

Efficient Calibration of Multi-Agent Market Simulators from Time Series with Bayesian Optimization

Yuanlu Bai
Henry Lam
Columbia University, USA

YB2436@COLUMBIA.EDU
HENRY.LAM@COLUMBIA.EDU

Svitlana Vyetenko
Tucker Balch
J.P.Morgan AI Research, USA

SVITLANA.S.VYETRENKO@JPMCHASE.COM
TUCKER.BALCH@JPMCHASE.COM

Abstract

Multi-agent market simulation is commonly used to create an environment for downstream machine learning or reinforcement learning tasks, such as training or testing trading strategies before deploying them to real-time trading. In electronic trading markets only the price or volume time series, that result from interaction of multiple market participants, are typically directly observable. Therefore, multi-agent market environments need to be calibrated so that the time series that result from interaction of simulated agents resemble historical – which amounts to solving a highly complex large-scale optimization problem. In this paper, we propose a simple and efficient framework for calibrating multi-agent market simulator parameters from historical time series observations. First, we consider a novel concept of eligibility set to bypass the potential non-identifiability issue. Second, we generalize the two-sample Kolmogorov-Smirnov (K-S) test with Bonferroni correction to test the similarity between two high-dimensional time series distributions, which gives a simple yet effective distance metric between the time series sample sets. Third, we suggest using Bayesian optimization (BO) and trust-region BO (TuRBO) to minimize the aforementioned distance metric. Finally, we demonstrate the efficiency of our framework using numerical experiments.

1. Introduction

Simulation is commonly used in financial markets to model counterfactual scenarios and support trading decisions. It sets up an infrastructural environment where machine learning or reinforcement learning experiments can be carried out. Experimentation with trading strategies in “live” environments can be costly or otherwise impossible, therefore, before they are released to production, trading strategies and algorithms need to undergo extensive testing inside simulated environments under a variety of markets scenarios. Additionally, it is also the case that realistic simulated environments are needed for training trading strategies using reinforcement learning. To ensure that the simulation tool is useful, it is of vital importance to calibrate the simulator parameters accurately, as the simulation model can be misleading if it is far off from reality. A peculiar calibration challenge is that in electronic trading markets, it is often the case that the agent-specific data is unavailable, and only the time series that result from interaction of multiple market agents (such as price and volume time series) are directly observable to the general public. In this paper, our main contribution is to propose an efficient learning-based framework to calibrate the multi-agent simulator parameters only from such output-level time series samples. We note that the methodology can be applied to calibrate large-scale or even black-box simulation models in general areas of study, as the

methodology does not assume specific structure of the simulators. The framework can be applied as long as the simulator generates output time series from a probability distribution determined by the input parameters and real output data are observable.

Market simulation systems usually have many parameters that are necessary to capture the market complexity. Calibrating a multi-agent simulator with respect to a large number of parameters can potentially cause the non-identifiability issue [14] – namely, a different configuration may have the same or very similar output distribution as the true configuration, so they are indistinguishable only from the output data. To address this issue, we employ a novel concept of eligibility set [1, 2], where we measure the discrepancy between the output sample distributions instead of the input parameter values and then accept or reject the candidate configuration using the idea of hypothesis testing. To be specific, with a distance metric between output time series sample sets as well as the critical value chosen accordingly, the candidate is not rejected if and only if the distance metric is below the critical value.

To define the distance metric and the critical value properly, we generalize the two-sample Kolmogorov-Smirnov (K-S) test [8] with Bonferroni correction [11] to test whether two time series sample sets are from the same distribution. As a distance metric, the K-S statistic is easy and convenient to compute, and the corresponding critical value with statistical guarantees can be computed or approximated explicitly. Before computing the K-S statistic, feature extraction can also be applied on the time series samples in order to reduce the sample dimension and hence reduce the critical value conservativeness. For simplicity, we do not apply feature extraction in our experiments, which turns out to perform well under our experimental setups. We will discuss feature extractors using stylized facts [5, 16] and learning-based methods [2, 13] and analyze their pros and cons in Appendix C.

In this paper, we choose the K-S statistic between the real and simulated time series (i.e. time series that result from interaction of multiple market agents) as an objective function for multi-agent simulator calibration. Multi-agent simulator calibration poses a difficult optimization problem. First, with thousands of (or even more) agents interacting in the system, the simulation tends to be slow, so the objective function is expensive and thus the budget of evaluations can be extremely limited. Second, the objective function is likely to lack both special structure (such as convexity) and derivative information, which impedes us from leveraging them to improve efficiency. Third, the objective function is noisy as the K-S statistic depends on random samples. To combat the above issues, we suggest using Bayesian optimization (BO) [7, 9], which has been a popular and competitive approach to optimize expensive black-box objective functions in recent years. However, it is well-known that BO does not scale well to high-dimensional cases, therefore, we recommend a state-of-the-art variant called trust region BO (TuRBO) [6], which adopts a local strategy instead.

To the best of our knowledge, we are the first to use the generalized K-S statistic as a distance between high-dimensional time series sample sets to calibrate multi-agent simulators. We are aware of its limitation of not using time series properties such as serial dependence information. We, however, believe that this approach is justified for the purpose of calibrating simulators instead of modeling time series. First, it enjoys statistical guarantees on correctness and conservativeness [2], which supports its effectiveness in high-dimensional usage. Second, compared to learning-based metrics, this statistic is more explainable, and also easier to evaluate without needing to be trained. Third, after eliminating the random noise, it is a relatively smooth function of the input configuration and thus enables us to use the Gaussian process (GP) surrogate model in BO algorithm, which is not only simple but also able to provide interval estimation besides point prediction. In the literature,

BO has been applied to calibrate agent-based models, but the objective function and the surrogate model are often chosen differently. [15] is the closest work to ours that we have found. However, they use the one-dimensional K-S statistic between daily return distributions, whereas we use a high-dimensional generalization of K-S statistic to compare the entire time series distributions. They also choose the tree-structured Parzen estimator as the surrogate model instead of GP.

As a result of numerical experiments, we observe that despite their simplicity combining the K-S statistic distance metric and the GP surrogate works well in practice – which makes us believe that our calibration framework would be of interest to the financial field.

2. Calibration Framework

2.1. Eligibility Set

Consider a simulation model which takes an input configuration $\theta \in \Theta \subset \mathbb{R}^I$ and outputs a time series sample in \mathbb{R}^O from distribution P^θ . Suppose that the real distribution is P^{real} , then ideally our goal is to find $\theta \in \Theta$ such that $P^\theta = P^{real}$, or equivalently, $d(P^\theta, P^{real}) = 0$ for any valid statistical distance $d(\cdot, \cdot)$ between probability distributions. However, the simulation model is not necessarily identifiable. That is, it is possible that there exist different configurations $\theta_1 \neq \theta_2$ such that $d(P^{\theta_1}, P^{real}) = d(P^{\theta_2}, P^{real}) = 0$. In this case, only given the output-level information, it is impossible to identify which one is indeed the “truth”. Therefore, instead of trying to recover the “best” configuration, we jointly consider the set $\{\theta \in \Theta : d(P^\theta, P^{real}) = 0\}$.

