# Active and Passive Learning in Agent-based Financial Markets

Blake LeBaron *

International Business School

Brandeis University

July 2010

**Abstract**

This short note compares and contrasts two forms of learning which are present in most agent-based financial markets. First, passive learning refers to a form of "as if rationality" where wealth accumulates on strategies which have done relatively well. Second, active learning refers to the active switching of agents across strategies. Most heterogeneous agent markets contain some form of both these types of learning. From what we know so far the dynamics of each may be quite different, and may yield a rich and complex joint dynamic.

*International Business School, Brandeis University, 415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. The author is also a research associate at the National Bureau of Economic Research.

# 1 Introduction

The construction of agent-based financial markets is now a field with nearly 20 years of experience to learn from. There are many basic principles of methodology and design that have been learned over the years. This short note will briefly comment on one aspect of markets and learning that is often ignored, the interaction between active and passive learning dynamics. I will define and argue that both these forms of learning are important to financial market dynamics. Both of these have been used by many authors, but rarely have the interactions between the two been explored. Further, few authors explore the relative strengths and weaknesses of using these forms of learning in a financial setting. In this way, this note serves as a quick reminder about what we are doing, and as a suggestion for important future research into how heterogeneous agent financial markets function.[1]

It will be important to first define active and passive learning in a agent-based market. Passive learning refers to the accumulation of wealth in strategies which have been successful. Good strategies thrive and become a larger part of the market, while weak strategies eventually die off. This is a version of "as if rationality" described originally in Friedman (1953). It is often used as a metaphor for convergence to some form of market efficiency, or at least a selection mechanism which would weed out ineffective strategies. The basic premise of who might survive in the long run out of a sea of different strategies is an old one in finance. It is tied to the original betting rules of Kelley (1956), and growth optimal portfolios which were debated in the the 1960's and 1970's in papers such as Samuelson (1971) and Hakansson (1971). A recent and up to date survey on this area is contained in Evstigneev, Hens & Schenk-Hoppe (2009).[2] In the next section I will describe some of the modeling features, strengths, and weaknesses of passive learning.

The other form of learning, active learning, may be closer in spirit to what people are thinking about when they imagine learning in a financial or economic setting. Agents actively chose strategies, with some well defined objective function in mind. This form of learning is part of almost all of the heterogeneous agent markets which consider dynamic strategy adjustment. Agents may be switching over fixed strategies, or over a set of evolving strategies as in markets built with genetic algorithms. In all cases, there is an active

[1] This is not a survey of learning, or heterogeneous agent models in finance. This is well beyond the scope of this short paper. On heterogeneous agent models many excellent surveys exist including, Chiarella, Dieci & He (2009), Hommes (2006), LeBaron (2006), and Lux (2009). On learning models in finance in general a recent survey of this large literature can be found in Pastor & Veronesi (2009).

[2] Another early theoretical derivation is in Breiman (1961). A nice summary of this is in Markowitz (1976). Blume & Easley (1990) and Blume & Easley (2006) state the problem in the context of a utility maximizing portfolio decision. The latter paper paper proves that in a complete market world the convergence to true beliefs will occur regardless of preference parameters. However, the authors point out that in an incomplete market world this convergence is not guaranteed. Evstigneev, Hens & Schenk-Hoppe (2006) look at an incomplete markets world with endogenous prices. In their framework the growth optimal strategy will dominate any other competing strategy in terms of acquiring all wealth in the long run.

attempt by agents to move their wealth into strategies that have performed well in the recent or distant past.

The next section will make these ideas clearer in a simple market framework. It will also go through some of what we know about these systems, and some conjectures about what we may find out in the future.

## 2 A simple model framework

First, I will describe a simple market framework, from which the principles of passive and active learning will be made clear. This is far from a fully developed market, and only represents a skeleton for a market representation. In the most basic of markets I will assume a world with a risky asset that pays a dividend at time $t$, $D_t$. The dividend will follow some arbitrary stochastic process.[3] Individual agents (indexed by $i$) are assumed to purchase shares in this risky asset, $S_{t,i}$. They also hold, $B_{t,i}$ units of a risk free asset which pays an interest rate $r_f$. The intertemporal budget constraint for agent $i$ is given by

$$W_{t,i} = P_t S_{t,i} + B_{t,i} = (P_t + D_t)S_{t-1,i} + B_{t-1,i}(1 + r_f) - C_{t,i} \tag{1}$$

