} = 0$ for having pure risk hedge. $\varsigma = 0.001$ is used to avoid the error of inverting a singular matrix. The optimal hedging position is shown in Figure (8a), where the value on the graph needs to be multiplied by 100 to correspond to 100 shares. It is exhibited that the option is more actively hedged when closer to the expiration date. The value of the replicating portfolio jumps near the maturity date and the hedging strategy effectively protects the portfolio against the risk of running out of money as shown in Figure (8b). The optimal Q -function value converges at -27.586 which means we better sell the option at the price of 2758.6 USD as illustrated in Figure (8d).

VI. CONCLUSION

In order to address the challenges posed by intrinsic low SNR and instability in financial data, we innovatively exploit imitative cross-sectional information to reinforcement learn option price and hedging position. When we turn to behavior finance, another challenge of identifying the leading investor's behavior and the stock price change comes along. To this end, we take advantage of the excellent features of B-DNN to explore non-deterministic relations by a market data driven fashion.

REFERENCES

- [1] Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–654.
- [2] Merton, R.C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125–144.
- [3] Heston, S., 1993. Closed-form solution for options with stochastic volatility, with application to bond and currency options. *Review of Financial Studies* 6, 327–343.
- [4] G. E. Monahan, “State of the art—a survey of partially observable markov decision processes: theory, models, and algorithms,” *Management science*, vol. 28, no. 1, pp. 1–16, 1982.
- [5] Hornik, K., 1991. Approximation capabilities of multilayer feedforward networks. *Neural Networks* 4, no. 2, 251–257.
- [6] Jason Yosinski, Jeff Clune, Yoshua Bengio, and Hod Lipson. 2014. How transferable are features in deep neural networks? In *Advances in neural information processing systems*, pages 3320–3328.
- [7] J. Schmidhuber, “Deep learning in neural networks: An overview,” *Neural Networks*, vol. 61, pp. 85–117, 2015.
- [8] Fan, J., Ma, C., and Zhong, Y. (2021). A Selective Overview of Deep Learning *Statistical Science* 36, 264–290.
- [9] Xin Jin, Note on Backpropagation in Neural Networks, <https://hal.archives-ouvertes.fr/hal-02265247>, 2019.
- [10] Jospin, L. V., Buntine, W., Boussaid, F., Laga, H., and Bennamoun, M., “Hands-on bayesian neural networks—a tutorial for deep learning users,” *arXiv preprint arXiv:2007.06823*, 2020.
- [11] Radford M. N., “Bayesian Learning for Neural Networks,” Springer-Verlag, Berlin, Heidelberg, 1996. ISBN 0387947248.
- [12] S. Sun, G. Zhang, J. Shi, and R. Grosse, “Functional Variational Bayesian Neural Networks,” *arXiv preprint arXiv:1903.05779*, 2019.
- [13] Wang, H.; Yeung, D.Y. A survey on Bayesian deep learning. *ACM Comput. Surv. (CSUR)* 2020, 53, 1–37.

- [14] Kailai Xu and Eric Darve. Calibrating multivariate lévy processes with neural networks. In *Mathematical and Scientific Machine Learning*, pages 207–220. PMLR, 2020.
- [15] Olivier Pironneau, Calibration of heston model with keras, <https://hal.sorbonne-universite.fr/hal-02273889>, 2019
- [16] C. Finn, S. Levine, and P. Abbeel. Guided cost learning: Deep inverse optimal control via policy optimization. In *Proceedings of the 33rd International Conference on Machine Learning*, 2016.
- [17] Ziebart, B., Maas, A., Bagnell, J. A., and Dey, A. K. Maximum entropy inverse reinforcement learning. In *AAAI Conference on Artificial Intelligence*, 2008.
- [18] J. Yang, X. Ye, R. Trivedi, H. Xu, and H. Zha. Learning deep mean field games for modeling large population behavior. *arXiv preprint arXiv:1711.03156*, 2017.
- [19] Faraz Torabi, Garrett Warnell, and Peter Stone. Behavioral cloning from observation. *arXiv preprint arXiv:1805.01954*, 2018.
- [20] Mohammad E. K., Alexander I., Ehsan A., and Maciej K., “Approximate inference turns deep networks into Gaussian processes,” *Advances in neural information processing systems*, 2019.
- [21] Rasmussen C. E. and Williams C. K. I., “Gaussian Processes for Machine Learning,” Cambridge, MA, USA: MIT Press, 2006.
- [22] Rue, H., Martino, S. and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 71 319–392.
- [23] K. Kristensen, A. Nielsen, C. W. Berg, H. Skaug, and B. Bell. TMB: Automatic Differentiation and Laplace Approximation. *Journal of Statistical Software*, 70(1):1–21, 2016.
- [24] L. Tierney and J. B. Kadane, “Accurate approximations for posterior moments and marginal densities,” *J. Amer. Statistical Assoc.*, vol. 81, no. 393, pp. 82–86, 1986.
- [25] Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money*. New York: Harcourt Brace and Co.
- [26] Gao, P. (2008). Keynesian beauty contest, accounting disclosure, and market efficiency. *Journal of Accounting Research*, 46:785–807.
- [27] Stamp, M.: A revealing introduction to hidden Markov models. <http://www.cs.sjsu.edu/~stamp/RUA/HMM.pdf> (2011)
- [28] I. Halperin. QLBS: Q-learner in the Black-Scholes (-Merton) worlds. *arXiv:1712.04609*, 2017.
- [29] J. Hull, *Options, Futures, and Other Derivatives*, 6th ed., Englewood Cliffs, NJ, USA: Prentice-Hall, 2006.
- [30] Richard Sutton and Andrew Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.