However, in practice, P^{real} is unknown and P^θ is hard to derive due to the model complexity. Thus, we can only estimate them using finite samples. Suppose real samples $X_1, \dots, X_N \sim P^{real}$ are available, and for each θ , we generate samples $Y_1^\theta, \dots, Y_n^\theta \sim P^\theta$ using the simulator. Then we can use the empirical distributions $P_N^{real}(\cdot) = \frac{1}{N} \sum_{i=1}^N \delta_{X_i}(\cdot)$ and $P_n^\theta(\cdot) = \frac{1}{n} \sum_{j=1}^n \delta_{Y_j^\theta}(\cdot)$ to approximate P^{real} and P^θ respectively. As a relaxation, we construct the statistically confident eligibility set as $\{\theta \in \Theta : d(P_n^\theta, P_N^{real}) < q\}$ where $q \in \mathbb{R}^+$ is a suitable constant threshold. In particular, when q is chosen as the critical value with significance level α , the eligibility set is actually a $(1 - \alpha)$ -level confidence region for the “true configuration”. We refer to [2] for more details about the construction and theoretical guarantees of eligibility set.

2.2. Choice of Distance Metric

The choice of distance metric d is critical for the success of calibration framework. First of all, given two sample sets from P^{real} and P^θ respectively, the distance $d(P_n^\theta, P_N^{real})$ should be easy to compute. Moreover, the choice of d affects the conservativeness which is intuitively measured by the “size” of the eligibility set. As an extreme example, if $d(P_n^\theta, P_N^{real}) \equiv 0$, then the eligibility set is extremely conservative as it contains any $\theta \in \Theta$. Thus, we aim to pick an easily-evaluated distance metric which well captures the discrepancy between two time series sample sets.

In general, we consider first applying a feature extractor $f : \mathbb{R}^O \rightarrow \mathbb{R}^K$ on each time series sample. That is, we compare K -dimensional sample sets $\tilde{X}_i = f(X_i), i = 1, \dots, N$ and $\tilde{Y}_j^\theta = f(Y_j^\theta), j = 1, \dots, n$. Note that setting f as the identity map $x \mapsto x$ implies no feature extraction. A good choice of feature extractor should reduce the dimension of the original time series without losing too much information.

Recall that the one-dimensional two-sample K-S test is a non-parametric test for equality of two distributions. To be specific, for any $k = 1, \dots, K$, we use $\tilde{X}_{i,k}$ and $\tilde{Y}_{j,k}^\theta$ to denote the k -th

component of \tilde{X}_i and \tilde{Y}_j^θ , and then $\tilde{X}_{i,k}$'s, $\tilde{Y}_{j,k}^\theta$'s are two one-dimensional sample sets. Denote their empirical CDF as $F_{N,k}$ and $F_{n,k}^\theta$ respectively, i.e. $F_{N,k}(x) = \frac{1}{N} \sum_{i=1}^N I(\tilde{X}_{i,k} \leq x)$, $F_{n,k}^\theta(x) = \frac{1}{n} \sum_{j=1}^n I(\tilde{Y}_{j,k}^\theta \leq x)$. In the one-dimensional K-S test, the statistic $\sup_{x \in \mathbb{R}} |F_{N,k}(x) - F_{n,k}^\theta(x)|$ is contrasted with the critical value $q_{N,n,\alpha}$. If the statistic is greater than the critical value, then we reject the null hypothesis that $\tilde{X}_{i,k}$'s and $\tilde{Y}_{j,k}^\theta$'s come from the same distribution with significance level α .

The K-S test cannot be easily generalized to high-dimensional case since the joint CDF needs to be considered. As a substitute, we propose pooling the discrepancies in each dimension with Bonferroni correction. That is, the high-dimensional K-S statistic is defined as

$$\max_{k=1,\dots,K} \sup_{x \in \mathbb{R}} |F_{N,k}(x) - F_{n,k}^\theta(x)|$$

and the α -level critical value is defined as $q_{N,n,\alpha/K}$. Under the null hypothesis $P^\theta = P^{real}$, we have that $P(\max_{k=1,\dots,K} \sup_{x \in \mathbb{R}} |F_{N,k}(x) - F_{n,k}^\theta(x)| > q_{N,n,\alpha/K}) \leq \sum_{k=1}^K P(\sup_{x \in \mathbb{R}} |F_{N,k}(x) - F_{n,k}^\theta(x)| > q_{N,n,\alpha/K}) \leq \sum_{k=1}^K \frac{\alpha}{K} = \alpha$. Hence, if the K-S statistic exceeds the critical value, then we can reject the null hypothesis with significance level α .

Thus, we suggest choosing the K-S statistic as the distance metric $d(P_n^\theta, P_N^{real})$ with $q = q_{N,n,\alpha/K}$. For simplicity, it is known that q can be approximated with $\sqrt{-\frac{(N+n)\log(\alpha/2K)}{2Nn}}$. From the above derivation, we know that a bad choice of feature extractor f does not harm the correctness of the eligibility set. That is, the eligibility set defined in this way is always a $(1 - \alpha)$ -level ‘‘confidence region’’ for the true configuration (if any). However, as we have explained, the choice of f affects the conservativeness. We will discuss different types of feature extractors in Appendix C and we refer to [2] for statistical guarantees on both correctness and conservativeness.

2.3. Bayesian Optimization

For a continuous configuration space $\Theta \subset \mathbb{R}^I$, we cannot generate a sample set from each $\theta \in \Theta$ and check whether it is rejected or not. Thus, the eligibility set can only be approximated with discrete points. If we find some points in the eligibility set, then we can jointly consider them in downstream tasks. That is, to estimate a target quantity of P^{real} , instead of evaluating it only with the ‘‘best’’ configuration, we use all the ‘‘eligible’’ configurations to provide a more robust estimation (see [2]), which enables us to bypass the non-identifiability issue. In the low-dimensional case such as $I = 1, 2$, the eligibility set can be approximated by grid search or random search. However, the parameter space grows exponentially in the number of parameters I . In order to find an eligible point in the higher-dimensional case, we consider minimizing the distance metric, i.e. the K-S statistic, over the entire parameter space.

As mentioned in Section 1, this global optimization problem is hard, since the objective function is expensive, black-box and noisy. To solve such problems, BO has gained popularity in recent years. The key idea of standard BO is to model the objective function as a realization of a prior GP model. After evaluating the objective values at some points, we gain some information on the objective function, and conditional on the information, the posterior GP model can be used to predict the objective function at any point. In the algorithm, we evaluate the objective values on randomly sampled initial points to gain initial information, and then repeatedly select the next point, evaluate the objective value and update the posterior model. See Appendix A for more detailed explanation

of the methodology as well as the discussion on why we stick to GP instead of more complicated alternatives as the surrogate model.

