The Optimal Investment Problem

Modern Models and Deep Learning

Vladimir Piterbarg
NatWest Markets

1 Outline of the problem

- Consider the classical optimal investment allocation problem of Merton [5]: dynamic allocation between risky versus riskless securities to optimize a given utility function
- Through the lens of some of the more modern approaches such as stochastic volatility and local volatility models
- There is a huge body of literature dedicated to the optimal investment problem, often focused on theoretical results of the related dynamic programming problem
- Here, our interest lies more in the practical domain following interesting questions posed/considered in the articles [14], [15], and others to be found in https://elmfunds.com/blog/
- Q1: How much value accurate modeling of the volatility process of a risky asset adds to the utility of an investor?
- Q2: Optimal investment as a reinforcement learning problem – towards model-free optimal investment policy

2 Our setup, processes

- We consider a general process for a risky asset in the "P", or real–world, measure,

$$dX(t)/X(t) = \mu(t)\ dt + \xi(t)\ dB(t), \tag{1}$$

  where both $\mu(t)$ and $\xi(t)$ are stochastic processes

- Let $r(t)$ be the risk-free rate (also stochastic in full generality), and

$$dR(t)/R(t) = r(t)\ dt$$

  be the risk-free asset

- Consider a self-financing wealth process $P(t)$ defined by

$$P(t) = \pi_0(t)R(t) + \pi(t)X(t)$$
$$dP(t) = \pi_0(t)\ dR(t) + \pi(t)\ dX(t)$$

- With some stochastic calculus manipulations and defining $\kappa(t) \triangleq \pi(t)X(t)/P(t)$ to be the proportion of wealth invested into the risky asset, we obtain

$$dP(t)/P(t) = (r(t) + \kappa(t)\,(\mu(t) - r(t)))\,dt + \kappa(t)\xi(t)\ dB(t) \tag{2}$$

- The problem is to choose $\kappa(t)$ to optimize "utility"

3 Our setup, utility

- Let $U(w)$ be a function that we call the utility function, and consider the investor's objective to maximize

$$\mathrm{E}\left(U(P(T))\right),$$

over all admissible choices of $\kappa(\cdot)$

- We maximize the terminal wealth and not the more general problem of maximizing the utility of consumption at the same time as in [7]

- We specialize the utility function to be the family of Constant Relative Risk Aversion (CRRA) utility functions parameterized by the risk aversion parameter $\gamma \geq 1$:

$$U(w) = U(w;\gamma) = \frac{w^{1-\gamma} - 1}{1-\gamma} \tag{3}$$

- Turnpike theorem [10] tells us that for sufficiently long investment time horizons the optimal allocation for a large class of general utility functions is approximated, arbitrarily closely, by that of type (3)

- CRRA utility functions for different γ's, as a function of w, terminal wealth relative to current wealth, are plotted below
- Intuitively, higher risk aversion penalizes lower terminal wealth more

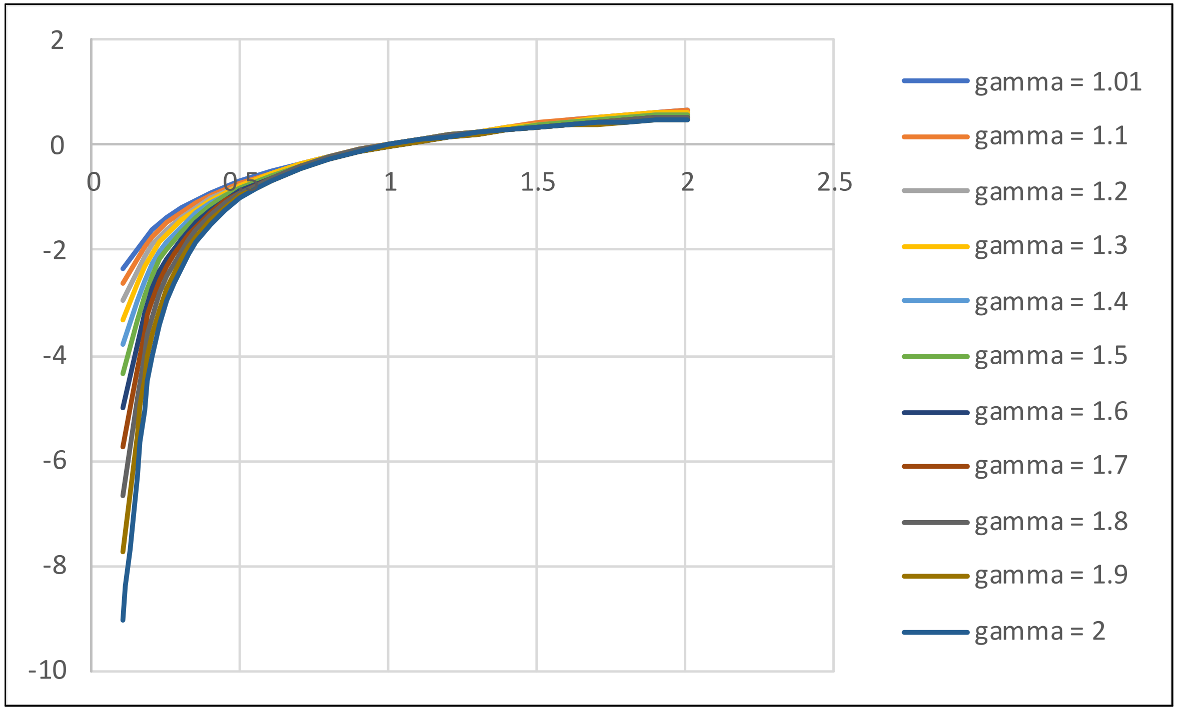

Figure 1: CRRA utility functions as a function of terminal wealth, for diferent risk aversion parameters γ.

