

Approximating Optimal Trading Strategies Under Parameter Uncertainty: A Monte Carlo Approach

Thomas Johnson

tjohnson2009@kellogg.northwestern.edu

June 1, 2009

1 Introduction

This paper considers the problem of a capital-limited investor with log utility who has the opportunity to invest in a security that follows a parametric price process. While the investor knows the form of the process, the exact parameter values are not known and must be inferred by observing the evolution of the security's price over time. The approach that will be described is applicable to any model and to single or synthetic securities. However, this paper will specifically consider a synthetic security that follows an Ornstein-Uhlenbeck process, $dS_t = \eta(\bar{x} - S_t)dt + \sigma dW_t$. The synthetic security will be formed by buying one asset and selling another.

The Ornstein-Uhlenbeck process was chosen for two reasons. First, it has real-world applicability, for example as a model for pair trading. Pair trading has been practiced in industry since at least 1985 (Pole, 2007) and the profitability of a pair-trading strategy has also been examined in the literature.

For example, (Gatev *et al.*, 2006) shows that a diversified pair-trading strategy is profitable and does not correlate with any well-known risk factor. Second, although this paper presents a Monte Carlo approach for finding an optimal trading strategy, there is a significant literature that focuses on analytic solutions to optimal trading strategy problems including the optimal strategy for trading an Ornstein-Uhlenbeck process.

2 The Kelly Criterion as Applied to Ornstein-Uhlenbeck Processes

The optimal strategy for investors with log utility facing an investment opportunity that has fixed, known odds and a discrete set of payoffs was presented in (Kelly Jr, 1956) and has since become known as the Kelly Criterion. Kelly originally investigated this problem in the context of the mathematical field of information theory, where he framed the problem in terms of a gambler who has access to occasionally incorrect early information about the outcomes of a horse race while he can still place bets on that race. A restatement of the problem in an investing framework would go as follows: Suppose that you are an investor with an investment strategy that outperforms the market on a risk-adjusted basis. For instance, you may have access to nonpublic information, a structural advantage in your access to the market, or merely a superior way of finding the relative value of stocks. However the strategy is not risk-free, and a certain percentage of the time you may lose money even though you will outperform the market in the long run. Assuming that you know (e.g., from historical analysis) how often your methodology generates positive returns, how much should you invest in a particular opportunity? Clearly, you should not invest all of your capital even if you have a greater than 50% chance of being

correct in your valuation, or indeed even if you have a 99% chance of being correct. Although investing all of your capital maximizes your expected wealth after the investment thesis plays out, it is also possible that this strategy will bankrupt you.

The Kelly Criterion says that rather than maximizing your expected wealth, you should maximize the expected growth rate of your wealth; that is, instead of maximizing the expected value of your portfolio at some future point in time you should instead maximize your portfolio's return over that time period. In order to do this, the Kelly Criterion states that the investor should invest a fraction of his wealth equal to $f^* = \frac{bp-q}{b}$ where b is the odds received, p is the probability of winning, and $q = (1-p)$ is the probability of losing. For instance, if we believe that our investment thesis has a 75% chance of being correct, and if we are correct then our investment will yield 10% while if we are incorrect we lose all of our capital, then we should invest $f^* = \frac{1.1*0.75-0.25}{1.1} \approx 52.3\%$ of our wealth. The Kelly Criterion has the attractive property of increasing an investor's wealth faster than any other strategy almost surely over the long run. It provides a risk-management tool for an investor who wants to maximize his long-term wealth while ensuring that he does not go bankrupt. Indeed, in a sense the recommended fraction of total wealth to invest can be thought of as one measure of the riskiness of an investment given the investment's return. However, it is important to recognize that while the Kelly criterion eliminates the risk of bankruptcy, it does not minimize volatility. For opportunities that provide very high average returns but which pay off only infrequently, drawdowns of arbitrary size are possible. The investor can reduce volatility by betting only fraction of f^* , and this will maximize the growth rate of the investor's wealth for that level of volatility. Thus the investment amount given by f^* should be considered an upper bound, and Kelly showed that if you invest more than this fraction you

are guaranteed to go bankrupt almost surely in the long run.

While the original Kelly formula is most useful for analyzing investment opportunities with discrete payoffs, many opportunities have an essentially continuous range of possible payoffs with a long-term average return. Investment strategies are also typically thought of as having a volatility rather than a specific probability of being correct for individual investments. The formula for extending the Kelly criterion to an investment having an average return μ and standard deviation σ with risk-free rate r was analyzed in (Thorp, 2006):

$$f^* = \frac{\mu - r}{\sigma^2} \tag{1}$$

It is important to note that there is a key practical difference between the continuous and discrete case when using the Kelly formula. The investor in the discrete case is guaranteed to avoid bankruptcy. In the continuous case however, unless the investor adjusts his investment continuously and instantaneously he is not guaranteed to avoid bankruptcy, or indeed even to avoid negative wealth, cf. (Hausmann and Sass, 2004).

The Kelly Criterion allows an investor to be completely myopic when making investment decisions, in that he does not have to consider future or past investment opportunities or results (Hakansson, 1971). But when the expected return changes over time, as it does in the Ornstein-Uhlenbeck process, the optimal fraction of wealth to invest will also change. For an Ornstein-Uhlenbeck process with known parameters the optimal fraction is (Lv *et al.*, 2009):

$$f_t^* = \frac{\eta(\bar{x} - \log(S_t)) + \frac{1}{2}\sigma^2 - r}{\sigma^2} \tag{2}$$

This analytical solution provides an optimal benchmark that can be used to evaluate the performance of the Monte Carlo approach.

3 Partial Information and Parameter Risk

Although (2) provides the optimal investment fraction when the investor knows the exact value of the Ornstein-Uhlenbeck parameters, in the real world parameter values may vary over time and must be inferred from price paths or fundamental relationships. Optimal investment under partial information (i.e., when only the security's price is observable) has been studied before. For example, (Lakner, 1995) considers the optimal investment problem under partial information when the asset price follows Brownian motion with drift and the drift itself (rather than the asset price as in the Ornstein-Uhlenbeck process) may be mean reverting. (Hahn *et al.*, 2007) examines asset allocation under partial information when the assets have finitely many possible rates of return and volatility varies over time. (Mudchanatongsuk *et al.*, 2008) studies the optimal strategy for trading a cointegrated pair under partial information by modeling the two assets explicitly and using stochastic control techniques to find an optimal trading strategy for an investor with power utility. This is very similar to the problem that this paper considers, however (Mudchanatongsuk *et al.*, 2008) uses maximum likelihood estimation to obtain parameter values. Maximum likelihood estimation provides point estimates for parameter values, and does not correctly account for an investor's aversion to parameter risk. This paper will show that if an investor uses the point estimates provided by maximum likelihood techniques, the estimated optimal investment fraction will often be unrealistically large, which leads to a severe negative impact on the investor's wealth.

To explicitly incorporate parameter uncertainty, particle filtering is used to approximate probability distributions for each parameter. Particle filters have been used successfully in computational finance before, where they are commonly referred to as Sequential Monte Carlo techniques, for example see

(Johannes *et al.*, 2002; Andersen *et al.*, 2008). A brief description of particle filtering is presented here; thorough tutorials are available in (Andersen *et al.*, 2008) and (Doucet and Johansen, 2009).

Particle filtering is a Monte Carlo-based method for estimating the values of hidden parameters. While particle filters are conceptually similar to Kalman filters in the way they are used, they have the advantage of being able to handle nonlinear and non-Gaussian models. In addition, particle filters may provide more accurate estimates than even unscented or extended Kalman filters (Daum and Co, 2005; Kim and Iltis, 2002).