While BO generally performs well in the low-dimensional case, it is well established that it converges more slowly when the dimension is higher. In our problem, the objective function is highly expensive to evaluate due to the complexity of the simulator, so the total budget evaluation tends to be small, say several hundred. To find an eligible point within such a limited budget, we apply a state-of-the-art variant of BO called TuRBO, which adopts a local strategy by maintaining a trust region. We refer to [6] for more details. An interesting observation from our numerical experiments is that TuRBO achieves an efficiency improvement within only 100 evaluations, even if it is originally designed for larger budgets such as thousands of or even tens of thousands of evaluations.

3. Experimental Results

We apply our calibration framework to the Agent-Based Interactive Discrete-Event Simulation (ABIDES) environment [4] which simulates limit order book (LOB) exchange markets. In this paper, we illustrate our proposed calibration framework with synthetic data. We designate a synthetic ground truth simulator configuration and use it to generate “real data” with respect to which we conduct calibration. In such an experimental setting, there is no error due to model assumptions, which makes it easier to assess the performance of our calibration framework. We further demonstrate our method performance on calibration examples for two-, six- and eighteen-parameter cases. In the below experiments, we apply the KS-statistic to the concatenated minute-level mid price return and volume time series samples that result from the market agent interaction directly without prior feature extraction. Please refer to Appendix B for a more detailed introduction to ABIDES and the experiment setup, and Appendix C for a discussion about feature extractors.

3.1. Calibration Performance of BO and Fitting Performance of GP

Figure 1 shows the experiment result of applying the BO algorithm on the two-parameter calibration problem. From Figure 1(a)subfigure, we see that the found optimal point is close to the truth. In particular, it is in the eligibility set, so the calibration is successful. The color represents the predicted K-S statistic by the posterior GP model, which looks similar to the grid search plot (Figure 6(a)subfigure in Appendix C). To further justify that the posterior GP model is a good fitting for the K-S statistic, we fix one parameter as the truth and perturb the other one, and then compare the true statistic values with the confidence band predicted by the posterior GP. Figure 1(b)subfigure shows that the band is not wide, especially near the true parameter value. Moreover, the true statistic values lie in the confidence band and basically around the predicted curve. Thus, in this example, GP is able to fit the K-S statistic well.

3.2. Non-Identifiability Issue

In the two-parameter example, the eligibility set is approximately an ellipsoid around the truth, so the problem seems to be identifiable. However, with more parameters incorporated, the non-identifiability issue appears. We collect all the eligible points found in the six-parameter experiments, which forms an approximation to the eligibility set. Figure 2 plots the arrival rate against the number of value agents, and we observe that these two parameters are approximately inversely

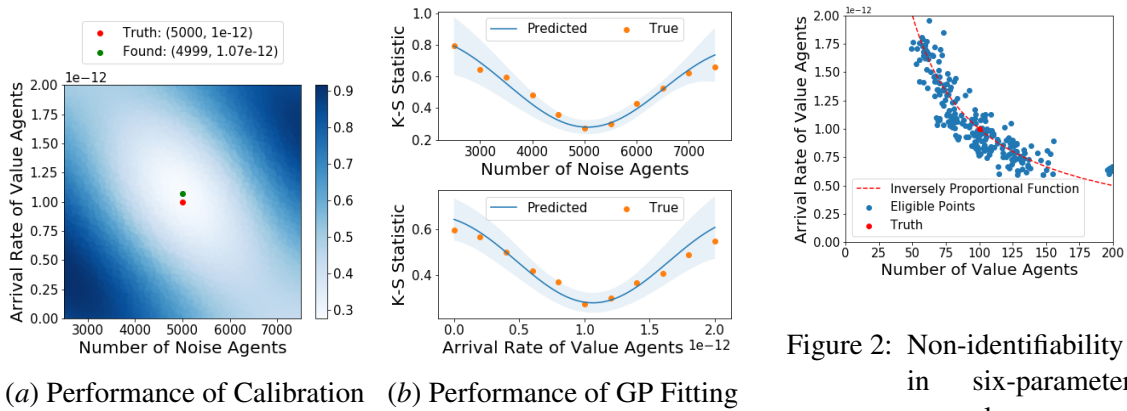


Figure 1: Performance on two-parameter example.

Figure 2: Non-identifiability in six-parameter example.

proportional to each other in these eligible points, which makes sense since the overall arrival rate of value agents is the product of these two parameters. The complete pairplot is in Appendix D.

3.3. Efficiency Improvement with TuRBO

Although the performance in the two-parameter calibration example is satisfactory, we observe that the standard BO algorithm converges more slowly as the number of parameters increases. In the six-parameter or eighteen-parameter experiments, we compare the efficiency of random search, standard BO and TuRBO with 20 or 10 random seeds to reduce randomness. Figures 3 and 4 show the experiment results. Figure 3(a)subfigure shows how the best objective value evolves as the number of evaluations increases, where we take an average over 20 experiments. Figure 3(b)subfigure compares the finally found best objective values under the 20 random seeds for each method. Moreover, if we regard finding an eligible point within 100 evaluations as a success, then Table 3(c)subfigure shows the success rate of each method. Overall, in terms of both convergence rate and success rate, TuRBO is the best while random search is the worst. Similarly, Figure 4 further supports that TuRBO still greatly outperforms the other two methods in such a high-dimensional problem.

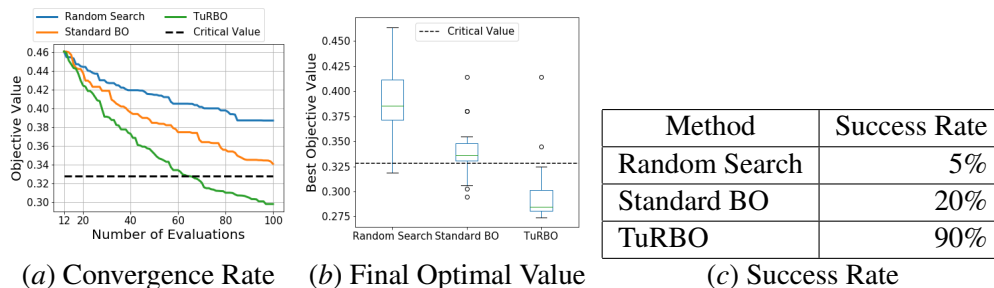


Figure 3: Comparison of random search, standard BO and TuRBO on six-parameter example.