4 Log Utility

- The case of $\gamma \to 1$ is illuminating and corresponds to the utility function being log-utility $U(w;1) = \log(w)$.
- In this case $U = \log(P)$ and we know $dP/P = \dots$ ([2]) so

$$\mathrm{E}\left(U(P(T);1)\right) = \mathrm{E}\left(\int_0^T \left(r(t) + \kappa(t)\left(\mu(t) - r(t)\right) - \frac{1}{2}\kappa(t)^2\xi(t)^2\right) dt\right)$$

- Quadratic problem; the optimal $\kappa(t)$ is given by

$$\kappa^*(t) = \frac{\mu(t) - r(t)}{\xi(t)^2}. \tag{4}$$

- This allocation is known as the *Merton ratio* (for $\gamma = 1$)
- With log-optimal allocation (4) in the portfolio process (2), the value of the optimal (under the log-utility) process $P^*_{\log}(t)$, the so-called *growth-optimal portfolio*, is

$$dP^*_{\log}(t)/P^*_{\log}(t) = \left(r(t) + \varsigma(t)^2\right)\, dt + \varsigma(t)\, dB(t),$$

where $\varsigma(t)$ is the Sharpe ratio:

$$\varsigma(t) = \frac{\mu(t) - r(t)}{\xi(t)}$$

5 Local Volatility

- To start, let us specialize the processes $\mu(t)$, $\xi(t)$ to be of the local volatility form
$$\mu(t) = \mu(t, X(t)), \quad \xi(t) = \xi(t, X(t)),$$
where $\mu(t,x)$, $\xi(t,x)$ are deterministic functions of t and x.

- Let us define the expected utility at time t, given the wealth process is at p, the risky asset process is at x, and the risky allocation function is $\kappa(\cdot)$, by
$$J(t,p,x;\kappa(\cdot)) \triangleq \mathrm{E}\left(U(P(T))\right| P(t) = p, X(t) = x).$$

- Given the functional form of the utility function (3), we can easily see that
$$J(t,p,x;\kappa(\cdot)) = \frac{p^{1-\gamma}}{1-\gamma} h(t,x;\kappa(\cdot)), \tag{5}$$
where
$$h(t,x;\kappa(\cdot)) \triangleq \mathrm{E}_{t,x} \exp\left((1-\gamma) \int_t^T \left(r(s) + \kappa(s)(\mu(s,X(s)) - r(s))\right) ds \right.$$
$$\left. - \frac{1}{2}(1-\gamma) \int_t^T \kappa(s)^2 \xi(s,X(s))^2 ds + (1-\gamma) \int_t^T \kappa(s)\xi(s,X(s))\, dB(s) \right),$$
as follows from (2) and (3). Using Girsanov's theorem and Feynman-Kac (see [12]) we can derive a PDE for h

6 LV, Feedback Control

- The function h satisfies the following PDE

$$h_t + x(\mu(t,x) + (1-\gamma)\,\kappa(t)\xi(t,x)^2)h_x + \frac{1}{2}x^2\xi(t,x)^2 h_{xx}$$
$$+ \left((1-\gamma)\,(r(t) + \kappa(t)(\mu(t,x) - r(t))) - \frac{1}{2}\,(1-\gamma)\,\gamma\kappa(t)^2\xi(t,x)^2\right) h = 0.$$

- The optimal $\kappa(t)$ is given by,

$$\kappa^*(t) = \arg\max\left(x\kappa(t)\xi(t,x)^2 h_x + \kappa(t)(\mu(t,x) - r(t))h - \frac{1}{2}\gamma\kappa(t)^2\xi(t,x)^2 h\right)$$

- Quadratic problem, solution in feedback form:

$$\kappa^*(t) = \frac{\mu(t,x) - r(t)}{\gamma\xi(t,x)^2} + \frac{x}{\gamma}\frac{h_x(t,x,\kappa^*(\cdot))}{h(t,x,\kappa^*(\cdot))}. \tag{6}$$

- The first term is the Merton ratio in the lognormal model/under log-utility (see [5]), also called the myopic allocation, static portfolio optimal allocation, or short-term optimal allocation

- The second term is a volatility (local or stochastic; depending on context) correction term. It depends on the dynamics of the volatility process

7 LV, Non-linear to Linear PDE for the Optimal Utility

- The optimal value function is defined as

$$g(t,x) \triangleq h(t,x,\kappa^*(\cdot))$$

- Plugging the expression for κ^* into the equation for h we obtain the following non-linear (quasi-linear) PDE

$$g_t + x\mu(t,x)g_x + \frac{1}{2}x^2\xi(t,x)^2 g_{xx} + (1-\gamma)\,r(t)g + \frac{1-\gamma}{2\gamma}\frac{\big(x\xi(t,x)^2 g_x + (\mu(t,x)-r)g\big)^2}{\xi(t,x)^2 g} = 0.$$

- Can transform to linear PDE (see [12]) using $g(t,x) = f(t,x)^\gamma$,

$$f_t + x\left(r(t) + \frac{1}{\gamma}(\mu(t,x) - r(t))\right) f_x + \frac{1}{2}x^2\xi(t,x)^2 f_{xx} + \frac{1-\gamma}{\gamma}\left(r(t) + \frac{1}{2\gamma}\frac{(\mu(t,x)-r(t))^2}{\xi(t,x)^2}\right) f = 0 \quad (7)$$

- The optimal allocation is given, in terms of f, as

$$\kappa^*(t) = \frac{\mu(t,x) - r(t)}{\gamma\xi(t,x)^2} + x\frac{f_x(t,x)}{f(t,x)} \quad (8)$$

8 Stochastic Volatility - Heston

- Heston model (similar results in general SV models):

$$dX(t)/X(t) = \mu\ dt + \lambda\sqrt{V(t)}\ dB^1(t), \tag{9}$$

$$dV(t) = \theta(V_0 - V(t))\ dt + \eta\sqrt{V(t)}dB^2(t), \quad V(0) = V_0 = 1, \tag{10}$$

with $\langle dB^1(t), dB^2(t)\rangle = \rho\ dt$ and μ, λ, θ, V_0, η constants

- The optimal allocation is given by

$$\kappa^*(t,v) = MertonRatio + Const \times \frac{f_v(t,v)}{f(t,v)} \tag{11}$$

where

$$\frac{f_v(t,v)}{f(t,v)} \sim \frac{d}{dv}\log\bar{\mathrm{E}}\left(\exp\left(\int_t^T \frac{ds}{\lambda^2 V(s)}\right)\middle| V(t) = v\right) \tag{12}$$