APPENDIX A

Denote likelihood times prior as $a(\theta) := p(\mathcal{D}|\theta)p(\theta)$ which satisfies 2nd-order sufficient conditions: $a'(\theta^*) = \mathbf{0}$ and $a''(\theta^*)$ is positive definite, where θ^* is a strict local maximizer. We do a Taylor series approximation of $\log a(\theta)$ around the location of its maximum to give:

$$\log a(\theta) \approx \log a(\theta^*) + \frac{1}{2}(\theta - \theta^*)^T \nabla_{\theta\theta}^2 \log a(\theta^*)(\theta - \theta^*) \quad (58)$$

Plugging this truncated Taylor expansion of $\log a(\theta)$ into the posterior $P(\theta|\mathcal{D})$ will demonstrate the posterior as the gaussian distribution $\mathcal{N}(\theta | \theta^*, [-\nabla_{\theta\theta}^2 \log a(\theta^*)]^{-1})$:

$$\begin{aligned}
P(\boldsymbol{\theta}|\mathcal{D}) &= \frac{P(\mathcal{D}|\boldsymbol{\theta})P(\boldsymbol{\theta})}{\int P(\mathcal{D}|\boldsymbol{\theta})P(\boldsymbol{\theta})d\boldsymbol{\theta}} \\
&= \frac{\exp(\log a(\boldsymbol{\theta}))}{\int \exp(\log a(\boldsymbol{\theta}))d\boldsymbol{\theta}} \\
&\approx \frac{\exp\left(-\frac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\theta}^*)^T[-\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \log a(\boldsymbol{\theta}^*)](\boldsymbol{\theta}-\boldsymbol{\theta}^*)\right)}{\int \exp\left(-\frac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\theta}^*)^T[-\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \log a(\boldsymbol{\theta}^*)](\boldsymbol{\theta}-\boldsymbol{\theta}^*)\right)d\boldsymbol{\theta}} \\
&= \frac{\exp\left(-\frac{1}{2}(\boldsymbol{\theta}-\boldsymbol{\theta}^*)^T[-\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \log a(\boldsymbol{\theta}^*)](\boldsymbol{\theta}-\boldsymbol{\theta}^*)\right)}{\sqrt{2\pi[-\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 \log a(\boldsymbol{\theta}^*)]^{-1}}}
\end{aligned}$$

(59)

When we enforce the likelihood and the prior belong to the exponential family with the following expression $P(\mathcal{D}|\boldsymbol{\theta}) := \frac{e^{-L(\boldsymbol{\theta}, \mathcal{D})}}{\int e^{-L(\boldsymbol{\theta}, \mathcal{D})}d\boldsymbol{\theta}}$ and $P(\boldsymbol{\theta}) := \frac{e^{-r(\boldsymbol{\theta})}}{\int e^{-r(\boldsymbol{\theta})}d\boldsymbol{\theta}}$, the gaussian distribution followed by the posterior becomes $\mathcal{N}(\boldsymbol{\theta}|\boldsymbol{\theta}^*, [\nabla_{\boldsymbol{\theta}\boldsymbol{\theta}}^2 L(\boldsymbol{\theta}^*, \mathcal{D}) + \lambda \mathbf{I}_M]^{-1})$, where $\mathbf{I}_M$ is an identity matrix of size $M \times M$.