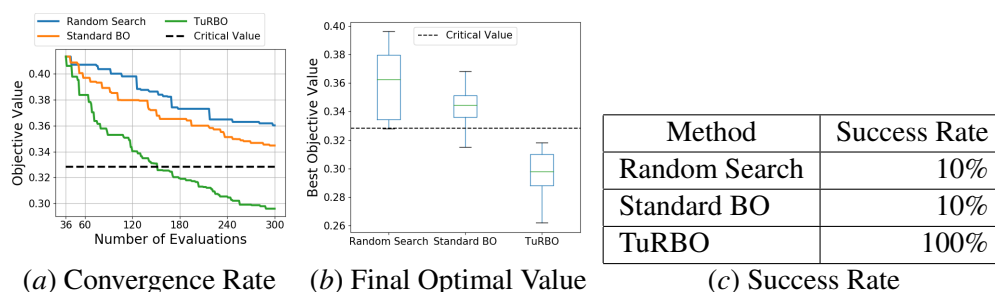


Figure 4: Comparison of random search, standard BO and TuRBO on eighteen-parameter example.

Acknowledgement

This work was done while Yuanlu Bai was working as a summer intern with J.P. Morgan AI Research, under supervision of her academic advisor Henry Lam and internship managers Svitlana Vyetenko and Tucker Balch. We also acknowledge support from the J. P. Morgan Chase Faculty Research Award.

Disclaimer

This paper was prepared for informational purposes [“in part” if the work is collaborative with external partners] by the Artificial Intelligence Research group of JPMorgan Chase and Co and its affiliates (“J.P. Morgan”), and is not a product of the Research Department of J.P. Morgan. J.P. Morgan makes no representation and warranty whatsoever and disclaims all liability, for the completeness, accuracy or reliability of the information contained herein. This document is not intended as investment research or investment advice, or a recommendation, offer or solicitation for the purchase or sale of any security, financial instrument, financial product or service, or to be used in any way for evaluating the merits of participating in any transaction, and shall not constitute a solicitation under any jurisdiction or to any person, if such solicitation under such jurisdiction or to such person would be unlawful.

References

- [1] Yuanlu Bai and Henry Lam. Calibrating Input Parameters via Eligibility Sets. In *Proceedings - Winter Simulation Conference*, volume 2020-Decem, pages 2114–2125, 2020.
- [2] Yuanlu Bai, Tucker Balch, Haoxian Chen, Danial Dervovic, Henry Lam, and Svitlana Vyetenko. Calibrating Over-Parametrized Simulation Models: A Framework via Eligibility Set. 2021. URL <https://arxiv.org/abs/2105.12893v1>.
- [3] David Byrd. Explaining Agent-Based Financial Market Simulation. 2019. URL <http://arxiv.org/abs/1909.11650>.
- [4] David Byrd, Maria Hybinette, and Tucker Hybinette Balch. ABIDES: Towards high-fidelity multi-agent market simulation. In *Proceedings of the 2020 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, pages 11–22, 2020.

- [5] Rama Cont. Empirical properties of asset returns: stylized facts and statistical issues, 2021.
- [6] David Eriksson, Michael Pearce, Jacob R Gardner, Ryan Turner, and Matthias Poloczek. Scalable global optimization via local Bayesian optimization. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- [7] Peter I. Frazier. A Tutorial on Bayesian Optimization. 2018. URL <http://arxiv.org/abs/1807.02811>.
- [8] J. L. Hodges. The significance probability of the smirnov two-sample test. *Arkiv för matematik*, 3(5):469–486, 1958.
- [9] Donald R. Jones, Matthias Schonlau, and William J. Welch. Efficient Global Optimization of Expensive Black-Box Functions. *Journal of Global Optimization*, 13(4):455–492, 1998.
- [10] Francesco Lamperti, Andrea Roventini, and Amir Sani. Agent-based model calibration using machine learning surrogates. *Journal of Economic Dynamics and Control*, 90:366–389, 2018.
- [11] Juliet Popper Shaffer. Multiple hypothesis testing. *Annual Review of Psychology*, 46(1):561–584, 1995.
- [12] Jasper Snoek, Oren Ripped, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Md Mostofa Ali Patwary, Prabhat, and Ryan P Adams. Scalable Bayesian optimization using deep neural networks. In *32nd International Conference on Machine Learning, ICML 2015*, volume 3, pages 2161–2170, 2015.
- [13] Victor Storchan, Svitlana Vyetenko, and Tucker Balch. Learning who is in the market from time series: market participant discovery through adversarial calibration of multi-agent simulators. In *ICML 2021 Workshop on Time Series*, 2021.
- [14] Albert Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 2005.
- [15] Minh Tran, Man Ngo, Duc Pham-Hi, and Marc Bui. Bayesian Calibration of Hyperparameters in Agent-Based Stock Market. In *Proceedings - 2020 RIVF International Conference on Computing and Communication Technologies, RIVF 2020*. Institute of Electrical and Electronics Engineers Inc., 2020.
- [16] Svitlana Vyetenko, David Byrd, Nick Petosa, Mahmoud Mahfouz, Danial Dervovic, Manuela Veloso, and Tucker Hybinette Balch. Get Real: Realism Metrics for Robust Limit Order Book Market Simulations. 2019. URL <http://arxiv.org/abs/1912.04941>.

Appendix A. Introduction to BO

A.1. Standard BO Methodology

Consider the optimization problem $\min_{x \in \mathcal{X}} g(x)$ where $g : \mathcal{X} \rightarrow \mathbb{R}$ is the objective function. In our settings, x is θ , \mathcal{X} is Θ and g is the distance metric $\theta \mapsto d(P_n^\theta, P_N^{real})$. The key idea of standard BO is to model $g(x), x \in \mathcal{X}$ as a realization of a GP. Note that this prior GP model

is uniquely defined by its mean structure $\mu(x) := E(g(x)), x \in \mathcal{X}$ and covariance structure $\Sigma(x, x') := cov(g(x), g(x')), x, x' \in \mathcal{X}$.

We fix the initial budget n_{init} and the total budget n_{total} in advance. The BO algorithm is initialized by randomly sampling n_{init} points $x_1, \dots, x_{n_{init}}$ in \mathcal{X} and evaluating their objective values $g(x_1), \dots, g(x_{n_{init}})$. At each iteration, suppose we have already evaluated the objective values at $x_1, \dots, x_{\tilde{n}}$. In the prior model, for any x , we know that $(g(x_1), \dots, g(x_{\tilde{n}}), g(x))$ follows a Gaussian distribution. Now conditional on the values of $g(x_1), \dots, g(x_{\tilde{n}})$, we get that $g(x) \sim N(\bar{\mu}(x), \bar{\sigma}^2(x))$ where the posterior mean and variance function $\bar{\mu}$ and $\bar{\sigma}^2$ can be explicitly computed. In particular, the posterior model $g(x)|(g(x_1), \dots, g(x_{\tilde{n}})), x \in \mathcal{X}$ is still a GP. If the total budget is reached, i.e. $\tilde{n} = n_{total}$, then we stop the algorithm and use the posterior GP model as a predictor for the objective function. Otherwise, with the posterior model, the next point $x_{\tilde{n}+1}$ is selected by optimizing the acquisition function, which can be interpreted as maximizing a reward function, and then $g(x_{\tilde{n}+1})$ is computed. In summary, we iteratively select the next point, evaluate the objective value and update the posterior model, until we reach the budget of evaluations.