- $V(\cdot)$ behavior around 0 is of significant importance and depends (see [2]) on the Feller ratio

$$2\theta/\eta^2 \lessgtr 1$$

  – Below critical: If $2\theta < \eta^2$, then $V(\cdot) = 0$ is attainable (reflecting)
  – Above critical: If $2\theta \geq \eta^2$, then $V(\cdot) = 0$ is unattainable with probability 1

- Different regimes lead to differences in how (12) behaves

9 Heston, Optimal Allocation

- Set the time horizon $T = 10$ years, the risk-free rate $r = 0$, the risky return rate $\mu = 4\%$, the volatility of the risky asset λ to be 20%, $\rho = -90\%$, and the risk aversion parameter $\gamma = 2$
- This selection of values implies the myopic Merton ratio, the baseline allocation to the risky asset, to be
$$\frac{\mu(t,v) - r(t)}{\gamma \xi(t,v)^2} = \frac{0.04}{2 \times 0.2^2} = 0.5,$$
i.e. half of the wealth should be invested in the risky asset
- We shall investigate the impact of stochastic volatility on the adjustment to this ratio, i.e. the second term in (11)

10 Heston, Dependence on the Cutoff in SV Model

- The adjustment depends critically on the behavior of the vol process around zero
- Introduce a low bound cutoff to the stochastic volatility process for some $v_{\text{cutoff}} > 0$:

$$\xi(t, v) = \lambda \sqrt{\max(v, v_{\text{cutoff}})} \tag{13}$$

- Below critical Feller ratio (0.25) with $\theta = 0.5$ and $\eta = 2$:

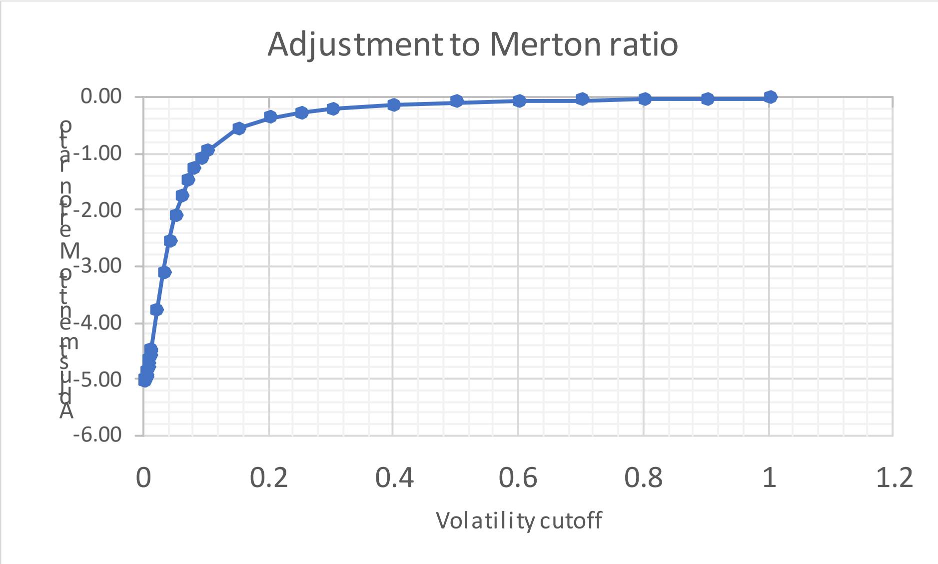

- Significant dependence on v_{cutoff} especially in $v_{\text{cutoff}} \in [0, 0.1]$ range

11 Heston, Dependence on the Cutoff in SV Model

- Feller ratio 1, with $\theta = 0.5$, $\eta = 1$
- Adjustments are not nearly as extreme, and the change in values with the cutoff parameter is much less dramatic:

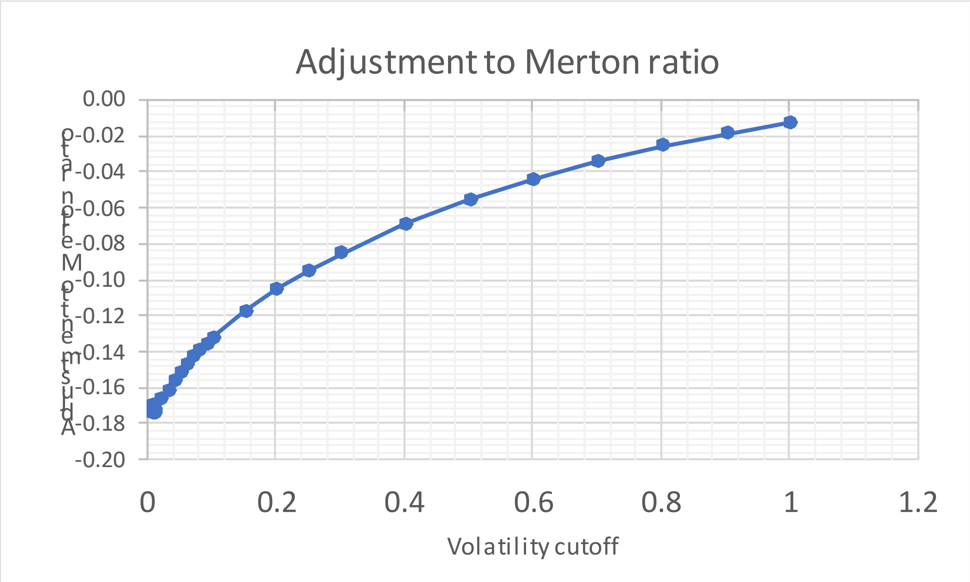

- As reported in ([12]), similar behavior with respect to cutoff parameters is observed in LV models calibrated (Dupire-style) to Heston