Particle filters use a set of point estimates to approximate the joint probability distribution of parameter values. Each point estimate is called a “particle,” and consists of a likelihood value and a parameter vector, in this case $(\hat{x}, \hat{\eta}, \hat{\sigma})$. The particles are updated each time a new observation is received by the filter. The update can be thought as a two-step process. In the first step, a new prior value for the particle’s parameter vector is calculated based only on the current parameter values. While this would normally involve updating the parameter values based on a system equation, in this case we are trying to estimate fixed parameter values that do not change over time. Although the true parameter vector does not change, skipping the movement step would rapidly lead to all of the probability weight being placed on the single most likely particle. While some sophisticated approaches such as (Doucet and Tadić, 2003; Johansen *et al.*, 2008) have been proposed to deal with this problem, in this project each parameter value simply has low-variance Gaussian noise added to it.

In the second step, a new posterior likelihood is calculated for the particle given the new system observation and updated parameter vector. This likelihood is multiplied by the particle’s previous likelihood so that a single outlier observation does not invalidate all of the filter’s previous estimates. Tradition-

ally, this step requires an analytical likelihood function. As mentioned in the introduction, the Ornstein-Uhlenbeck process was chosen in part because it is easy to derive such functions. However, even if an exact likelihood function is not available, one can be approximated using approximate Bayesian computation techniques, for example as described in (Toni *et al.*, 2009), although of course this increases the computational requirements of the algorithm.

Since the time $t + 1$ value of an Ornstein-Uhlenbeck process is normally distributed, the likelihood function of an Ornstein-Uhlenbeck process is simply an adaptation of the likelihood function of a normal distribution. The adapted likelihood is

$$L(x, \mu, \hat{\eta}, \hat{\sigma}) = \frac{1}{\hat{\sigma}\sqrt{2\pi}} \exp\left(-\frac{(x - \mu)^2}{2\hat{\sigma}^2}\right) \quad (3)$$

with $x = \log(S_t)$, $\mu = \log(S_{t-1})e^{-\delta\hat{\eta}} + \hat{x}(1 - e^{-\delta\hat{\eta}})$, $\sigma = \hat{\sigma}\sqrt{\frac{1 - e^{-2\delta\hat{\eta}}}{2\hat{\eta}}}$, and δ is the amount of time that has passed between the previous observation and the current observation.

Even if the particles have Gaussian noise added at each step, particle diversity may suffer as most of the probability density becomes concentrated in a few particles. To mitigate this effect, particle filters use a technique called resampling that removes the lowest-probability particles and probabilistically clones high-probability particles. Resampling occurs only when the effective sample size, $ESS = (\sum_{i=1}^N (W^i)^2)^{-1}$ where W_i is the normalized weight of particle i , drops below a certain threshold. For this study, the threshold $ESS \leq \frac{N}{2}$ was used to trigger resampling.

Once a parameter probability distribution has been estimated using the particle filter, a series of Monte Carlo simulations can be used to estimate an optimal investment fraction. To do this, the algorithm performs repeated simulations to estimate the mean and variance of the returns. For each simulation,

a particle is first selected from the particle filter. The probability of selecting a particular particle is proportional to its weight. Once the particle is selected, an Ornstein-Uhlenbeck process is created with parameter values $(\hat{x}, \hat{\eta}, \hat{\sigma})$ equal to the values of the particle's parameter vector, and S_t equal to the value of the latest observation. The process is simulated for δ units of time, and after performing repeated simulations the return mean and variance are estimated. With these estimates, \hat{f}^* can be estimated using (1).

It should be noted that the parameter and optimal investment fraction estimation procedures described above are not limited to Ornstein-Uhlenbeck processes, and in fact are completely general and can be applied to any process that can be simulated in a Monte Carlo manner.

4 Simulated Data Experiments

4.1 Experimental Design

To evaluate the quality of the Monte Carlo technique described above, the technique was compared against three other strategies. The first was an optimal strategy that had perfect knowledge of the Ornstein-Uhlenbeck parameters and calculated the fraction to invest at each time step using (2). The second strategy also calculated the fraction to invest using (2), but used maximum-likelihood estimation to determine the parameter estimates. The third strategy was identical to the second, but was prohibited from trading for ten time steps in order to allow the maximum-likelihood estimates to start to converge. Each strategy traded a security whose log price followed an Ornstein-Uhlenbeck process. The strategies were run 1000 times, and each run consisted of 100 time steps. The parameters of the Ornstein-Uhlenbeck process were arbitrarily chosen to be $\bar{x} = 0$, $\eta = 1$, $\sigma = 0.25$, $\delta = 0.1$. The relative performance of the strategies

are robust to variation in parameter values. If the strategy’s wealth during a particular simulation became negative, the strategy was given no further opportunity to trade during that simulation. The particle filter was implemented using (Johansen, 2009). The filter used 1000 particles, and 100,000 Monte Carlo samples each timestep to approximate the optimal fraction of wealth to invest. All initial particle parameter estimates were drawn from a Normal(0,1) distribution, with the exception that $\hat{\eta}$ and $\hat{\sigma}$ were constrained to be greater than 0.001. The full source code for this experiment is available upon request.

4.2 Results

Statistics for single-period returns are shown in Table 1 and statistics for total returns are shown in Table 2. Sharpe ratios are not calculated due to the artificial nature of the timescale and process being simulated. Average single-period returns are not reported for the MLE strategy since the high number of bankruptcies biases the single-period returns due to the fact all single-period returns after a bankruptcy are zero for the remainder of the run. The nearly immediate deterioration of the wealth of the MLE trader is shown in Figure 4. The terrible results of the MLE strategy are due to two factors. First, the MLE strategy’s initial parameter estimates are highly variable. This often leads to the strategy making bets that are more than millions of times its bankroll during early periods, which leads to almost immediate bankruptcy. When bankruptcy does not occur immediately, the maximum-likelihood estimator’s bias towards overestimating the mean-reversion speed (Andersen, 2009) leads to consistent overbetting. This overbetting is devastating to the strategy’s returns. An example can be seen in Figure 5, which plots the mean wealth of the maximum-likelihood trader that was prohibited from trading for the first ten time steps. The outsized gains of the strategy are offset by even larger losses; the break in

the plot comes from an especially large loss that occurred shortly after $t = 0.6$. (Boguslavsky and Boguslavskaya, 2004) investigates the performance of a MLE-based trader on Ornstein-Uhlenbeck processes and provides a chart showing the effect of overestimating the mean-reversion speed. The chart is reproduced in Figure 1. The details of the simulation are slightly different, but the negative impact of overestimating the mean-reversion speed for an Ornstein-Uhlenbeck trader is clear.

In contrast, the smooth and consistent gains of the optimal trader shown in Figure 2. The mean wealth path of the Monte Carlo trader shown in Figure 3 is also relatively smooth, but of course does not grow as rapidly since the Monte Carlo trader must attempt to learn the parameters over time. Figures 6, 7, and 8 show the mean parameter estimates of the Monte Carlo trader's particle filter over time with error bars indicating ± 1 standard deviation of the particle filter's estimates. It is worth clarifying that these the error bars do not show the variability of a point estimate over all runs, but rather the variability estimated by the particle filter during the course of a run averaged over all runs. That is, these plots show how uncertain the Monte Carlo trader is about its own estimates of the parameters. The speed at which the parameter estimates converge are interesting. The estimates for $\hat{\sigma}$ converge almost immediately to the proper value, since $\hat{\sigma}$ can be estimated independently of \hat{x} and $\hat{\eta}$. The estimates of \hat{x} and $\hat{\eta}$ are not independent, however. In general, a good estimate of $\hat{\eta}$ requires a good estimate of \hat{x} . In fact, the variance of $\hat{\eta}$ first decreases as the filter begins to converge prematurely, then increases while the filter attempts to learn \hat{x} more accurately.