A.2. GP Surrogate

We discuss the pros and cons of using GP as the surrogate model. GP is not only simple, but also possesses elegant probabilistic properties, and thus we can make statistical inference with the posterior model. Nevertheless, the workload of computing the covariance matrix grows cubically with the number of evaluations, and hence it cannot be scaled to large budget. Moreover, it is often questioned whether this simple model is capable of approximating the complicated objective function, especially if the objective function is rugged. Based on these limitations, surrogate models using various machine learning tools have been developed [10, 12]. However, we still decide to use GP as it fits our problem setting. On the one hand, the simulation itself takes so much time that we can only afford a limited total budget, with which computing the covariance matrix is not yet a concern. On the other hand, without the observation noise, the K-S statistic is relatively smooth in the input configuration θ , so GP is capable of approximating this particular objective function.

Appendix B. Experiment Details

B.1. Introduction to ABIDES Simulator

In this paper, we try to calibrate ABIDES, which simulates limit order book (LOB) exchange markets with various types of background agents, a NASDAQ-like exchange agent and a simulation kernel managing the flow of time as well as all the agent interactions. More specifically, we consider three types of basic trading agents here. Value agents are designed to simulate the fundamental traders who trade based on their judgment of the fundamental value of the asset. In the simulation system, the fundamental price is modeled as a discrete-time mean-reverting Ornstein-Uhlenbeck process [3]. Value agents arrive to the market following a Poisson process and trade a stock depending on whether it is cheap or expensive relative to their own noisy observation of the fundamental price. Besides, noise agents emulate the retail traders who trade on demand without other considerations. Each noise agent arrives to the market at a time uniformly distributed throughout the trading day and place an order of a random size in a random direction. Finally, market maker agents arrive with a constant rate and place limit orders on both sides of the LOB to provide liquidity.

While the simulator is able to output all the agent activities including bid prices and ask prices, such labeled data are typically unavailable. Thus, we extract minute mid price returns and minute traded volumes from the simulated limit order book, and then concatenate them as the output time series. We only keep the data from 10:00am to 4:00pm. Our goal is to calibrate the parameters of this simulation system only from the output time series sample set.

B.2. Experiment Setups

In this paper, we pick a “true” configuration and generate 1000 real time series samples using this configuration and different random seeds. Then for each candidate configuration in the parameter space, we generate 50 time series samples to compute the K-S statistic. The significance level α is chosen as 0.05. That is, the eligibility set is an approximate 95% confidence region.

To assess the calibration performance under different dimensions, we gradually increase the number of flexible parameters while fixing the other parameters to be same as the true configuration. In the experiments, we respectively consider calibrating two, six and eighteen parameters.

B.2.1. TWO-PARAMETER CALIBRATION EXAMPLE

To start with, we apply our framework to calibrate two parameters summarized in Table 1. In this low-dimensional example, we could use grid search as a baseline to investigate how the distance metric varies with the parameter values.

Table 1: Parameters to calibrate in two-parameter calibration example.

Parameter Symbol	Meaning	True Value	Range
num_noise	Number of Noise Agents	5000	[2500,7500]
lambda_a	Arrival rate of Value Agents	1e-12	[1e-16,2e-12]

To run the BO algorithm, we set the initial budget n_{init} as 10 and the total budget n_{total} as 100. We use the commonly-used RBF kernel as the GP kernel. White noise is added to the kernel considering that our distance metric is actually noisy. We use a mixture of expected improvement (EI), probability of improvement (PI) and lower confidence bound (LCB) as the acquisition function, which is the default choice of the skopt package.

B.2.2. SIX-PARAMETER CALIBRATION EXAMPLE

In the six-parameter calibration example, we try to calibrate the parameters described in Table 2. Note that for the last parameter, we take logarithm before scaling the parameter space to $[0, 1]^6$ since the output time series distribution is not sensitive to this parameter.

We compare the performance of three algorithms: random search, standard BO and TuRBO. The total budget is still fixed as 100. Random search is simply randomly sampling points in the parameter space and evaluating the objective values. We try to keep the settings in standard BO and TuRBO identical as much as possible to be fair. For both algorithms, we set the initial budget as 12, use the Matérn-5/2 kernel with white noise as the GP covariance kernel and choose Thompson sampling (TS) with 500 candidate points as the acquisition function. Specifically for TuRBO, the hyperparameters are chosen as: number of trust regions $m = 1$, batch size $q = 1$, $\tau_{succ} = 1$, $\tau_{fail} = 6$, $L_{min} = 2^{-7}$, $L_{max} = 1.6$, $L_{init} = 0.8$. We refer to [6] for more details about the

Table 2: Parameters to calibrate in six-parameter calibration example.

Parameter Symbol	Meaning	True Value	Range
num_value	Number of Value Agents	100	[0,200]
num_noise	Number of Noise Agents	5000	[2500,7500]
lambda_a	Arrival rate of Value Agents	1e-12	[1e-16,2e-12]
r_bar	Mean of Fundamental	1e5	[1e2,2e5]
kappa	Mean Reversion Rate of Fundamental	1.67e-12	[1e-16,3e-12]
fund_vol	Fundamental Volatility	1e-4	[1e-8,1]

hyperparameters and we note that our choice is mostly the same as their recommendation, except a little adjustment due to our much more limited budget of evaluations than theirs.

To reduce randomness of experiments, we test the three algorithms with 20 different random seeds. Under the same seed, the sampled initial points and the corresponding K-S statistic values are the same for each algorithm, so they “start off” with the same information about the objective function.

B.2.3. EIGHTEEN-PARAMETER CALIBRATION EXAMPLE

Now we consider an eighteen-parameter calibration problem to further justify the effectiveness of TuRBO. We introduce two additional types of value agents to the system and also add the order size upper and lower bounds for each type of value agent or noise agent as parameters. The parameters are documented in Table 3.

Table 3: Parameters to calibrate in eighteen-parameter calibration example.

