Adaptive Control Applied to Financial Market Data

J.Sindelar
Charles University, Faculty of Mathematics and Physics and Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Prague, Czech Republic.
jan.sindelar@matfyz.cz

M.Kárný
Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Prague, Czech Republic.
school@utia.cas.cz

Abstract. This article describes a formal approach to decision making optimization in commodity futures markets. Our aim is to plan optimal decision at a given time, where we could decide to buy or sell a commodity contract or stay out of the market. The decision is made using dynamic programming with loss function equal to negative profit measured in money, where the probability density functions(PDF) are estimated using Bayesian learning. Predictive PDFs are computed using parametric models from exponential family, giving us easy to adapt systems. Trading costs (slippage and commission) are taken into account. The theory is supported by a series of experiments indicating the measure of success in predicting the market price movements.

Addressed problem

Using mathematical methods for prediction of financial markets has become popular at the end of last century with evolution of Quantitative Finance [Shreve(2004a); Shreve(2004b)]. Our main goal is to exploit dependence of price evolution on previous state of the world. To describe the state we use many different quantities - previous price maxima and minima, variance, commitment of traders information or own engineered quantities taken out of trading experience. The apparatus used is general and normalized data of any kind can be used. Later on, in the experimental section, we present experiments using generally obtainable data only, mainly past prices. In this general setup we are allowed to buy a contract, sell a contract or hold our position and we try to maximize our profit. For the simplicity of the problem, we can have a maximal position of +1 contract and minimal position of -1 contract (for details about a futures contract see [Jílek(2004)] or [Lee et al.(1990)Lee, Finnerty, and Wort]). The experiments should show if we are engaged in a game with positive or negative gain expectation and we can use money management rules to enhance our profit. Here we don't enhance the model and use only the simplest setup. For the problem to be computationally solvable, we supply several restrictions and approximations specified carefully in third section of this article.

Theoretical background

For optimization of the position held in the market we use the mathematical apparatus of stochastic dynamic programming, we will give a brief summary of here. At first, we suppose a simple binomial model of stock price evolution. At a given time $t \in 0, 1, 2, ..$ we choose an action a_t - we buy ($a_t = +1$), sell ($a_t = -1$) or hold ($a_t = 0$). Then, between t and *t+1* the price Δ_t either rises or falls by 1 and again we choose our action. Therefore we get a sequence $Q = (a_0, \Delta_0, a_1, \Delta_1, ..)$. If we are at a time t in the middle of such sequence, we can split it into two or more parts as $(P_t, .., F_t)$, P standing for the past and F for the future. In place of the dots can be any number of listed members of Q. We mark sets with an asterisk so for example set of possible pasts at time t is P_t^*.

Theorem 1 (Stochastic dynamic programming) *Let P_t^* be a set of possible paths the price of an asset could have gone through to reach a given price at time t and following actions taken by the decision maker and let $R_t^* = \{R_t : P_t \to a_t\}$ be a set of possible strategies(e.g.set of sequences of mappings of the pasts into the actions). An optimal causal strategy $\bar{R}_t \in R_t^*$ is a strategy that satisfies*

$$\bar{R}_t = \operatorname{argmin} \left\{ \min_{R_t^*} E\left[Z\left(Q\right)\right] \right\}$$

where Z(Q) is a loss function chosen by the decision maker. As long as the optimization has a finite horizon $\bar{t}$ such a strategy can be constructed in a value-wise way against the course of time. For every t and each $P_t \in P_t^*$ an optimal a_t is chosen so that

$$\nu\left(P_t\right) = \min_{a_t \in a_t^*} E\left[\nu\left(P_{t+1}\right) | a_t, P_t\right]$$

where $\nu\left(P_t\right)$ is a so called Bellman function. The recursion starts with