One of the key goals of this project was to determine whether the parameter uncertainty risk implied by the particle filter could be integrated into a trading strategy. To investigate this, the actual invested fractions of both the Monte

Carlo and maximum-likelihood trader as compared to the true optimal fraction were calculated, i.e. $\frac{\hat{f}_t^*}{f_t^*}$. Figure 9 shows the median fraction invested by the Monte Carlo strategy as a proportion of the optimal fraction. The Monte Carlo trader initially bets a very small portion of its capital in the wrong direction due to early mis-estimates of \hat{x} . However, as the particle filter’s parameter estimates begin to converge, the Monte Carlo trading strategy becomes “more confident” and begins to increase its investment fraction. Despite the filter eventually converging on relatively accurate estimates for all of the parameters however, the Monte Carlo strategy never increases its fraction much above 0.5. This gap can be interpreted as the strategy discounting its estimate of $\frac{\hat{\mu}}{\hat{\sigma}^2}$ due to continued parameter uncertainty. Betting only a portion of the estimated Kelly investment is known as “fractional Kelly” investing and is widely practiced in the sports wagering community. (Thorp, 2006) notes that “because ‘over-betting’ is much more harmful than underbetting, ‘fractional Kelly’ is prudent to the extent the results of the Kelly calculations reflect uncertainties.”

In contrast, Figure 10 shows the median values of $\frac{\hat{f}_t^*}{f_t^*}$ for the maximum-likelihood trading strategy. While it initially may seem that the MLE strategy should generate better returns since it invests closer to the optimal fraction, these higher returns are more than offset by the frequent bankruptcy that the strategy encounters when it overbets. As Table 2 shows, the MLE strategy goes bankrupt over 50% of the time even when it is allowed to improve its estimates for ten periods before trading.

5 Real Data Experiment

5.1 Experimental Design

While the Monte Carlo trading strategy generally outperformed the maximum-likelihood strategy on simulated data, it is not always the case that simulated results translate into the real-world. To test the quality of the Monte Carlo trader on real-world data, the trader was run on three pairs of cointegrated equities. The equities pairs each consisted of two different classes of shares from the same company (see Table 3). For ease of implementation and to match the setup of the simulated data experiments, rather than trading the spread directly a synthetic security was constructed for each pair whose price was $P_t = \exp(\frac{P_t^A}{P_t^B})$. The pricing data for each stock was taken from the daily closing prices reported in the CRSP database. Plots of the price of each synthetic security are shown in Figures 11, 12, and 13. Clearly, the synthetic securities exhibit more complicated dynamics than the simple Ornstein-Uhlenbeck process that was used in the simulated data experiments. However, the parameters and algorithms were left exactly the same for the real data experiments, with the exception that only 10 runs were performed on each pair instead of 1000 due to time constraints.

5.2 Results

On real data, the Monte Carlo trading strategy unequivocally outperformed the maximum-likelihood strategy, even when the MLE strategy was given up to a full year to calibrate before trading. Since the maximum-likelihood strategy is deterministic, it is not possible to perform repeated runs on a static dataset. However, various calibration times were tried, and in every case the maximum-likelihood strategy went bankrupt within one year of trading. Just as in the

simulated data experiments, the maximum-likelihood trader overinvested heavily, racking up outsized gains for a few periods and then going bankrupt.

In contrast, the Monte Carlo trader was much more conservative in its investing, even though the parameter estimates converged quite quickly. The Monte Carlo trader tended to estimate both significantly higher values for $\hat{\sigma}$ and lower values for $\hat{\mu}$ when calculating \hat{f}^* compared to the maximum-likelihood trader. The difference in $\hat{\mu}$ values was driven largely by the maximum-likelihood's high $\hat{\eta}$ estimates, which were sometimes as much as ten times higher than comparable particle filter estimates by the Monte Carlo trader.

The trading results for the Monte Carlo strategy are reported in Table 4. The mean wealth paths for the strategy are shown in Figures 14, 15 and 16 along with the path of $\frac{P_t^A}{P_t^B}$. Because the results are averaged over a relatively low number of runs, there are some discontinuities in the graphs due to a particularly good or bad results at certain time periods during a single run. In general, the strategy performed extremely well, achieving extremely high Sharpe ratios as well as rapid capital growth. Only when trading Berkshire Hathaway did the strategy lose money on average, possibly because of the extremely tight bounds that Berkshire trades within and the needs to learn a very high value for η . Even in this worst case, the strategy lost only 1.14% of its wealth over 11 years.

6 Conclusion

This study investigated the performance of a novel Monte Carlo-based system for approximating optimal trading strategies under parameter uncertainty. The system uses particle filters to explicitly model parameter uncertainty, and combines both this parameter uncertainty and price path uncertainty when estimating optimal investment levels for a log utility investor. This strategy outperforms a strategy based on maximum-likelihood parameter estimation when tested on

both simulated and real data by avoiding the error of overleveraging its investments. Because overleveraging leads to almost certain ruin for an investor that follows the Kelly criterion, outperformance through lower average investments is a valid method for growing wealth while avoiding bankruptcy, and does not merely reflect a lower risk/reward preference.

The Monte Carlo system described is extremely general in nature, and does not require that the underlying process be modeled as an Ornstein-Uhlenbeck process. At most a likelihood function for the model parameters needs to be provided, and even this may be avoided by using Approximate Bayesian Computation techniques. Importantly for practitioners, the system is highly parallelizable and can be easily modified to take advantage of modern multi-processor computers or computing clusters.

There are many opportunities for future investigations. The particle filter could be improved by using modified particle methods that are more focused on parameter estimation as mentioned on page 6. A more complex model involving price jumps or additional parameters should be evaluated. For instance, a modified Ornstein-Uhlenbeck process that allows jumps in \bar{x} might be used to model credit spreads. This kind of model could be tested by trading the TED spread contract on the CME. Although trading costs were not part of this study, transaction costs could also be integrated into the trading model by prohibiting trades when the expected next-period gain is larger than the trading cost. In summary, the preliminary results presented using this model appear promising and the model provides a rich framework for future research.

Acknowledgments

I would like to thank Prof. Torben Andersen, Prof. Deborah Lucas, Prof. Robert McDonald, and Prof. Dimitris Papanikolaou for their invaluable help

and advice throughout this project. Even though I did not participate in a formal independent study with any of these professors this quarter, none of them hesitated to generously give their time whenever I had a question. Any errors are of course my own.

	μ	σ
Optimal Strategy	0.0446	0.2167
Monte Carlo Strategy	0.0153	0.1663

Table 1: Period-over-period return statistics

	μ	σ	Percentage of bankruptcies
Optimal strategy	16.3103	17.3753	2.8%
Monte Carlo Strategy	2.1603	8.6984	5.3%
MLE Strategy	-8.3108×10^{19}	1.738×10^{21}	76.9%
MLE, 10-Period Delay	-27.6486	787.3313	53.3%

Table 2: Total return statistics

Company	Ticker A	Ticker B	Start Date	End Date
Berkshire Hathaway	BRK.B	BRK.A	Jan 1, 1997	Sep 29, 2008
Liberty Global	LBTYA	LTYKV	Jan 1, 2006	Sep 29, 2008
News Corp	NWS	NWSA	Jan 1, 2005	Sep 29, 2008

Table 3: Equity Pairs Tested

	μ	σ	Sharpe Ratio
Liberty Global	0.6976	0.0681	10.2429
News Corp	0.3802	0.0463	8.2117
Berkshire	-0.0114	0.0021	-5.5498