12 Stochastic (Heston) versus Local Volatility Model

- Let us now examine the impact of the volatility dynamics on the adjustment to the Merton ratio
- Our procedure is as follows
  - Fix Heston model parameters
  - Calculate option prices in Heston model
  - Calibrate a local volatility model, Dupire style, to those option prices
  - Calculate the Merton ratio adjustment in the original SV model and in the calibrated LV model
  - The difference quantifies the volatility dynamics vs. the observed option prices impact on optimal allocation
- Heston model (9)–(10), $\lambda = 20\%$ and $\rho = -90\%$
- Apply the volatility cutoff around $V = 0$ as in (13) for different v_{cutoff}
- Dupire's local volatility:

$$\xi_{\text{Dupire}}(t,x)^2 = \mathrm{E}\left(V(t)^2 \middle| X(t) = x \right)$$

13 S/LV, Below Critical Feller Ratio

- Case $\theta = 0.5$ and $\eta = 2$, the Feller ratio of 0.25, well below 1. The impact of the volatility cutoff (on the Heston model) on the difference between the two models

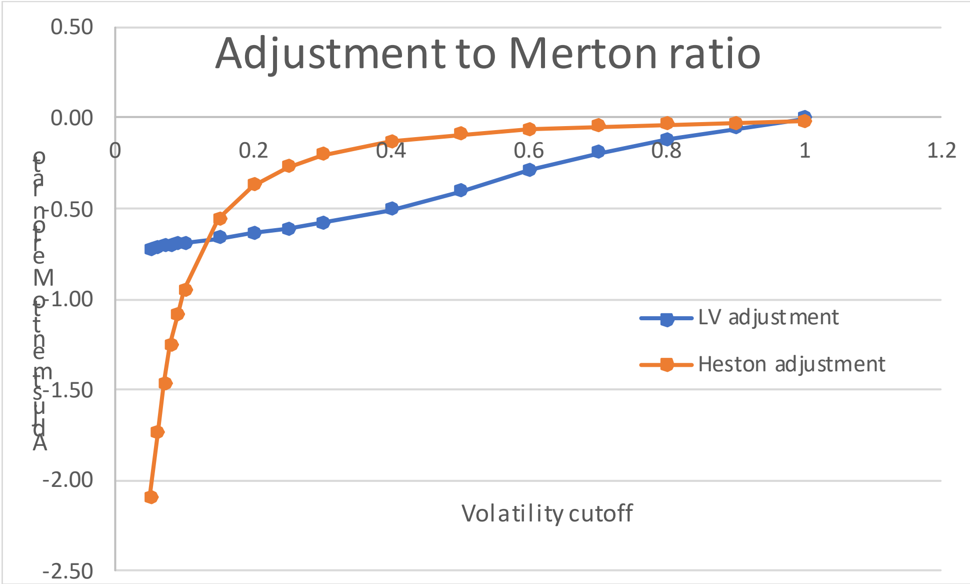

- Significant divergence between the SV and the LV model for most values of the volatility cutoff. Largely explained by the high level of dependence of the Heston model results on the cutoff parameter as discussed before

14 S/LV, At Critical Feller Ratio

- Case $\theta = 0.5$, $\eta = 1$, the Feller ratio $= 1$, at the critical level

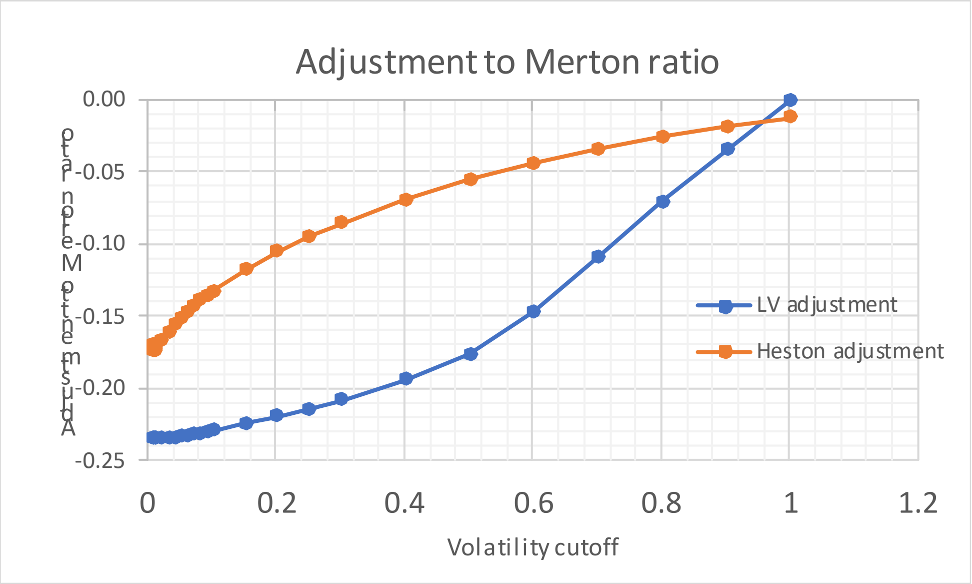

- The agreement with the local volatility model is not perfect but it is much closer
- Much less variability with respect to the volatility cutoff parameter (for reasonable values) than in the sub-Feller case

15 Towards Model-Free Optimal Asset Allocation

- Results so far relied on a specific model of the risky asset dynamics
- We have shown that, after controlling for the behavior at very low volatilities, the results are reasonably model-independent
- Can we build on this to obtain (nearly) model independent results? Using historical data?
- Techniques of deep learning helpful here? Motivation from [13] (Ritter on optimal trading strategies) and [4] (Buehler at al on deep hedging)
  - (Much) longer time horizon here, approx 10 years+
- Optimal control problem/American Monte-Carlo re-cast (re-branded?) as a reinforcement learning problem (see also Henry-Labordere's [9])
  - We start with *some* model, to have enough path realizations for "training"
  - Use Reinforcement Learning ideas/terminology to parameterize the "action-value function" and then use the "greedy" allocation algorithm