Parameter Symbol	Meaning	True Value	Range
num_value_1	Number of Type 1 Value Agents	100	[0,200]
lambda_a_1	Arrival Rate of Type 1 Value Agents	1e-12	[1e-16,2e-12]
min_size_value_1	Minimum Order Size of Type 1 Value Agents	20	[6,34]
max_size_value_1	Maximum Order Size of Type 1 Value Agents	50	[36,64]
num_value_2	Number of Type 2 Value Agents	0	[0,50]
lambda_a_2	Arrival Rate of Type 2 Value Agents	1e-12	[1e-16,2e-12]
min_size_value_2	Minimum Order Size of Type 2 Value Agents	100	[76,124]
max_size_value_2	Maximum Order Size of Type 2 Value Agents	150	[126,174]
num_value_3	Number of Type 3 Value Agents	0	[0,50]
lambda_a_3	Arrival Rate of Type 3 Value Agents	1e-12	[1e-16,2e-12]
min_size_value_3	Minimum Order Size of Type 3 Value Agents	200	[176,224]
max_size_value_3	Maximum Order Size of Type 3 Value Agents	250	[226,274]
num_noise	Number of Noise Agents	5000	[2500,7500]
min_size_noise	Minimum Order Size of Noise Agents	20	[6,34]
max_size_noise	Maximum Order Size of Noise Agents	50	[36,64]
r_bar	Mean of Fundamental	1e5	[1e2,2e5]
kappa	Mean Reversion Rate of Fundamental	1.67e-12	[1e-16,3e-12]
fund_vol	Fundamental Volatility	1e-4	[1e-8,1]

Similar to the six-parameter experiment, we compare the performance of random search, standard BO and TuRBO. We raise the total budget to 300 and the initial budget for standard BO or TuRBO to 36. We still use the Matérn-5/2 kernel with white noise as the GP kernel. The acquisition function is still TS but with 3600 candidate points. The hyperparameters for TuRBO are number of trust regions $m = 1$, batch size $q = 1$, $\tau_{succ} = 1$, $\tau_{fail} = 18$, $L_{min} = 2^{-7}$, $L_{max} = 1.6$, $L_{init} = 0.8$. We repeat the experiment with 10 random seeds to reduce randomness.

Appendix C. Discussion about Feature Extractors

With the two-parameter calibration example, we investigate how to choose the feature extractor f . In the experiments, we adopt the simplest idea which is not to apply any feature extraction at all. We will verify the effectiveness of this plain K-S statistic and discuss about possible selections of feature extractors.

C.1. Effectiveness of Plain K-S Statistic with No Feature Extraction

We use two simple examples to visualize the K-S statistic and show the effectiveness. In the first example, we compare the true configuration (5000, $1e - 16$) with a fake configuration (3000, $1e - 16$). Figure 5(a)subfigure plots 10 time series samples for each configuration. We observe that the mid price return time series distributions look similar while the volume time series seem to come from different distributions. In fact, the fake configuration cannot be rejected by the K-S test applied only on the mid price return time series, but can be rejected only with the volume time series. In Figure 5(b)subfigure, we find the dimensions in the mid price return as well as the volume time series that maximize the one-dimensional K-S statistic and plot the discrepancy in the distributions over these dimensions. It is verified that the discrepancy in the volume distribution is more significant. In the second example, we select another fake configuration (7000, $1e - 16$). In this case, the volume time series look similar while the mid price return time series look obviously different (see Figure 5(c)subfigure). The K-S statistic also coincides with our observation in this example (see Figure 5(d)subfigure). Combining the two examples, we conclude that both price and volume time series are useful in calibration, and the K-S statistic is consistent with intuitive judgments.

More systematically, we compute the K-S statistic at some grid points in the parameter space and plot the heatmap in Figure 6(a)subfigure. Note that the true configuration achieves the minimum K-S statistic among all the grid points. We also mark the points in the eligibility set. We find that this approximate eligibility set is relatively small compared to the entire space and it looks like an ellipsoid around the truth. Generally, the K-S statistic looks smooth with respect to the parameter value, though there is indeed some randomness. Overall, we decide to directly use the K-S statistic without feature extraction in the experiments.

C.2. Conservativeness of K-S statistic with High-Dimensional Time Series

However, we note that the K-S statistic can be conservative especially when the length of time series is too large due to the nature of Bonferroni correction. Compared with the minute data, we try to increase the sampling frequency to 10 seconds and hence the length of the time series increases to over 4,000. We plot the heatmap in Figure 6(b)subfigure and we observe that the eligibility set is much larger. Some configurations could have been rejected with minute data but fail to be rejected

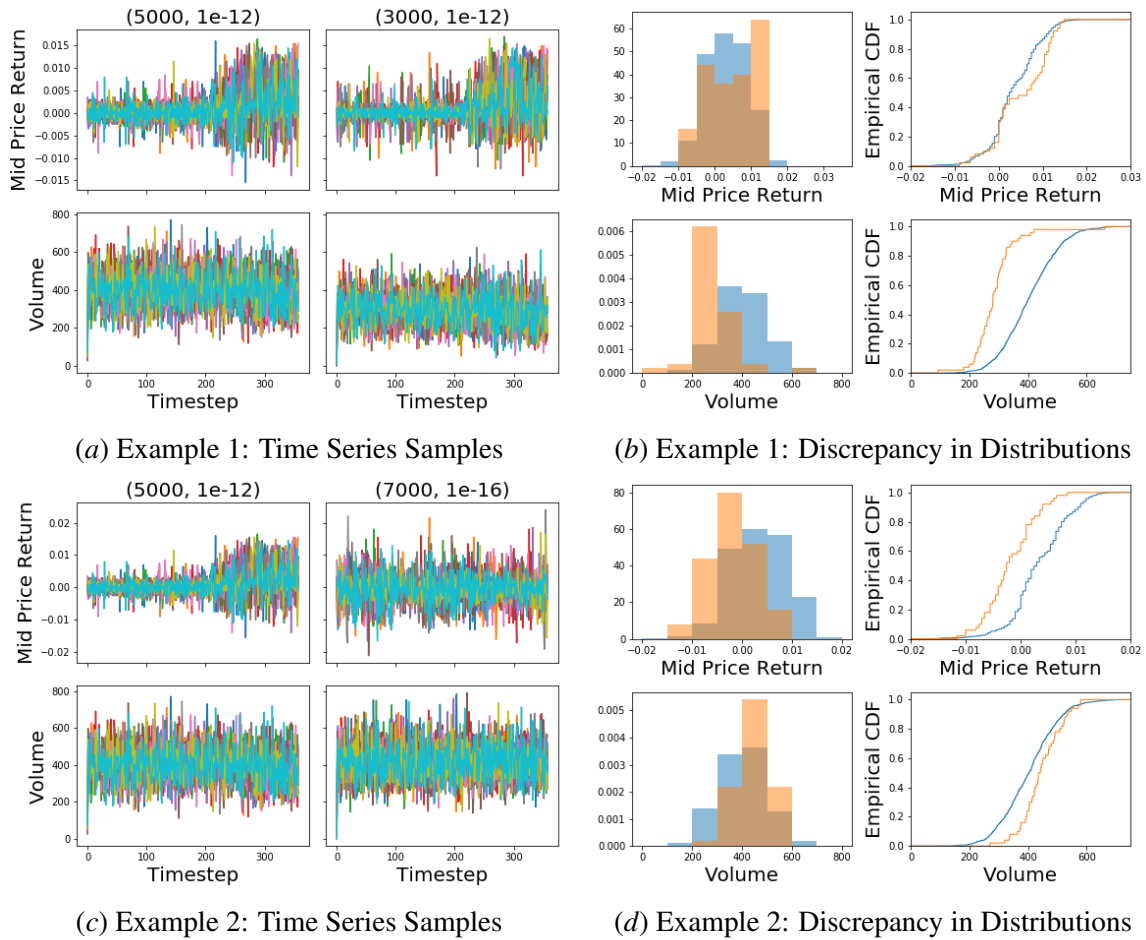


Figure 5: Visualization of the K-S statistic with two examples. (a) and (c) plot 10 time series samples under each configuration. (b) and (d) plot the discrepancy in the distributions.

with the 10 seconds data. In this case, applying a suitable feature extractor before computing the K-S statistic can be helpful.