Table 4: Total return statistics for real data experiments

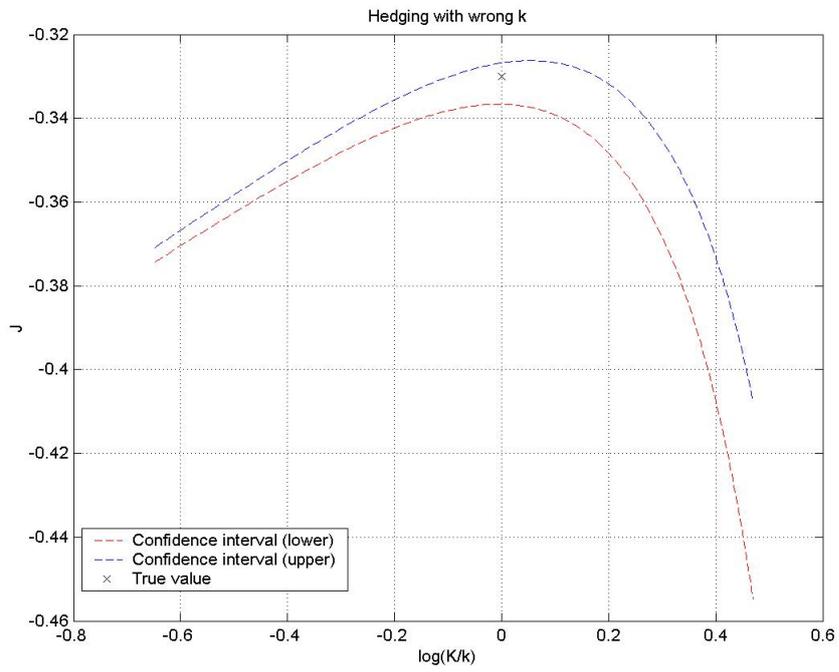


Figure 1: Effect of incorrect estimation of mean-reversion speed. K is the estimated mean-reversion speed; k is the true mean-reversion speed, and J is the log of terminal wealth. Reproduced from (Boguslavsky and Boguslavskaya, 2004), see that paper for details.