16 Discretizing the Problem

- Machine Learning requires time discretization
- Timeline $0 = T_0 < \cdots < T_N = T$, where T is the final horizon. Denote

$$X_n = X(T_n), \quad R_n = R(T_n), \quad P_n = P(T_n), \quad \pi_n = \pi(T_n), \quad \kappa_n = \kappa(T_n)$$

- Assume can only rebalance portfolio at times $\{T_n\}$, i.e. the control/allocation policy is now given by a sequence $(\kappa_0, \ldots, \kappa_{N-1})$ not continuous function $\kappa(\cdot)$
- E_n is the expected value conditional on the information at time T_n
- Discretized returns

$$\xi_n = X_n/X_{n-1} - 1, \quad \beta_n = R_n/R_{n-1} - 1 \tag{14}$$

- After some manipulations we can derive an intuitive discretized version of our portfolio process

$$P_n/P_{n-1} - 1 = \kappa_{n-1}\xi_n + (1 - \kappa_{n-1})\,\beta_n \tag{15}$$

- The return on the portfolio is equal to the weighted average of returns on the risk-free and risky asset, weighted by the allocation policy

17 Myopic Optimal Allocation

- With $U(\cdot)$ the utility function (3) we denote

$$V_n(\kappa) = U\left(P_n/P_{n-1}\right) = U\left(1 + \kappa\xi_n + (1-\kappa)\,\beta_n\right) \tag{16}$$

- Simple problem of myopic optimal allocation: choose policy κ to maximize

$$\mathrm{E}_{n-1}V_n(\kappa) = \mathrm{E}_{n-1}U\left(1 + \kappa\xi_n + (1-\kappa)\,\beta_n\right) \tag{17}$$

- Suppose we have $i = 1, \ldots, K$ realizations of our model. For each value of κ we can then easily construct the realized utility over the time period $[T_{n-1}, T_n]$,

$$V_{n,i}(\kappa) = U\left(1 + (1-\kappa)\,\beta_{n,i} + \kappa\xi_{n,i}\right), \quad i = 1, \ldots, K$$

- ML for myopic optimal allocation: $\mathrm{E}_{n-1}V_n$ is the reward and κ is the policy, in ML-speak. Choose κ (state-dependent) to maximize the reward
- This is the setup of reinforcement learning
- Specify artificial neural network (ANN) for $\mathrm{E}_{n-1}V_n(\kappa)$, the action-value function, for each κ. Then apply greedy algorithm to find optimal κ

18 Myopic Optimal Allocation via ML

- Let $\omega_1, \ldots, \omega_K \in \Omega$ represent simulated states at time T_{n-1} (asset values, volatilities, previous values of assets, etc)
- Goal: determine $\kappa = \kappa(\omega)$ ($\mathcal{F}_{n-1}$-measurable) that maximizes $\mathrm{E}_{n-1} V_n(\kappa)$ in (17):

1. Select a grid of values κ^j, $j = 1, \ldots, J$, that spans the domain of interest
2. For each j, calculate realized values of the utility on all paths

$$v_i^j \triangleq U\left(1 + \left(1 - \kappa^j\right) \beta_{n,i} + \kappa^j \xi_{n,i}\right), \quad i = 1, \ldots, K$$

3. Use your favorite ANN to "learn" the function $v^j(\omega)$ that is the best fit to the observed mapping of state to the action-value function

$$\omega_i \to v_i^j, \quad i = 1, \ldots, K$$

4. Once done for all $j = 1, \ldots, J$, use the greedy policy

$$\kappa_i^* = \arg\max_j v^j(\omega_i), \quad i = 1, \ldots, K$$

- This determines the optimal policy $\kappa^* : \Omega \to \mathbb{R}$ for the time T_{n-1}, i.e. the maximizer of $\mathrm{E}_{n-1} V_n(\kappa)$ which is the solution to the myopic optimal allocation
- A value-based optimization; policy-based optimization also possible

19 Utility – More Notations

- Recall for $n = 1, \ldots, N$,

$$P_n/P_{n-1} - 1 = \kappa_{n-1}\xi_n + (1 - \kappa_{n-1})\,\beta_n$$

- Denote

$$U_n = U_n(\boldsymbol{\kappa}_{n-1}) = U\left(P_N/P_{n-1}\right)$$

where $\boldsymbol{\kappa}_{n-1} = (\kappa_{n-1}, \ldots, \kappa_{N-1})$

- U_n depends on the series of decisions on allocation at $T_{n-1}, \ldots, T_{N-1}$
- Consider long-term (non-myopic) optimal allocations problem to maximize $\mathrm{E}_{n-1}U_n\left(\boldsymbol{\kappa}_{n-1}\right)$ and, in particular, $\mathrm{E}_0U_1\left(\boldsymbol{\kappa}_0\right)$
- As before the optimal allocation policy/optimal value function are denoted by a star,

$$\boldsymbol{\kappa}^*_{n-1} = \arg\max_{\boldsymbol{\kappa}_{n-1}} \mathrm{E}_{n-1}U_n\left(\boldsymbol{\kappa}_{n-1}\right), \tag{18}$$

$$U^*_n = \max_{\boldsymbol{\kappa}_{n-1}} \mathrm{E}_{n-1}U_n\left(\boldsymbol{\kappa}_{n-1}\right) = \mathrm{E}_{n-1}U_n\left(\boldsymbol{\kappa}^*_{n-1}\right) \tag{19}$$

(note that U^*_n is $\mathcal{F}_{T_{n-1}}$-measurable)

- Our choice of the utility function allows for *a very useful recursive relation* in U_n's,

$$U_n = V_n + \left(1 + (1 - \gamma)\,V_n\right)U_{n+1}, \quad n = N, \ldots, 0, \quad U_{N+1} \equiv 0 \tag{20}$$

20 Bellman Principle (≡"Q-Learning")

- Bellman's principle of optimality:

An optimal policy has the property that whatever the initial decision is, the remaining decisions must constitute an optimal policy with regard to the state resulting from the previous decision

- Suppose we have the optimal policy starting from time T_n, denoted by $\boldsymbol{\kappa}_n^*$. Let us consider all policies starting from time T_{n-1} of the form

$$\boldsymbol{\kappa}_{n-1} = (\kappa, \boldsymbol{\kappa}_n^*)$$

for various κ allocations at time T_{n-1}

- Recall the recursion (20),

$$U_n = V_n + (1 + (1-\gamma)\, V_n)\, U_{n+1}, \quad n = N, \ldots, 0, \quad U_{N+1} \equiv 0$$

- Hence the (recursive) equation for the value function is

$$U_n^* = \max_{\kappa} \mathrm{E}_{n-1} \left\{ V_n(\kappa) + ((1 + (1-\gamma)\, V_n(\kappa))\, U_{n+1}^*) \right\}, \quad n = N, \ldots, 1 \qquad (21)$$

21 Long-Term Optimal Allocation via ML

- Recursion (21) leads to a straightforward extension of the myopic algorithm!
- Set $U_N^* \equiv 0$, assume $U_k^* = U_k^*(\omega_i)$ for all $k = n+1, \ldots, N$ are known
- To calculate $\kappa_n^*(\omega)$ (policy) and $U_n^*(\omega)$ (optimal value function) for time T_n we use the recursion (21)
- Modify Step 2 of the algorithm in the previous section to be

$$v_i^j \triangleq U\left(1 + \left(1 - \kappa^j\right)\beta_{n,i} + \kappa^j \xi_{n,i}\right),$$
$$u_i^j = v_i^j(\kappa) + \left(\left(1 + (1-\gamma)\, v_i^j(\kappa)\right) U_{n+1}^*(\omega_i)\right), \quad i = 1, \ldots, K$$

- Then proceed as before with the ANN now representing a function $u^j(\omega)$ to be "learned" from the mapping

$$\omega_i \mapsto u_i^j, \quad i = 1, \ldots, K$$

- Optimal policy κ for time T_{n-1} is defined by

$$\kappa_{n,i}^* = \arg\max_j u^j(\omega_i), \quad i = 1, \ldots, K$$

22 Example

- LV model: regularized CEV with 50% skew
- $X(0) = 100$, $B(0) = 1$
- $\mu = 4\%$, $r = 0\%$, $\lambda = 20\%$
- $\gamma = 2$
- Action policy κ discretized over $[-2, 2]$ i.e. maximum leverage of ± 2 ($\kappa = 2$ means 200% of wealth in risky asset, -100% in bond)
- 5 years into a 10 year allocation problem

23 Observed vs Learned for Specific Action

- The action-value function for a specific action of $\kappa = 0.5$, i.e. 50% allocation to the risky asset, as a function of the asset value

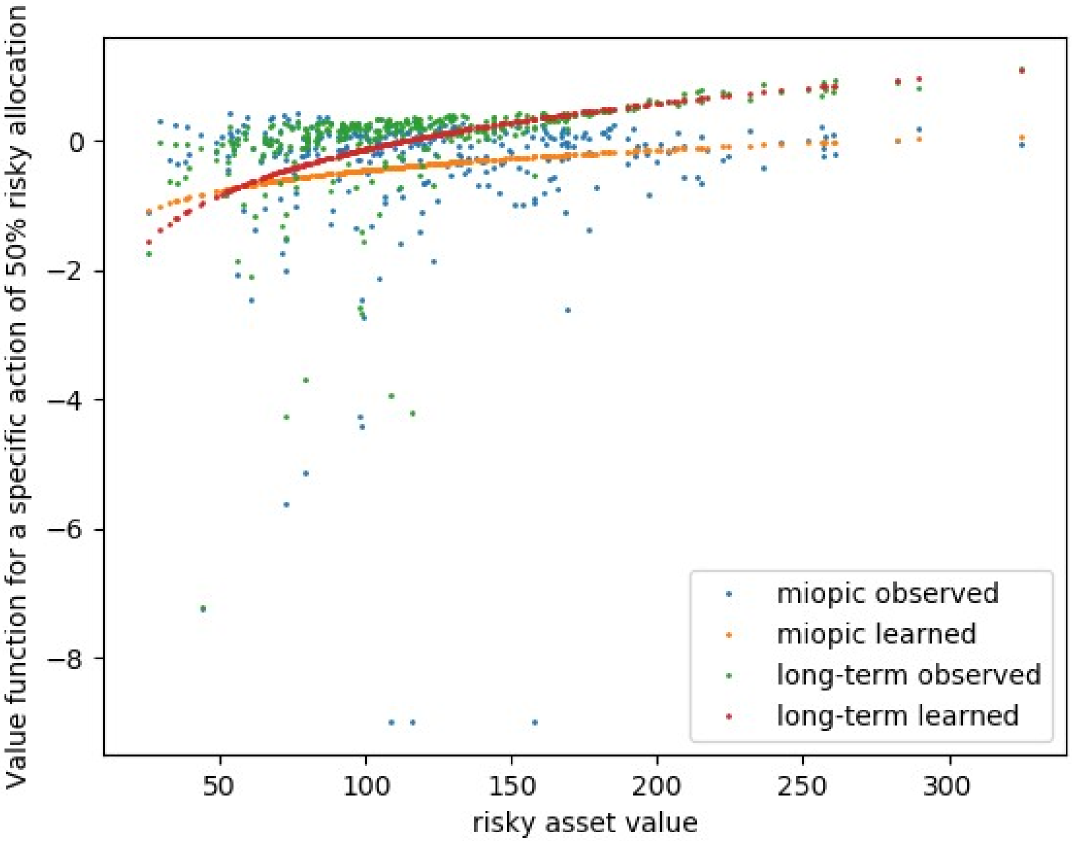

24 Action Value Function for Different Actions

- Learned action-value function for a selection of different actions that span the action space