$$\nu\left(P_{\bar{t}+1}\right) = Z(P_{\bar{t}+1})$$

and the reached minimum has the value

$$E\left[\nu\left(P_1\right)\right] = \min_{R^*} E\left[Z(F_1)\right]$$

For the proof and further details see [Nagy et al.(2005)Nagy, Pavelková, Suzdaleva, Homolová, and Kárný]. $E\left[\cdot\right]$ can be a general isotonic functional of its argument, but under some quite general conditions [Nagy et al.(2005)Nagy, Pavelková, Suzdaleva, Homolová, and Kárný] and if the decision maker is risk-indifferent (doesn't use utility function to show preference of certain loss levels) it coincides with mathematical expectation. Such conditions are supposed from now on. The finiteness of horizon doesn't hold for general market data, as markets don't have an ending date, but as we will see later on for the problem to be solvable, we have to trim the horizon at a finite point in the future. The choice of proper loss function is one of the most difficult tasks for the decision maker. Common sense tells us that such a loss function should closely correspond to the negative net profit in the market. It shows not to be the only criterion for a trader, because we also have to somehow involve a measure of losses through the trading period (a system, even with high profit, that reached a certain loss is not tradable) and we need to take the time value of money into the account (later profits should be discounted by a given interest rate) [Föllmer and Schied(2002)]. In presented experiments we use a loss function that is simply the negative of a profit measured in money and the improvements should be implemented in the future. The loss function we use is additive

$$Z(Q) = \sum_{t \in t^*} z\left(\Delta_t, a_t\right)$$

$z\left(\Delta_t, a_t\right)$ being a partial loss at time t and we can therefore use

Theorem 2 (Stochastic dynamic programming for additive loss function) *Let the condition stated in Theorem 1 hold and let the loss function be additive. Then, the optimal strategy $\bar{R}_t \in R_t^*$ can be constructed value-wise against the course of time by taking*

$$\nu\left(P_t\right) = \min_{a_t \in a_t^*} E\left[z\left(\Delta_t, a_t\right) + \nu\left(P_{t+1}\right) | a_t, P_t\right]$$

starting from

$$\nu\left(P_{\bar{t}+1}\right) = 0$$

In the case of simple binomial model (or even other model of discrete nature) the conditional expected loss has the form

$$E\left[Z(P_t, a_t, \Delta_t, F_t)\right] = \sum_{\Delta_t} p\left(\Delta_t | a_t, P_t\right) \cdot \left[z\left(a_t, \Delta_t\right) + \nu\left(P_t, a_t, \Delta_t\right)\right]$$

but because we suppose that our capital is small compared to the capital traded in the market, we assume the probabilities are independent of our actions and therefore

$$p\left(\Delta_t | a_t, P_t\right) = p\left(\Delta_t | P_t\right)$$

If this condition was not met, we would deal with a control problem, where we could influence the market by trading. Such a strategy is often used in the markets and could be of interest to large financial institutions. In the market the price can move in ticks, which are usually very small compared to the daily movement of the price. Because the optimization is done on a daily basis, we use probability distribution functions to describe the probabilities of asset movements and in such a case

$$E\left[\nu\left(P_{t+1}\right) | a_t, P_t\right] = \int_{\Delta_t^*} \nu\left(P_t, a_t, \Delta_t\right) dp\left(\Delta_t | P_t\right) = \int_{\Delta_t^*} \nu\left(P_t, a_t, \Delta_t\right) f\left(\Delta_t | P_t\right) d\Delta_t$$

If the probabilities p (discrete case) or PDFs f (continuous case) were known, we could simply choose the best path at every time t. Because we don't know the exact probabilities, we try to learn them from the outside world, from the data we obtain. We first split the probability density function $f\left(\Delta_t | P_t\right)$ into a part called *parametric model* and a part called *posterior parameter PDF* as follows

$$f\left(\Delta_t | P_t\right) = \int_{\Theta^*} f\left(\Delta_t | P_t, \Theta\right) f\left(\Theta | P_t\right) d\Theta$$

so that we reduce the uncertainty to only a vector of parameters Θ and we choose the parametric model in a friendly form we will show later on. We suppose the parameters don't change in time. This is a very rough approximation and should be altered in more sophisticated models. Thanks to this condition, we can use the apparatus of *Bayesian estimation* [Kárný et al.(2004)Kárný, Böhm, Guy, Jirsa, Nagy, Nedoma, and Tesar] and under a requirement, that chosen parametric model corresponds to the real world the following holds