$W_{t,i}$ represents the wealth of agent $i$ at time $t$, and $C_{t,i}$ is consumption at time $t$. If consumption is assumed to be some fraction of wealth determined by, $\lambda(I_t)$, a function of information at time t, the above budget becomes

$$W_{t,i} = P_t S_{t,i} + B_{t,i} = (1 - \lambda(I_t))((P_t + D_t)S_{t-1,i} + B_{t-1,i}(1 + r_f)). \tag{2}$$

Two further assumptions can be useful in modeling. First, simplifying the consumption decision to a constant fraction of wealth gives,

$$W_{t,i} = P_t S_{t,i} + B_{t,i} = (1 - \lambda)((P_t + D_t)S_{t-1,i} + B_{t-1,i}(1 + r_f)). \tag{3}$$

A second, and less used, assumption, is to set $r_f = 0$. This can be done to restrict the incoming resources to the economy to the dividend stream alone which makes the model a simple general equilibrium economy with costless storage in the consumption good.

Learning in this world takes place in the portfolio choice of individual agents. Assume that $\alpha_j(I_t)$ is a investment strategy (indexed by $j$) that yields the fraction of wealth to put into the risky asset. In general,

[3] Often this can be calibrated to some actual macro series.

an agent could spread a fraction of wealth over different strategies. Let $\omega_{i,j}$ be the fraction of wealth of agent $i$ in strategy $j$. In most agent-based models this value is either zero or one as agents concentrate their wealth in only one strategy. Both of these are functions of information at time t, $I_t$. Share demand for an agent $i$ summed over strategies $j$ is given by

$$S_{t,i} = \frac{\sum_{j=1}^{J} \omega_{i,j}(I_t)\alpha_j(I_t)(1-\lambda)W_{t,i}}{P_t}. \tag{4}$$

This share demand, and strategy is important in exploring active and passive learning. The key feature is that the demand for shares is proportional to wealth. This would be the outcome of most constant relative risk aversion preferences (not constant absolute risk aversion). The economy is closed by setting the total supply of shares to 1,

$$1 = \sum_{i=1}^{I} S_{t,i}. \tag{5}$$

It is important to note that pricing in this market depends not on the number of traders using a given strategy $j$, but on the wealth in strategy $j$ which would be written as

$$Z_{t,j} = \sum_{i=1}^{I} \omega_{i,j}(I_t)W_{t,i}. \tag{6}$$

Further learning dynamics can take place through the adaptive learning of the rules rules themselves. In this case, the rules become dynamic, and parameterized by $\theta_t$, giving $\alpha(I_t, \theta_t)$. Learning occurs in the rules as $\theta_t$ moves through time. This corresponds to classic learning models such as Bayesian updating or recursive least squares.[4]

This now forms the skeleton for a simple agent-based economy with a working financial market. Details of agent learning and behavior go into building sets of strategies, $\alpha_j(I_t)$, and methods for agents choosing strategies over time. A model of this form would have both active and passive learning, and I will use its structure to clearly define the concepts.

# 3 Passive learning

This market could represent pure passive learning with no active learning. This case would correspond to $\omega_{i,j}$ being constant, and agents stay with fixed strategies. Their strategies may be dynamic, in that $\alpha_j(I_t)$

[4]Some agent-based learning models go further in that the functional forms of the rules themselves are allowed to change over time as in Chen & Yeh (2001), or Arthur, Holland, LeBaron, Palmer & Tayler (1997).

is allowed to depend on current and past information in complex fashions, but the agents all stay with their given strategies no matter how poorly they are doing. For model design this is a powerful learning concept. As long as there is at least some persistence in the agents' decisions, $\omega_{i,j}$, there will be some form of passive learning, or wealth adjustment to successful strategies in the market.[5] So a strength of this form of learning, is that it is easy to model, and probably somewhat ubiquitous in all real and artificial markets.

Unfortunately, it comes with several drawbacks that are important to consider. First, passive learning is not equivalent to utility maximization. Wealth does not select utility maximizing strategies except in particular cases.[6] Many authors have made this point, but one of the sharpest examples is Blume & Easley (2006). The key result there is that with incomplete markets, and preferences that deviate from log, wealth will move to strategies with beliefs that deviate from true probabilities.