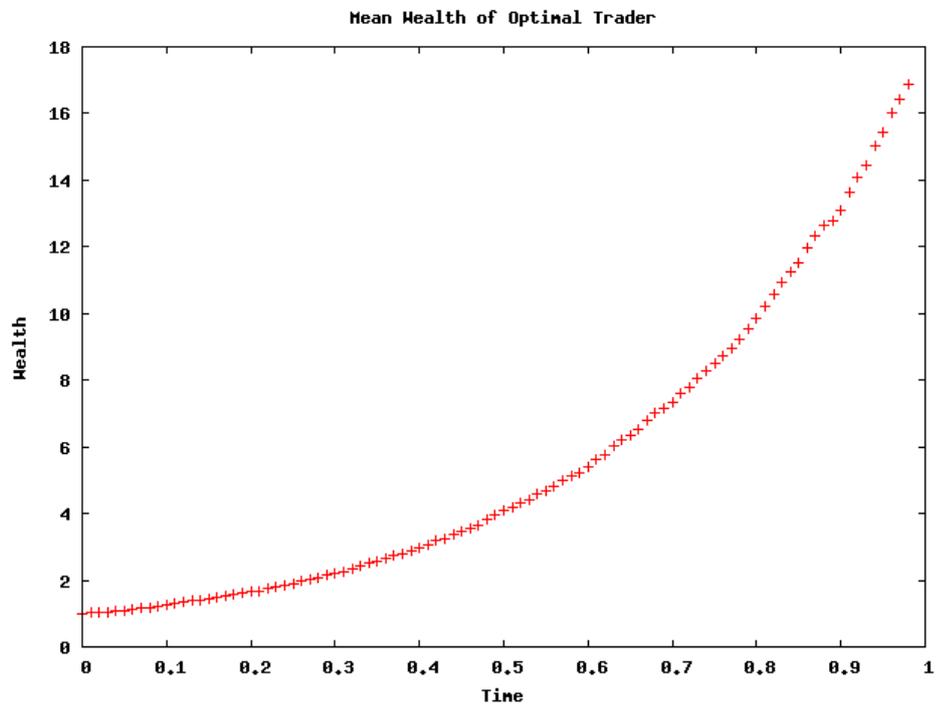


Figure 2: Mean Wealth of Optimal Trader

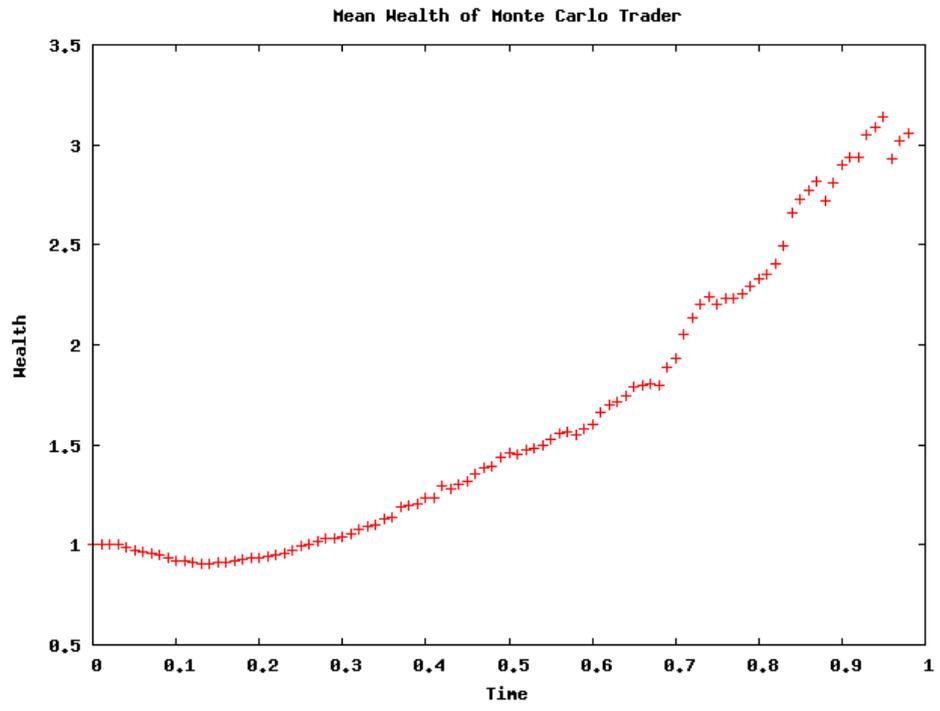


Figure 3: Mean Wealth of Monte Carlo Trader

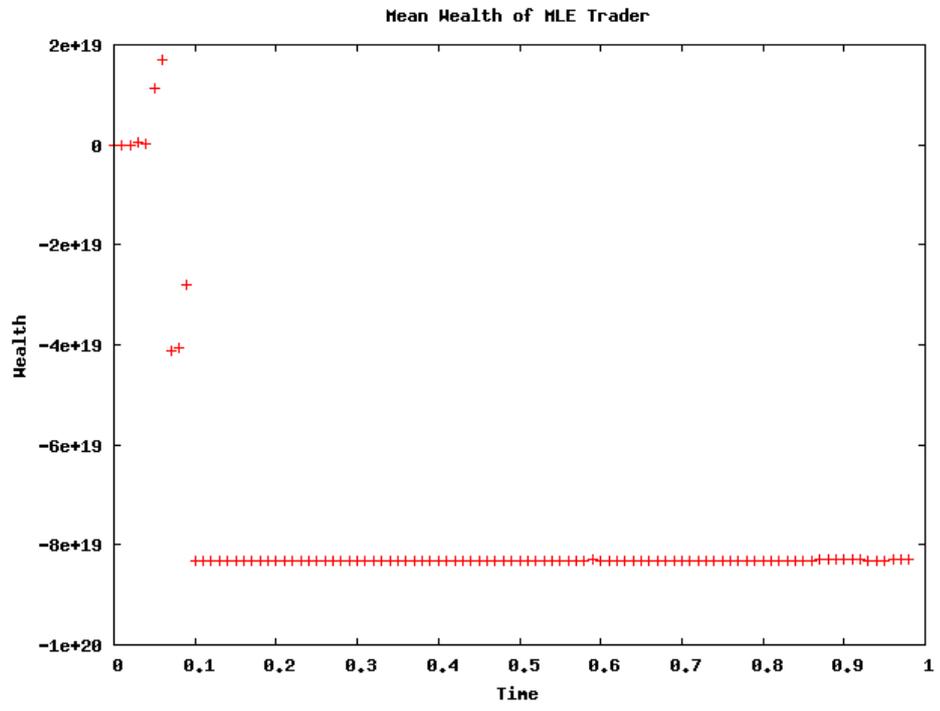


Figure 4: Mean Wealth of MLE Trader

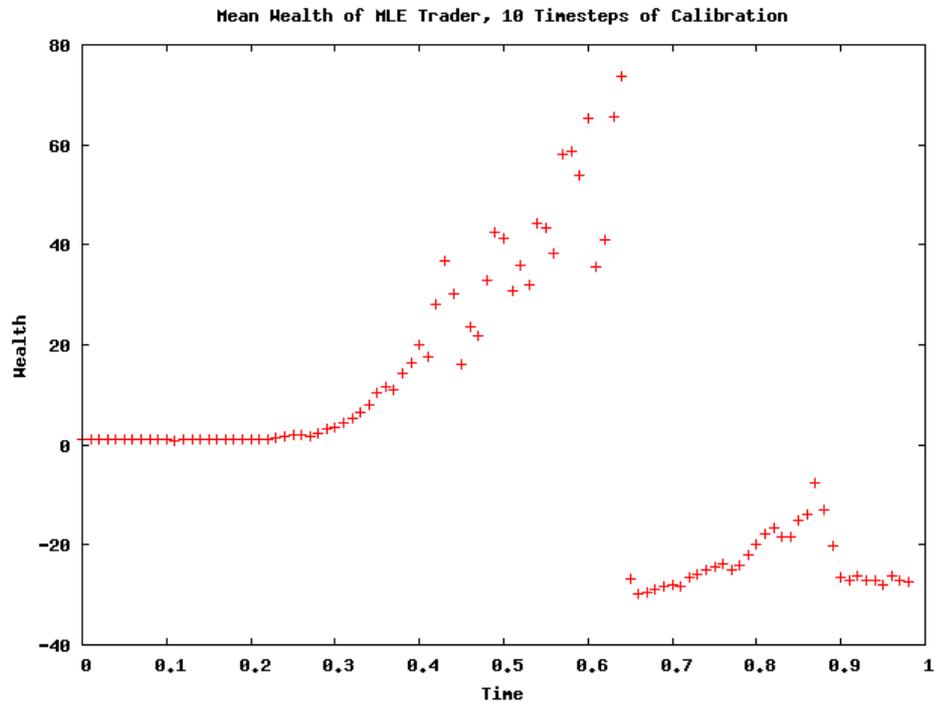


Figure 5: Mean wealth of MLE trader with 10 timesteps of calibration

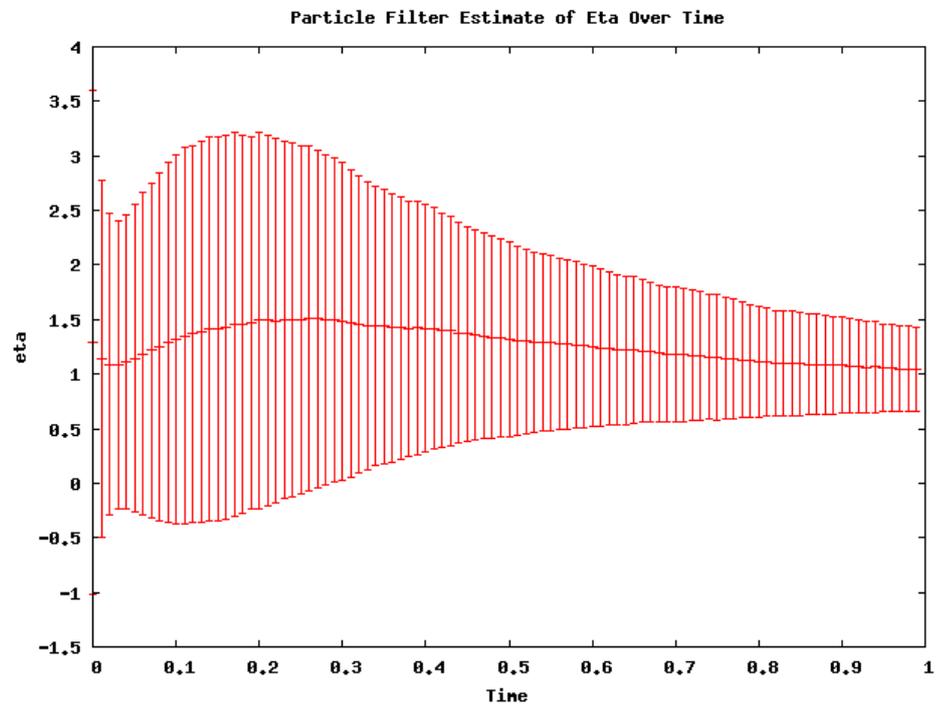


Figure 6: Particle Filter Estimate of $\hat{\eta}$ Over Time