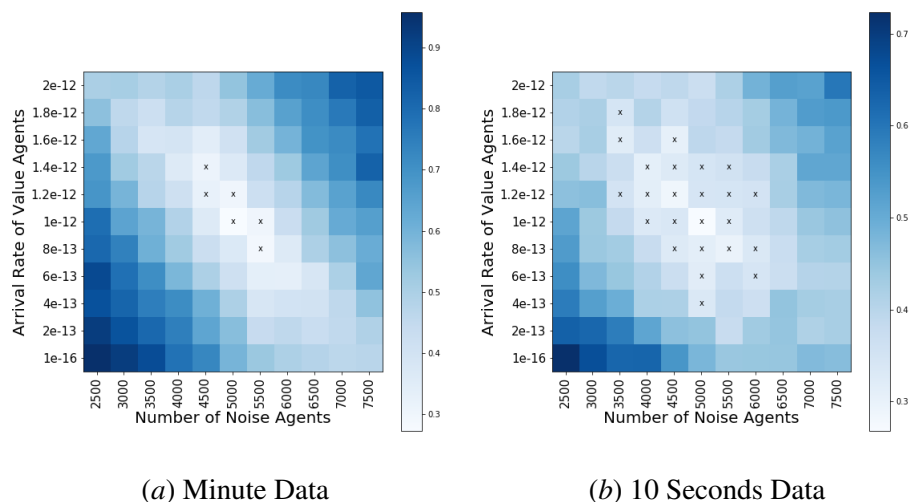


Figure 6: Heatmap of K-S statistic without feature extraction with respect to the number of noise agents and the arrival rate of value agents. The grid points with “x” marks are in the eligibility set.

C.3. Feature Extractors using Learning-Based Methods

Various learning-based feature extractors have been developed. We skip the detailed discussion here and refer to [2, 13] for feature extractors of high-dimensional financial time series using autoencoder, GAN or WGAN, which can effectively reduce the dimension and hence the conservativeness of the K-S statistic. However, while these machine learning models are flexible and powerful, they usually lack interpretability and also they need to be retrained whenever they are applied to a different problem setting.

C.4. Feature Extractors using Stylized Fact Metrics

Finally, we investigate whether we can use stylized fact metrics with physical meanings to construct the feature extractor. We consider four metrics discussed in [16]: autocorrelation, kurtosis, volatility clustering and volume volatility correlation. To be specific, we denote the minute log return as r_t and the volume as V_t . Then they are respectively the correlation between r_t and r_{t+1} , the kurtosis of r_t , the correlation between r_t^2 and r_{t+1}^2 , and the correlation coefficient of V_t against $|r_t|$. With each time series sample, we compute an estimate for each metric, and then use these estimates as the extracted features. Figure 7 show the heatmaps of the K-S statistic using each metric alone and also the corresponding eligible grid points. We see that for each of them, the K-S statistic seems not to vary with the parameter value regularly. If we combine the four metrics as a single feature extractor (i.e. $f : \mathbb{R}^O \rightarrow \mathbb{R}^4$), then as shown in Figure 7(e)subfigure, the eligibility set looks small, but there seem to be a lot of local minima, which will increase the difficulty of optimization. Based on our investigation, despite their usefulness in general, the stylized fact metrics may not be sufficient

to work as the feature extractor since too much information is lost when the entire time series is summarized into a few numbers.

Appendix D. Non-Identifiability

Figure 8 plots all the eligible points that we have found in our six-parameter calibration experiments. They form an approximation to the eligibility set, which roughly covers the true configuration. This plot suggests that the non-identifiability issue exists for our simulator. For instance, we observe that for the eligible points, the arrival rate of value agents is inversely proportional to the number of value agents. That is, as long as the total arrival rate is equal to the truth, we cannot really identify the number or the arrival rate of a single agent.