LeBaron (2007) provides a simple real world calibrated example showing how this bias may be important in asset pricing and observed investor behavior. The experiment considers investors constructing portfolios from a risky asset yielding an exogenous returns process, and a risk free bond. Returns are constructed as in Campbell & Viceira (2002),

$$r_{t+1} = x_{t+1} + e_{t+1} \tag{7}$$

$$x_{t+1} = \mu + \rho(x_t - \mu) + \eta_{t+1}. \tag{8}$$

which is a common representation which contains a predictable component, $x_t$, which is only observed subject to observational noise $e_{t+1}$.[7] The return parameters are calibrated to replicate actual financial return series. The parameters for the process are given in table 1, and the simulations (and portfolio adjustments) are made at weekly frequencies, and are run for the equivalent of 500 years.

The optimal forecast is given by the Kalman filter, and has the form,

$$E_t(r_{t+1}) = f_{t+1} = \mu + \rho_j(f_t - \mu) + w_j(r_t - f_t). \tag{9}$$

The parameter $w_j$ is the critical gain parameter which controls how recent observations should be weighted when building forecasts. Using the true time series parameters the optimal parameters, $(w^*, \rho^*)$ can be

[5] There is one important class of models where passive learning is inactive. Models with CARA utility and adaptive rule selection generally have no passive learning component. Two very different examples of this are Brock & Hommes (1998) and Arthur et al. (1997). Price formation depends on the fraction of traders in a given strategy, and not on their wealth.

[6] The best known case would be log utility.

[7] See Pastor & Stambaugh (2009) for a more complete treatment of systems of this form in finance.

determined. For the parameters used here these are $(0.0164, 0.95)$.

Assume investors determine their portfolio fraction using a standard mean variance decision rule,

$$\alpha_{t,j} = \frac{E_t^j(r_{t+1}) - r_f + \sigma_t^2/2}{\gamma\sigma_t^2}, \tag{10}$$

where the expectation corresponds to a given parameter pair $(w_i, \rho_i)$.[8] $\gamma$ is the coefficient of relative risk aversion. The variance, $\sigma_t^2$, is assumed to be constant and known to all strategies. In the following experiments $\alpha$ is bounded between $-0.5$ and 2 allowing for some short sales and leverage. Wealth evolution will be simulated for a grid of different forecast strategies.[9]

When $\gamma = 1$ preferences correspond to log preferences and the passive wealth evolution selects the optimal forecast parameters. However, when $\gamma$ differs from one, interesting results are observed in terms of who survives in the long run. This can be seen in figure 1. The lower panel displays the utility contours across the different strategies measured as the annual certainty equivalent return.[10] The utility maximizing strategy is centered on the true Kalman forecast parameters. The upper panel displays contours based on the final wealth distribution after 500 years of simulated data and portfolio strategies. It shows the clear bias in passive learning. The maximum wealth forecast corresponds to a parameter pair of $(\omega, \rho) = (0.06, 1.00)$ which is far from the true parameters, both in terms of actual values, and in terms of expected utility. The estimated certainty equivalent return at this point is only 2.9 percent per year which compares to 5.25 at the maximum utility point.

It is also interesting that the gain parameter is biased high. This would correspond to agents putting too much weight on the recent past than they optimally should. If one were to look the the behavior of surviving agents in this world relative to their observed time series, they would be deemed irrational. The large gain parameters might even suggest they were "momentum traders," putting a large amount of weight on recent trends. Effectively, the biased parameters generate agents who behave closer to log utility. The key point here is that wealth evolution alone selects for something other than rationality, and therefore it should not be confused with rationality.[11]

A second, but much less explored, feature of passive learning, is that it may be very slow. Few models

[8] See LeBaron (2007), or Campbell & Viceira (2002) for derivations and connections to intertemporal preferences. The variance term in the numerator can be thought of as an adjustment for the fact that these are log returns.

[9] The consumption fraction, $\lambda$, is irrelevant for wealth races of this form where it is considered to be the same across all agents. Each period all agents consume the same fraction of wealth, so the relative performance is not affected by $\lambda$.

[10] This is the risk free return which would generate the same utility as the return on the risky asset.

[11] This point has been made by a large number of papers. For a result directly tied to Friedman's examples of firms and profit maximization see Radner (1998), and also Winter (1982).

try to assess the speed of adjustment, since this would depend on calibrating models to real data, and real strategies. However, Berrada (2009), LeBaron (2007), and Yan (2008) all suggest caution on the ability of of this form of learning to be relevant in real data due to its very slow speed. They show that in reasonable financial models convergence may be measured in units of decades, so that extreme patience may be required for this form of learning to be relevant. This is an important question for learning researchers to be concerned with, and should be further explored.