Theorem 3 (Bayesian estimation) *The evolution of the PDF $f\left(\Theta | P_t\right)$ generating a so called posterior PDF of unknown parameter is described by recursion*

$$f\left(\Theta | P_{t+1}\right) = \frac{f\left(\Delta_t | P_t, \Theta\right) f\left(\Theta | P_t\right)}{f\left(\Delta_t | P_t\right)}$$

and the result for time t is

$$f\left(\Theta | P_{t+1}\right) = \frac{f\left(\Theta\right) \prod_{\tau \leq t} f\left(\Delta_\tau | P_\tau, \Theta\right)}{\int_{\Theta^*} f\left(\Theta\right) \prod_{\tau \leq t} f\left(\Delta_\tau | P_\tau, \Theta\right) d\Theta}$$

where $f\left(\Theta\right)$ a so called prior PDF is a parameter PDF at time $t = 0$ that has to be chosen from expert information.

Choice of the model

In general the above presented general setup of decision making is rather too complex for solving if the model chosen is not friendly for computations. Because of that we choose the PDFs from an exponential family, where the estimation multiplication of exponential functions changes into addition of their exponents. We can write the parametric model as

$$f(\Delta_t|P_t, \Theta) = A(\Theta) e^{B^T(\Psi_t)C(\Theta)}$$

where Ψ_t is a finite dimensional data vector $\left(\Delta_{t-dim(\Psi_t)}, a_{t-dim(\Psi_t)}, \Delta_{t-dim(\Psi_t)+1}, .., \Delta_t, a_t\right)$, A is a non-negative scalar function and B, C are vectors of functions and T denotes transposition. Such a choice of a model is again a very rough approximation of financial time series. In a dynamic case it restricts us to a Gaussian type probability density functions, which are rarely observed in the markets. Financial data usually have heavier tails and can be skew [Tsay(2005)]. We hope to improve our model by employing Gaussian mixture PDFs in the future. For the needs of this article, we stay with the simple model. In such a setup the estimation task becomes straightforward

Theorem 4 (Bayesian estimation in exponential family) *If the prior PDF $f(\Theta)$ is chosen from conjugate family, e.g. is of the form*

$$f(\Theta) \propto A^{\rho_0}(\Theta) e^{V_0^T C(\Theta)}$$

where V_0 is a $dim(\Psi_t) \times dim(\Psi_t)$ matrix, ρ_0 a so called sample counter, the time evolution of posterior PDF is simplified to addition of exponents

$$V_t = V_{t-1} + B(\Psi_t)$$

and increment

$$\rho_t = \rho_{t-1} + 1$$

the recursion starting from ρ_0, V_0. The model of outer world, needed for dynamic programming, therefore evolves as

$$f(\Delta_t|P_t) = \frac{I[V_{t-1} + B(\Psi_t), \rho_{t-1} + 1]}{I[V_{t-1}, \rho_{t-1}]}$$

and

$$I(V, \rho) = \int_{\Theta^*} A^{\rho}(\Theta) e^{V^T C(\Theta)} f(\Theta) d\Theta$$

Experimental results

For our last experiments we have chosen an autoregressive model to predict price evolution from historical prices. We have used Matlab software together with Mixtools libraries [Nedoma et al.(2005)Nedoma, Kárný, Böhm, and Guy] developed at Institute of Information Theory and Automation CAS. Daily data from 49 futures markets were used, collected from 1990 to 2005. We have used a very simple one dimensional setup where only closing prices were considered. Various auto-regression orders were tested and in accordance to previous research [Tsay(2005)], the AR(2) model of the form:

$$f(\Delta_t|P_t, \Theta) = N(\theta_1 \Delta_{t-1} + \theta_2 \Delta_{t-2}, \sigma), \ \Theta = (\theta_1, \theta_2, \sigma)$$

Figure 1. An example of outcome for Australian Government Bonds. In the first frame, we can see the price evolution second picture shows actions taken over time, the third one is position held and the last one is profit in U.S. dollars. Time is labelled in days.