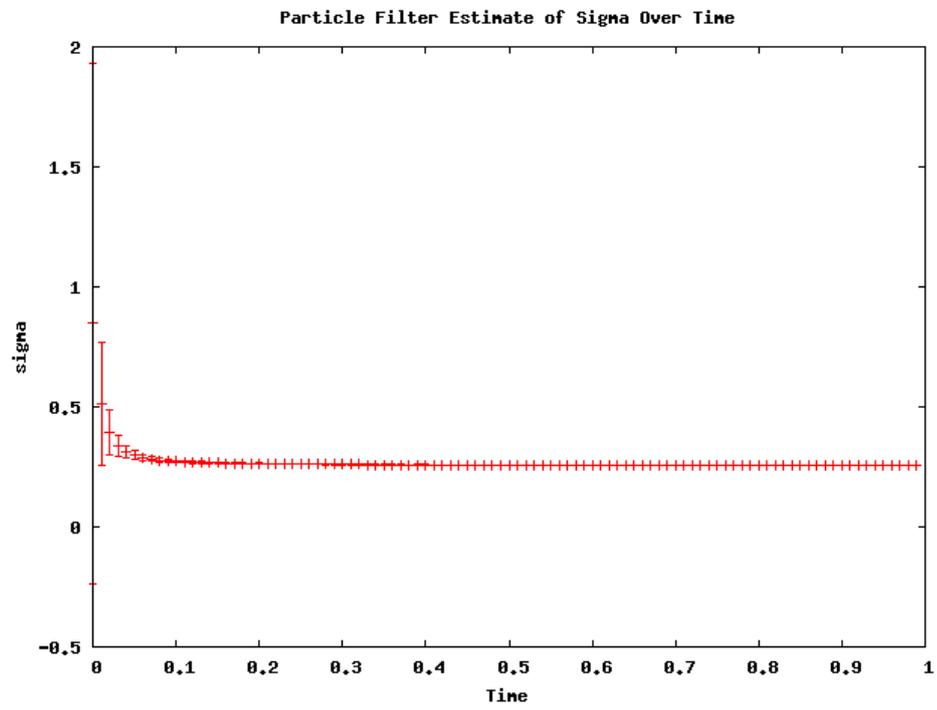


Figure 7: Particle Filter Estimate of $\hat{\sigma}$ Over Time

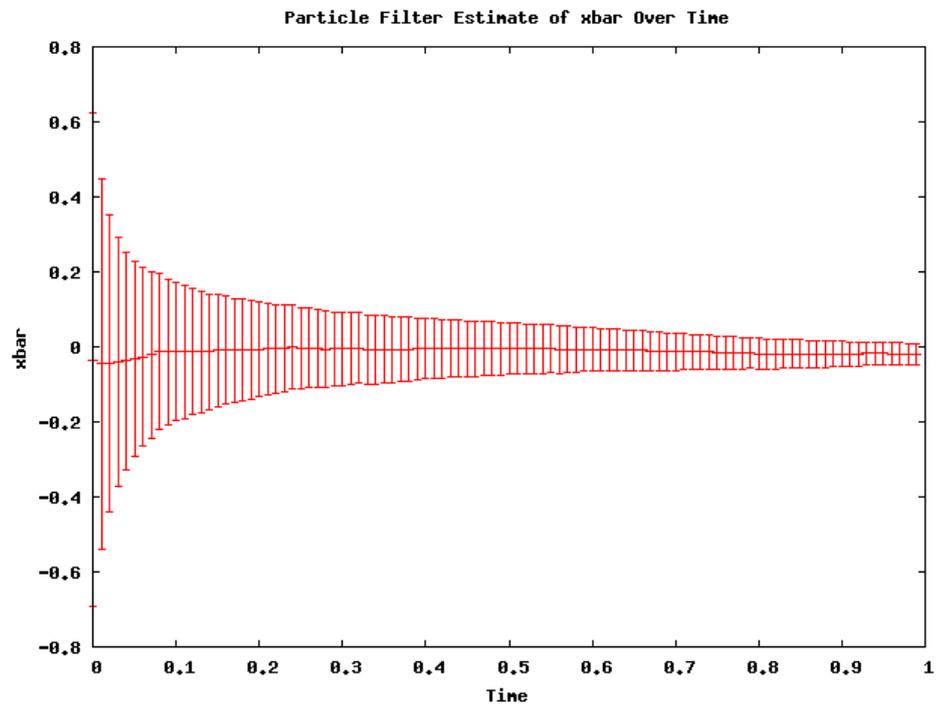


Figure 8: Particle Filter Estimate of \hat{x} Over Time

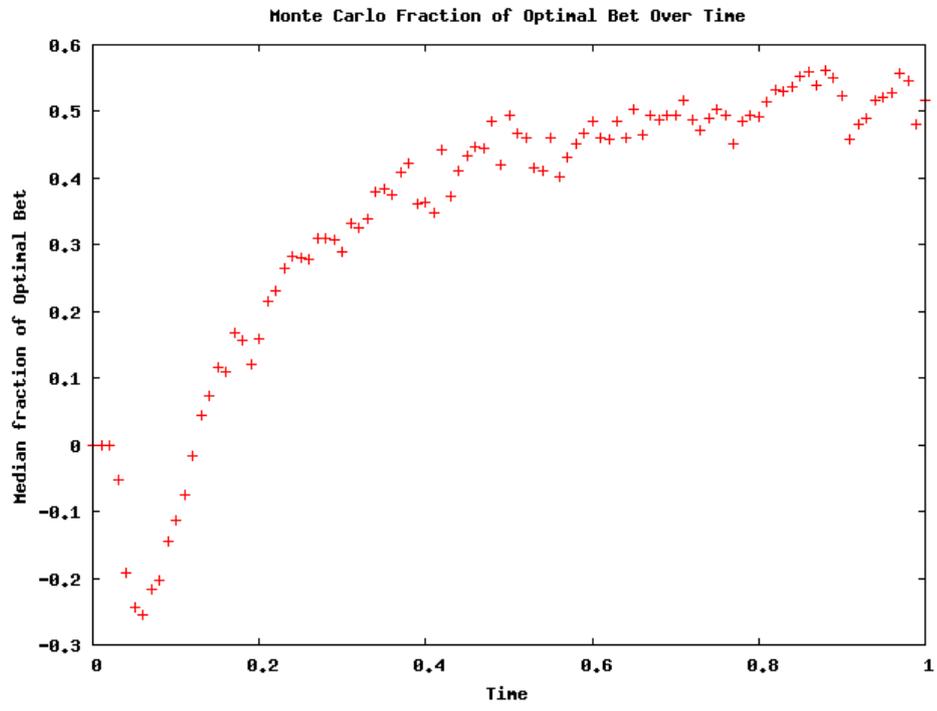


Figure 9: Median Monte Carlo Trader Bet as Fraction of Optimal Bet

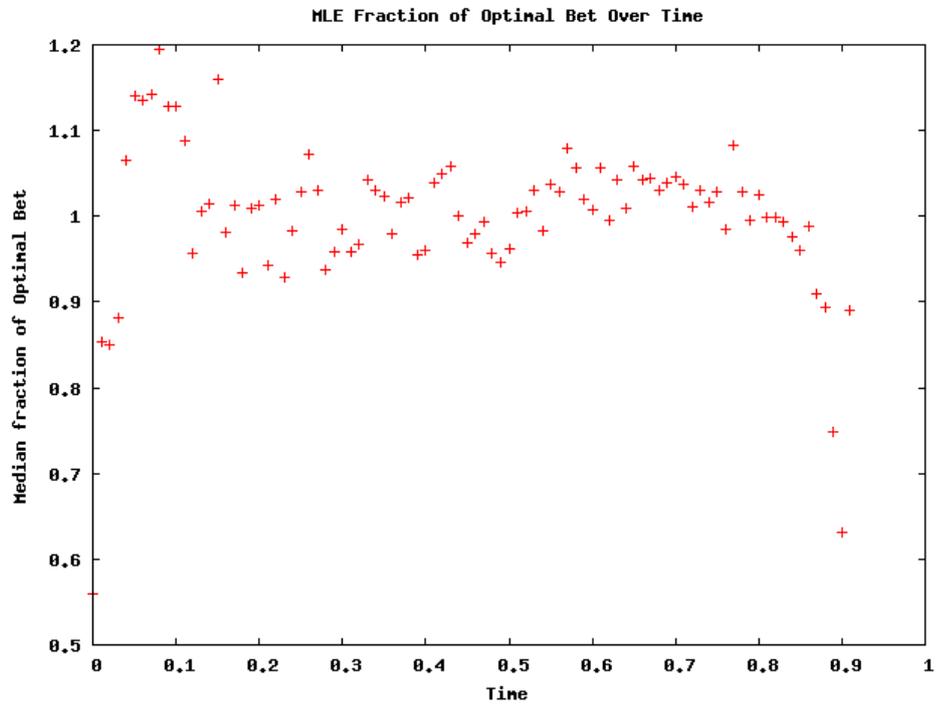


Figure 10: Median Bet of MLE Trader with Forced 10-Period Calibration as Fraction of Optimal Bet