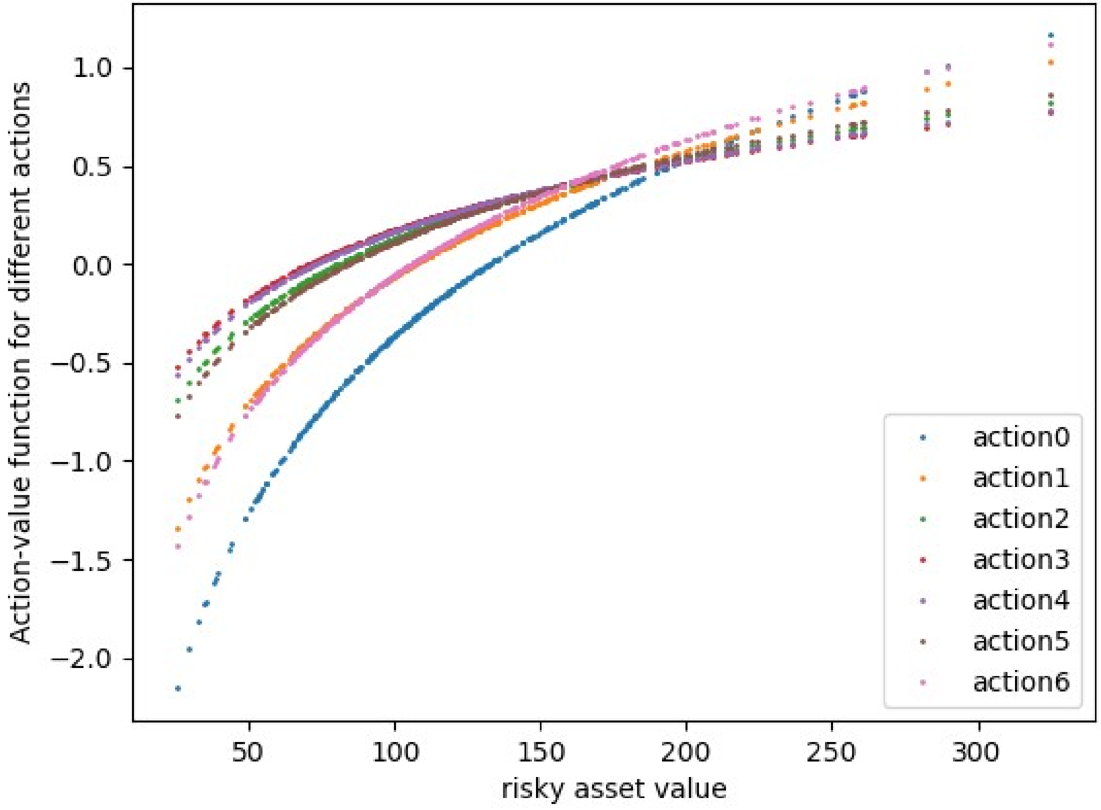

25 Optimal Value (Utility) Function

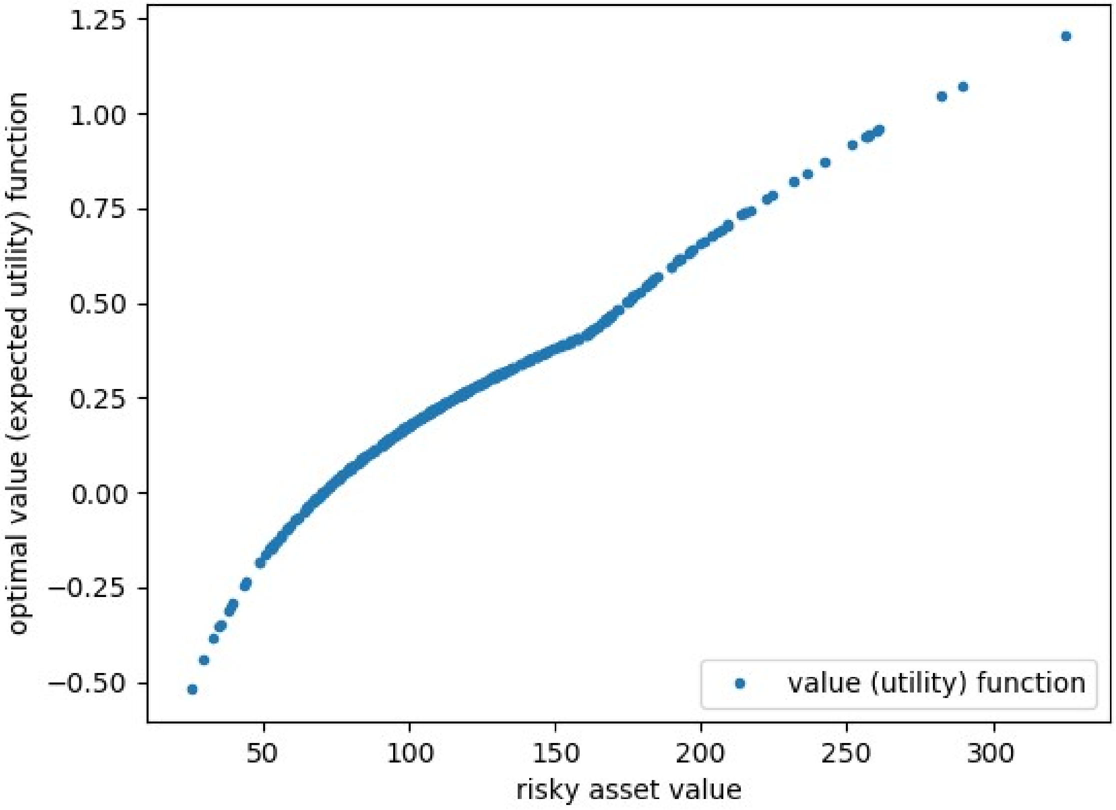

26 Optimal Policy vs Risky Asset Value

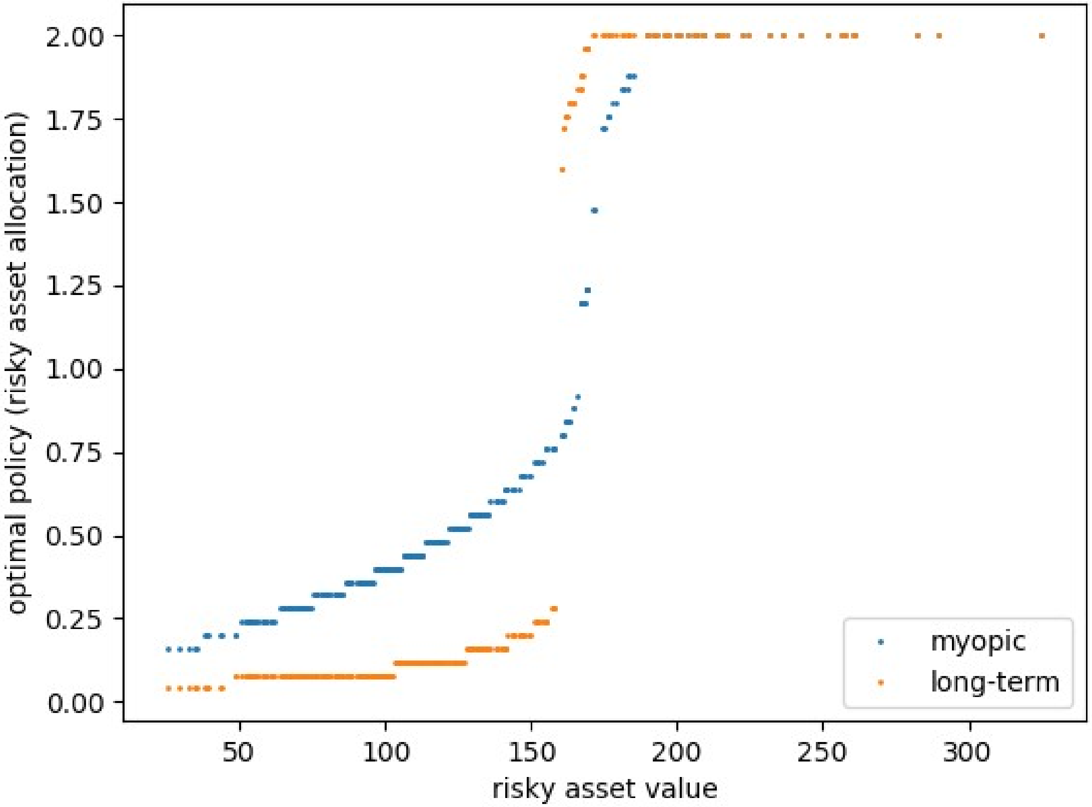

27 Discussion of ML

- An algorithm similar to an American Monte-Carlo with states (κ^j)
  - ANN: more flexible regression functions
- As presented requires many simulated paths for optimization, so a model of asset(s) evolution is still needed
  - We have tried it for some models and it works
  - We can use a much richer model than we can in PDE
- Can we use historical data?
  - Only one path of SPX
  - Resample? E.g. in the spirit of ML, use Restricted Boltzmann Machine as in Kondratyev-Schwarz [8]
- Alternatively have a very rich model
  - Fits historical path
  - Makes economic sense
- Plenty of scope for further research

28 Conclusions

- In the context of various models for the risky asset, the largely unobservable behavior of the volatility process around zero has a disproportionately large impact on optimal asset allocation
- Removing the effects around zero volatility, the near-optimal allocation can be deduced largely from the option prices and the local volatility process that fits them, i.e. in an almost model free way
- Approximate model independence is an inspiration to look into truly model-free approaches such as Reinforcement Learning
  - We derive an important recursion for the action-value function that enables reinforcement learning on simulated data
- Further research is needed to look into historical simulations/resampling for a truly model-free approach
- Some of this is covered in detail in [12]

References