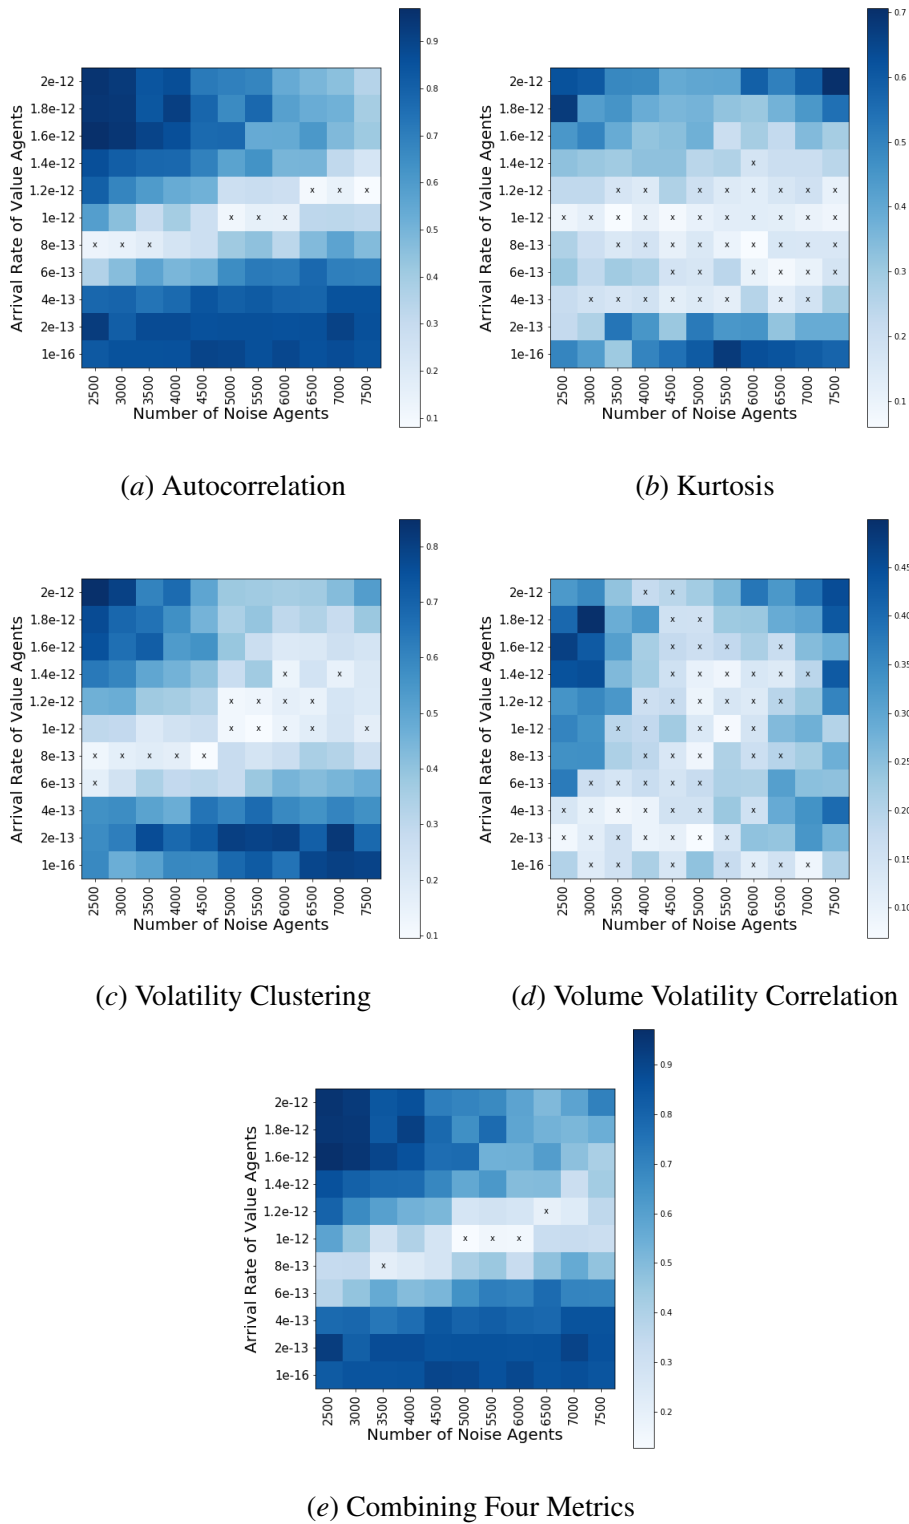


Figure 7: Heatmap of K-S statistic using stylized fact metrics as the feature extractor.

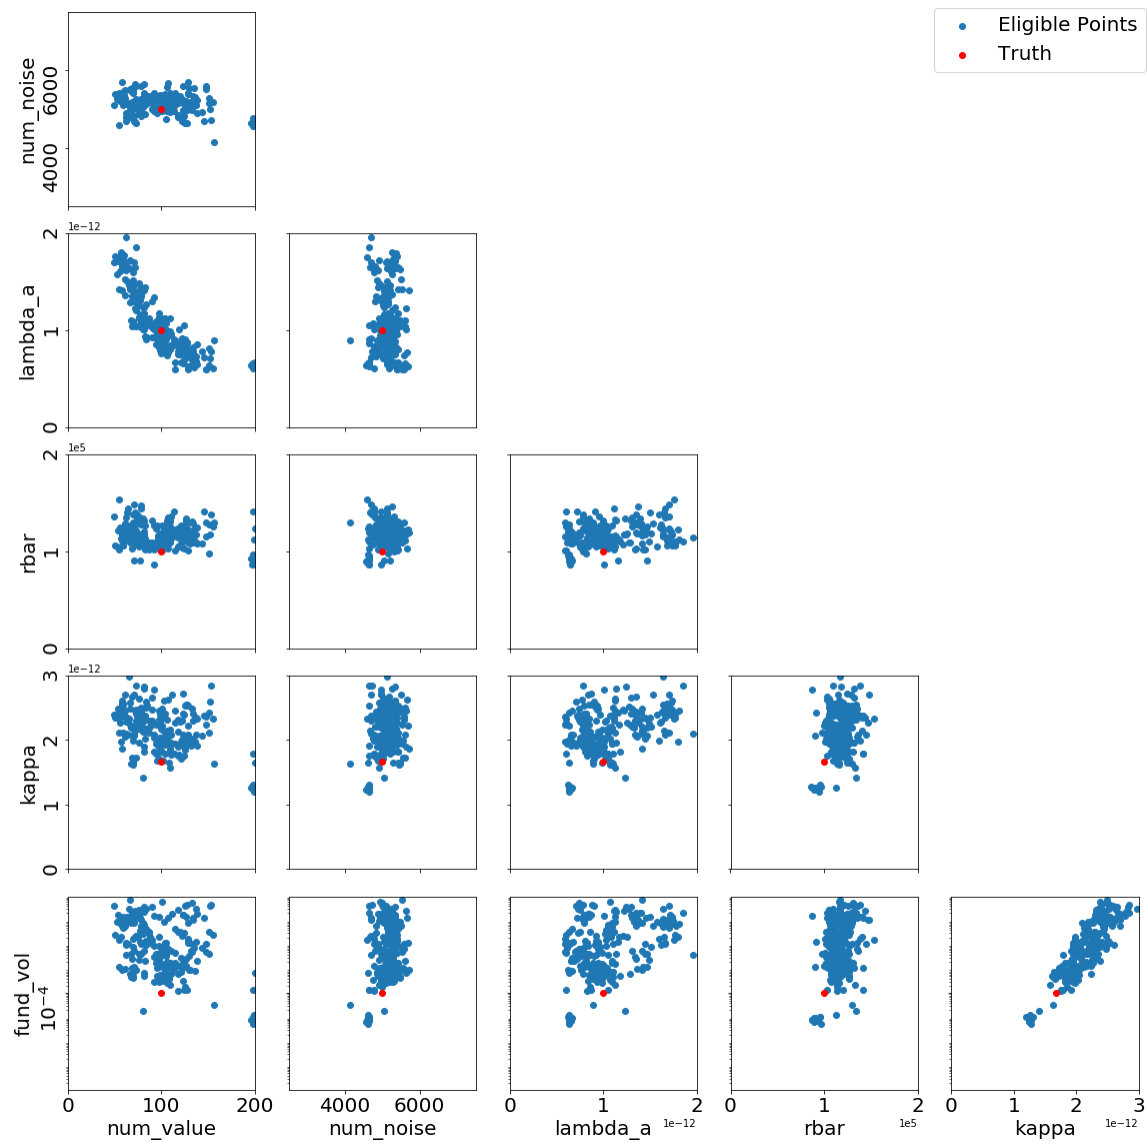


Figure 8: Pairplot of the eligible points found in the six-parameter experiments.