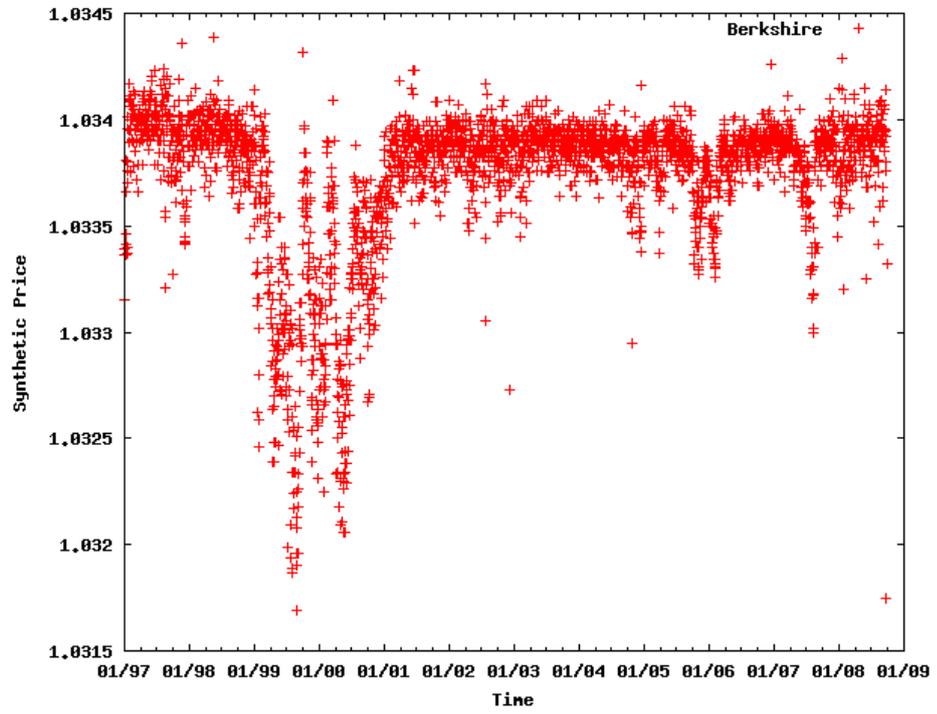


Figure 11: Price of Synthetic Berkshire Pair Security

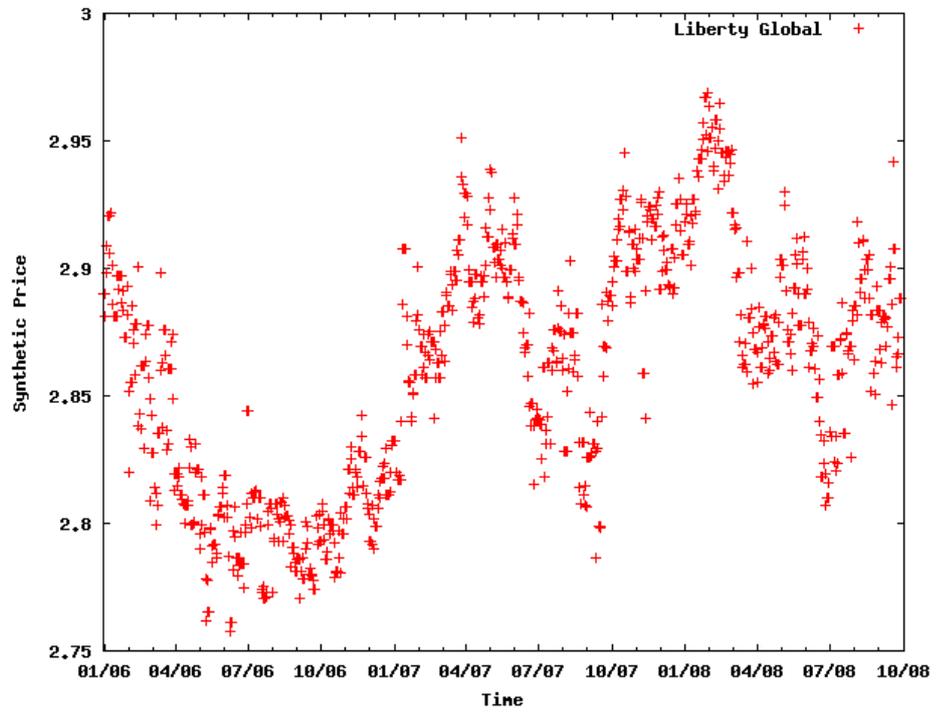


Figure 12: Price of Synthetic Liberty Pair Security

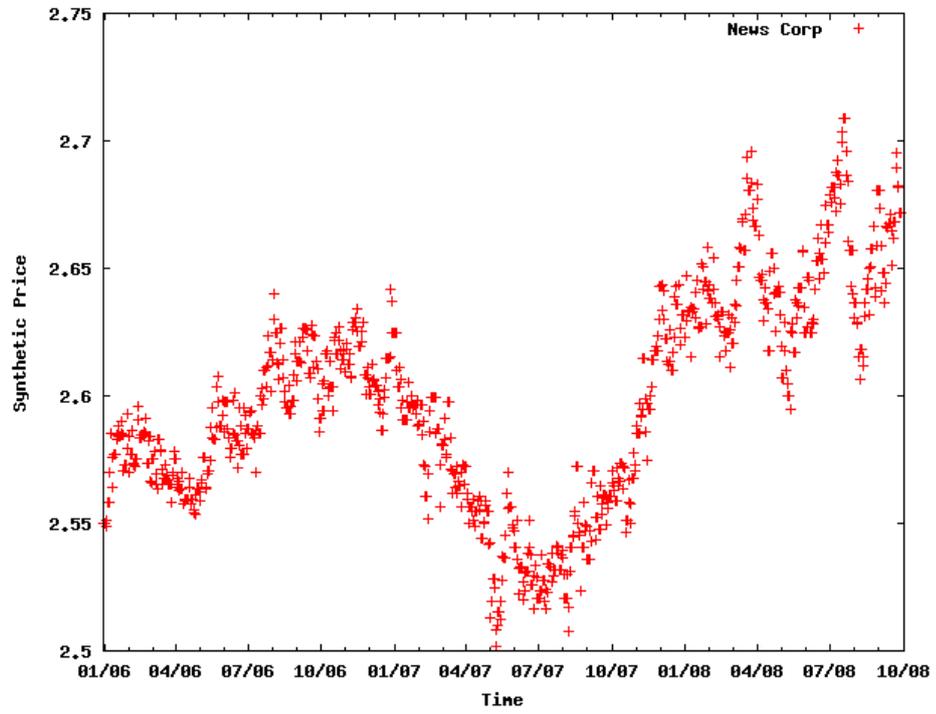


Figure 13: Price of Synthetic News Corp Pair Security

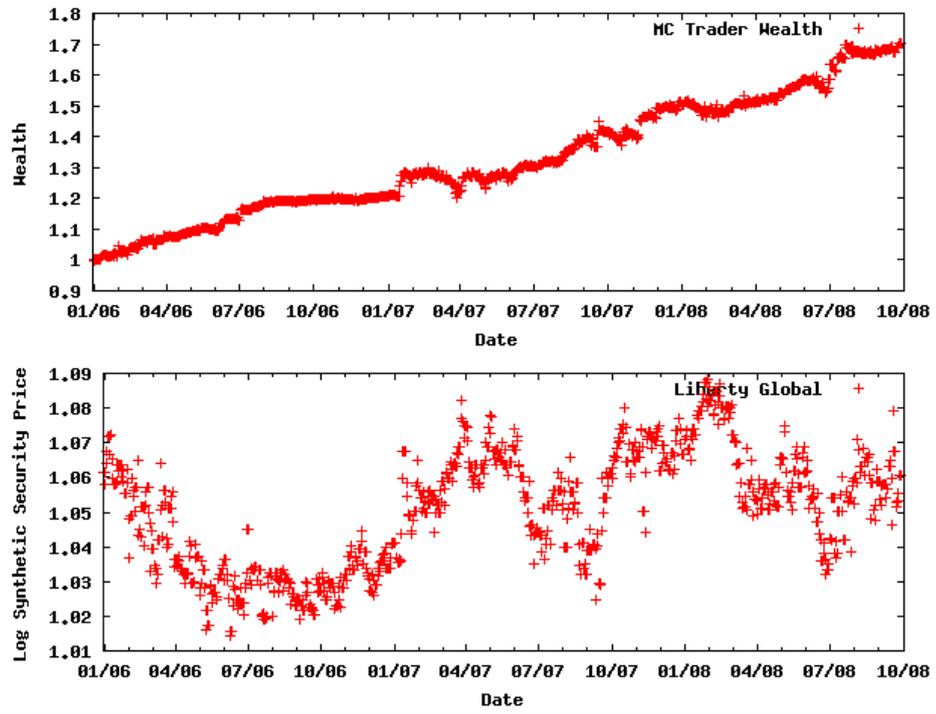


Figure 14: Monte Carlo Trader Results for Liberty Pair

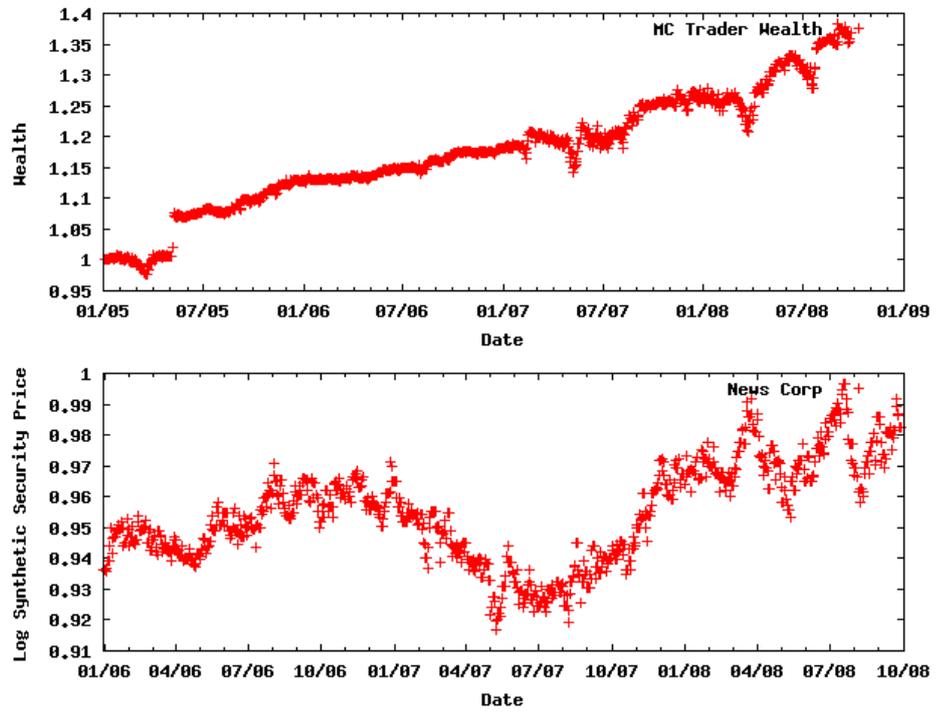


Figure 15: Monte Carlo Trader Results for News Corp

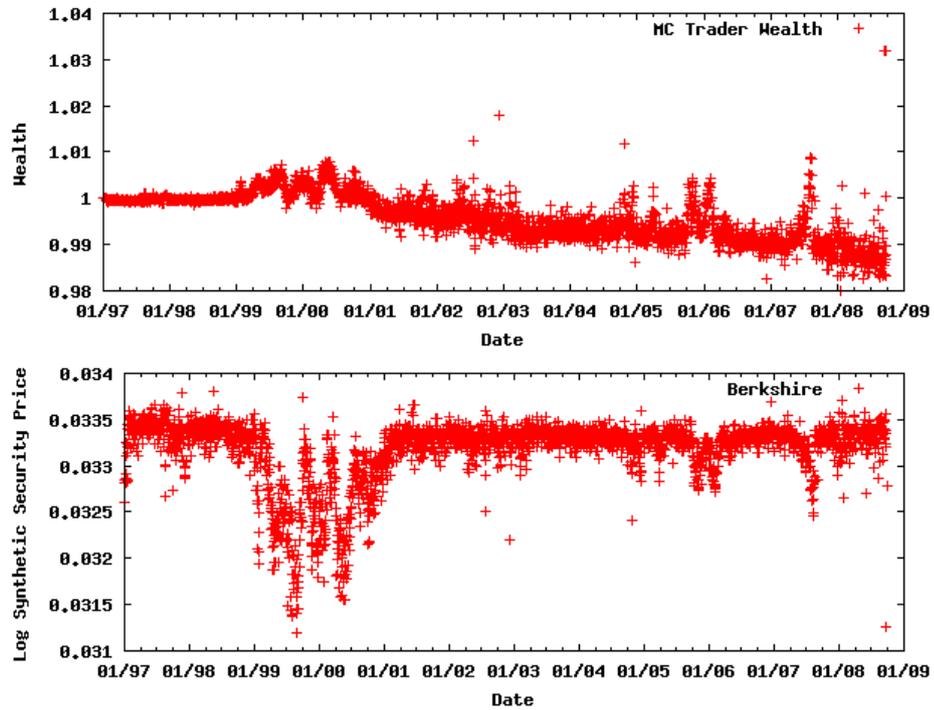


Figure 16: Monte Carlo Trader Results for Berkshire

References

- Andersen, T., 'Personal communication', (2009).
- Andersen, T.G., Davis, R.A., Kreiß, J.P. and Mikosch, T., *Handbook of Financial Time Series* (Springer Berlin, 2008).
- Boguslavsky, M. and Boguslavskaya, E., 'Arbitrage Under Power', *Risk*, vol. 17, no. 6, 69–73 (2004).
- Daum, F. and Co, R., 'Nonlinear filters: beyond the Kalman filter', *IEEE Aerospace and Electronic Systems Magazine*, vol. 20, no. 8 Part 2, 57–69 (2005).

- Doucet, A. and Johansen, A.M., 'A Tutorial on Particle Filtering and Smoothing: Fifteen years later', *The Oxford Handbook of Nonlinear Filtering*, Oxford University Press. To appear (2009).
- Doucet, A. and Tadić, V.B., 'Parameter estimation in general state-space models using particle methods', *Annals of the Institute of Statistical Mathematics*, vol. 55, no. 2, 409–422 (2003).
- Gatev, E., Goetzmann, W.N. and Rouwenhorst, K.G., 'Pairs trading: Performance of a relative-value arbitrage rule', *Review of Financial Studies*, vol. 19, no. 3, 797–827 (2006).
- Hahn, M., Putschogl, W. and Sass, J., 'Portfolio optimization with non-constant volatility and partial information', *Brazilian Journal of Probability and Statistics*, vol. 21, no. 1, 27–61 (2007).
- Hakansson, N.H., 'On optimal myopic portfolio policies, with and without serial correlation of yields', *Journal of Business*, pp. 324–334 (1971).