[1] M. Abramovitz, I.A.Stegun. "Handbook of Mathematical Functions", 1964, National Bureau of Standards

[2] L. Andersen, V. Piterbarg. "Interest Rate Modeling", 2010, Atlantic Financial Press

[3] E. Boguslavskaya, D. Muravey. "An explicit solution for optimal investment in Heston model", 2015, arxiv

[4] H. Buehler, L. Gonon, J. Teichmann, B. Wood. "Deep Hedging", 2018, arxiv

[5] R. Merton. "Continuous-time Finance", 199l, Blackwell Publishing Press

[6] H. Kraft. "Optimal Portfolios and Heston's Stochastic Volatility Model", 2003, working paper

[7] I. Karatzas, S. Shreve. "Brownian Motion and Stochastic Calculus", 1997, Springer

[8] A. Kondratyev, C. Schwarz. "The Market Generator", 2019,. SSRN working paper, https://ssrn.com/abstract=3384948

[9] P. Henry-Labordere. "CVA and IM: Welcome to the Machine", March 2019, Risk Magazine

[10] C. Huang and T. Zariphopoulou. "Turnpike Behavior of Long-term Investments", 1999, Finance and Stochastics 3, pp. 15-34.

[11] J. L. Kelly. "A New Interpretation of Information Rate", 1956, Bell System Technical Journal, 35 (4): 917–926.

[12] V. Piterbarg. "The Optimal Investment Problem in Stochastic and Local Volatility Models", The Journal of Investment Strategies, December 2018, Volume 7, Number 4, pp.1–26

[13] G. Ritter. Machine learning for trading. October 2017, Risk Magazine, pp. 84-89

[14] J. White. "Part I: Practical Utility, Risk Aversion, and Investment Sizing", 2017, Elm funds working paper, https://elmfunds.com

[15] J. White, V. Haghani. "Measuring the Fabric of Felicity", 2018, Elm funds working paper, https://elmfunds.com/measuring-the-fabric-of-felicity/

[16] T. Zariphopoulou. "A solution approach to valuation with unhedgeable risks", 2001, Finance and Stochastics, 5, 61-82