has shown to be the most suitable choice. Because an infinite horizon for dynamic programming would be to difficult to compute, we use a window of the size of 7 trading days for planning. To consider that the probability density functions are not stationary over a period of 20 years, we employ a forgetting factor $\lambda \in (0, 1)$ and multiply the exponents of $f\left(\Delta_{\tau \leq t} | P_\tau, \Theta\right)$ by $\lambda^{t-\tau}$ [Kárný et al.(2004)Kárný, Böhm, Guy, Jirsa, Nagy, Nedoma, and Tesar]. In agreement with the theory the estimation of probability distribution functions should be performed also for the individual steps of dynamic programming. The probability densities should therefore be different for times $t, t+1, .., t+7$. Again to simplify the calculations we use a single PDF for each time in the moving 7-day window and update the PDF as soon as the window moves. For simplicity point estimates of the parameters are used for the auto-regression. The results of such model are shown in Figure 1 for one of the markets. The Australian Government Bonds are a random representative of markets observed (A differentiation of the markets by type is planned in the future, to evaluate the quality of prediction depending on type). In the first frame is price evolution in time. Second frame shows performed actions in time. Third frame is position held at time t and the last frame shows gains from trading. The reason, these profits jump from positive to negative is, we have used a loss function, that evaluates our position only if we perform an action and not continuously as would be needed, if we wanted to incorporate risk into the loss. In Figure 2 we see a histogram of profitable markets in the experiment.

Conclusion

From Figure 2 we see our prediction is not very good under given conditions. We hope to improve the result by building a more sophisticated model. More data channels should be used, especially those containing fundamental information about traders as commitment of traders data. We have performed certain experiments in the past using such channels and the

Figure 2. A histogram of profits in individual markets over the period of years from 1990 to 2005. 49 markets were tested mostly from U.S. futures markets. Profits and losses are shown in U.S. dollars.

information value seemed quite high. Also we want to improve the prediction by selecting a *mixture of models*

$$f(\Delta_t|P_t,\Theta) = \sum_{c\in c^*} \alpha_c f_c\left(\Delta_t|P_t,\Theta_c,\alpha_c\right)$$
$$\sum \alpha_c = 1,\ \alpha_c \geq 0,\ \Theta = (\alpha_c,\Theta_c)$$

especially a mixture of models from the exponential family (Gaussian). Using such a mixture carries computational difficulties, because multiples of additions are needed for learning, but the problem is generally solvable [Kárný et al.(2004)Kárný, Böhm, Guy, Jirsa, Nagy, Nedoma, and Tesar].

Acknowledgments. The authors thank Colosseum a.s. for the supply of data and cooperation in the fields of testing and data preprocessing. The present work was supported by MŠMT ČR under Contract 2C06001 and by the Academy of Sciences of the Czech Republic under Contract 1ET10050401.

References

Föllmer, H. and Schied, A., *Stochastic Finance*, Walter de Gruyter, 2002.

Jílek, J., *Finanční a komoditní deriváty*, Grada, 2004.

Kárný, M., Böhm, J., Guy, T. V., Jirsa, L., Nagy, I., Nedoma, P., and Tesar, L., *Optimized Bayesian Dynamic Advising*, Springer, 2004.

Lee, C. F., Finnerty, J. E., and Wort, D. H., *Security Analysis and Portfolio Management*, Scott, Foresman and Co., 1990.

Nagy, I., Pavelková, L., Suzdaleva, E., Homolová, J., and Kárný, M., *Bayesian Decision Making*, UTIA AV ČR, 2005.

Nedoma, P., Kárný, M., Böhm, J., and Guy, T. V., Mixtools Interactive User's Guide, Tech. Rep. 2143, ÚTIA AV ČR, Praha, 2005.

Shreve, S. E., *Stochastic Calculus for Finance I.*, Springer, 2004a.

Shreve, S. E., *Stochastic Calculus for Finance II.*, Springer, 2004b.

Tsay, R. S., *Analysis of Financial Time Series*, John Wiley and Sons, second edn., 2005.