APPENDIX B

From Equation (45) and Equation (47), we know $Q_t^\pi(\tilde{S}_t, \tilde{a}_t)$ is a quadratic function of $\tilde{a}_t$ and $\tilde{a}_t$ is a linear function of w_{nt} . Therefore, finding the optimal coefficients w_{nt}^* which maximize Q -function value is equivalent to sloving the following equation $\frac{-\partial Q_t^\pi(\tilde{S}_t, \tilde{a}_t)}{\partial w_{nt}}|_{w_{nt}=w_{nt}^*} = 0$:

$$\begin{aligned}
&= \mathbb{E}_t \left[(\Delta S_t + \kappa f_{\boldsymbol{\theta}^*}(F) S_t) \sum_{n=1}^N \psi_n(\tilde{S}_t) \right] - 2\eta\gamma \mathbb{E}_t \left[\left(\dot{\Pi}_{t+1} - (\tilde{a}_t \triangle \dot{S}_t + \kappa f_{\boldsymbol{\theta}^*}(F) \tilde{a}_t \dot{S}_t) \right) \left(-\triangle \dot{S}_t - \kappa f_{\boldsymbol{\theta}^*}(F) \dot{S}_t \right) \sum_{n=1}^{\mathcal{B}} \psi_{nt}(\tilde{S}_t) \right] \\
&\approx \frac{1}{U} \sum_{u=1}^U \left[(\Delta S_t^u + \kappa f_{\boldsymbol{\theta}^*}(F) S_t^u) \sum_{n=1}^{\mathcal{B}} \psi_{nt}(\tilde{S}_t^u) \right] \\
&\quad - \frac{2\eta\gamma}{U} \sum_{u=1}^U \left[\left(\dot{\Pi}_{t+1} - (\tilde{a}_t \triangle \dot{S}_t^u + \kappa f_{\boldsymbol{\theta}^*}(F) \tilde{a}_t \dot{S}_t^u) \right) \left(-\triangle \dot{S}_t^u - \kappa f_{\boldsymbol{\theta}^*}(F) \dot{S}_t^u \right) \sum_{n=1}^{\mathcal{B}} \psi_{nt}(\tilde{S}_t^u) \right],
\end{aligned} \tag{60}$$

$$\begin{aligned}
&\frac{-\partial Q_t^\pi(\tilde{S}_t, \tilde{a}_t)}{\partial w_{nt}}|_{w_{nt}=w_{nt}^*} = 0 \Rightarrow \\
&\sum_{u=1}^U \left[\left(\frac{\Delta S_t^u}{2\eta\gamma} + \dot{\Pi}_{t+1}\Xi + \frac{\kappa f_{\boldsymbol{\theta}^*}(F) S_t^u}{2\eta\gamma} \right) \sum_{n=1}^{\mathcal{B}} \psi_n(\tilde{S}_t^u) \right] = \\
&\sum_{u=1}^U \left[\left(\sum_{m=1}^{\mathcal{B}} w_{mt}^* \psi_m(\tilde{S}_t^u) \triangle \dot{S}_t \Xi + \kappa f_{\boldsymbol{\theta}^*}(F) \sum_{m=1}^{\mathcal{B}} w_{mt}^* \psi_m(\tilde{S}_t^u) \dot{S}_t \Xi \right) \sum_{n=1}^{\mathcal{B}} \psi_n(\tilde{S}_t^u) \right] \Rightarrow \\
&D_n^t = \sum_{m=1}^{\mathcal{B}} E_{nm}^t w_{mt}^*, n = 1 \dots \mathcal{B} \Rightarrow \\
&w_t^* = \mathbf{E}_t^{-1} \mathbf{D}_t,
\end{aligned} \tag{61}$$

where w_t^* is a $\mathcal{B}$ -dimensional vector with entry w_{nt}^* . $\mathbf{E}_t$ and $\mathbf{D}_t$ are respectively a matrix of size $\mathcal{B} \times \mathcal{B}$ and a $\mathcal{B}$ -dimensional vector. Their entries E_{nm}^t and D_n^t are formulated below:

$$\begin{aligned}
E_{nm}^t &= \sum_{u=1}^U \left[\left(\psi_m(\tilde{S}_t^u) \triangle \dot{S}_t \Xi + \kappa f_{\boldsymbol{\theta}^*}(F) \psi_m(\tilde{S}_t^u) \dot{S}_t \Xi \right) \sum_{n=1}^{\mathcal{B}} \psi_n(\tilde{S}_t^u) \right], \\
D_n^t &= \sum_{u=1}^U \left[\left(\frac{\Delta S_t^u}{2\eta\gamma} + \dot{\Pi}_{t+1}\Xi + \frac{\kappa f_{\boldsymbol{\theta}^*}(F) S_t^u}{2\eta\gamma} \right) \sum_{n=1}^{\mathcal{B}} \psi_n(\tilde{S}_t^u) \right], \\
\Xi &= (\triangle \dot{S}_t + \kappa f_{\boldsymbol{\theta}^*}(F) \dot{S}_t).
\end{aligned}$$