- Haussmann, U.G. and Sass, J., 'Optimal terminal wealth under partial information for HMM stock returns', in *Mathematics of Finance: Proceedings of an Ams-IMS-Siam Joint Summer Research Conference on Mathematics of Finance, June 22-26, 2003, Snowbird, Utah*, vol. 351, p. 171, American Mathematical Society (2004).
- Johannes, M.S., Polson, N. and Stroud, J.R., 'Nonlinear filtering of stochastic differential equations with jumps', (2002).
- Johansen, Adam M., 'SMCTC: Sequential Monte Carlo in C++', *Journal of Statistical Software*, vol. 30, no. 6, 1–41 (2009).
URL: <http://www.jstatsoft.org/v30/i06>

- Johansen, A.M., Doucet, A. and Davy, M., 'Particle methods for maximum likelihood estimation in latent variable models', *Statistics and Computing*, vol. 18, no. 1, 47–57 (2008).
- Kelly Jr, J., 'A new interpretation of information rate', *Information Theory, IRE Transactions on*, vol. 2, no. 3, 185–189 (1956).
- Kim, SJ and Iltis, RA, 'Performance comparison of particle and extended kalman filter algorithms for GPS C/A code tracking and interference rejection', in *Proc. of Conf. Information Sciences and Systems* (2002).
- Lakner, P., 'Utility maximization with partial information', *Stochastic processes and their applications*, vol. 56, no. 2, 247–273 (1995).
- Ly, Y., Meister, B.K., Weinberger, E.D., Zumbach, G., Eom, C., Park, J., Jung, W.S., Kaizoji, T., Kim, Y.H., Peters, O. *et al.*, 'Application of the Kelly Criterion to Ornstein-Uhlenbeck Processes', *Arxiv preprint arXiv:0903.2910* (2009).
- Mudchanatongsuk, S., Primbs, JA and Wong, W., 'Optimal pairs trading: A stochastic control approach', in *American Control Conference, 2008*, pp. 1035–1039 (2008).
- Pole, A., *Statistical arbitrage: Algorithmic trading insights and techniques* (Wiley, 2007).
- Thorp, E.O., 'The Kelly criterion in blackjack, sports betting, and the stock market', *Handbook of Asset and Liability Management*, vol. 1, 385–428 (2006).
- Toni, T., Welch, D., Strelkowa, N., Ipsen, A. and Stumpf, M.P.H., 'Approximate Bayesian computation scheme for parameter inference and model selection in dynamical systems', *Journal of The Royal Society Interface*, vol. 6, no. 31, 187–202 (2009).