Several early papers also looked at pure wealth evolution, or passive learning, across simple trading strategies. Chiarella & He (2001) and Levy, Levy & Solomon (1994) are both good examples of this. It is interesting to note that both use log preferences (or demands which are closely related to them), so in both cases wealth selection and utility maximization coincide. More recent papers have tried to expand these to include an active learning channel, but few papers have tried to address the possible deviations in learning objectives that might occur when the utilities deviate from log.

## 4 Active learning

Allowing agents to begin adjusting their strategy choices $\omega_{i,j}(I_t)$ changes this to a model incorporating active learning. Active learning is intuitively appealing. It seems like something agents are doing in the real world as they adjust behavior to new information.[12] However, unlike passive learning, modeling this type of learning is challenging, and there are no clear paths for the agent-based model builder.[13] The builder needs to decide on many aspects of how agents select optimal rules. First, what sort of objective function should be used? Should it be profits, or some estimate of expected utility? Second, how much past data, or memory, should this estimator work with when building these estimates? Finally, what fraction of agents should be considering changing rules each period? Should it be a small fraction, or all agents? Should the decision to update depend on current market activity? These are only a few of the many open design questions that have to be answered to model active learning.

As mentioned earlier, active learning can also involve adjusting the forecasting rules over time by changing the parameter, $\theta_t$, in $\alpha(I_t, \theta_t)$. This drives a second learning dynamic beyond agents adapting over rules. Depending on how the model is built this might also follow a utility based gradient, but the speed relative to dynamic rule selection is not clear. It would depend on the structure of the learning model, and how much

[12] The evidence in support of various forms of active learning extends beyond casual introspection. Laboratory evidence shows some support for various forms of active learning. Some of this work in financial markets is surveyed in Hommes (2010).

[13] This is where Sims (1980)'s critique of deviations from rationality is in full force.

they allow the parameters to change given recent observations from the time series.

In the literature on active learning, some frameworks have proved useful and relatively easily applied in many different cases. A good example of this is the simple discrete choice model originally popularized by Brock & Hommes (1997). It is straight forward, yields strong analytic results, and has good micro foundations. However, even in this framework several of these design questions are still open, such as memory, and the fraction of the population updating. Furthermore, the dynamics depends on a crucial parameter, the intensity of choice, that needs to be pinned down.[14]

Models which use only constant absolute risk aversion along with some form of adaptive learning are pure active learning because the accumulation of wealth does not impact the results. Agents share demands do not depend on their relative wealth, and therefore the dynamics of these markets don't depend on any sort of passive wealth accumulation operating in the background. These purely active learning models don't fit into the wealth share demand framework outlined above which is inspired by constant relative risk aversion preferences. More recently wealth has been added to the Brock & Hommes (1998) framework as in Anufriev & Dindo (2010) Another example which combines both active and passive learning in a small set of trading strategies is Chiarella & He (2008). The models in LeBaron (2001) or LeBaron (2010) both include passive and active learning, and use a framework designed to eventually untangle their impact. In all these models learning is a hybrid between both active a passive forms, but there are still few general results on how the different forms of learning interact.

An important issue for adaptive learning, is whether its dynamics are driven more by noise, than actual fitness of various forecasting rules. Given that financial data are very noisy, and attempts to evaluate relative forecasts are often not conclusive, it is likely that in an agent-based market, generating realistic data, the adaptive learning process may be adapting to noise. Movements in strategy space may be more due to genetic drift than actual fitness differences. While this is something researchers should be aware of, it may not be a big problem, since this noisy rule adjustment may be quite realistic.[15] It would correspond to investors shifting funds over mutual funds in response to recent performance.

Figure 1 also gives a glimpse of active learning. Most all active learning systems operate on some expected utility gradient. Therefore, if agents were to make choices across rules based on expected utility they would be approaching the utility maximization point in the lower panel that represents the true Kalman filter parameters. The key question for active learning would be how fast do they get there, and how does this

[14]Important current work has moved in the direction of estimating the intensity of choice as in Goldbaum & Mizrach (2008) or Boswijk, Hommes & Manzan (2007).

[15]It reminds one of Fisher Black's discussions in Black (1986).

learning mechanism coexist with the passive wealth evolution in the top panel. Finally, it is important to stress that this model with its completely exogenous returns process is only a rough thought experiment. In actual markets the endogenous response of prices, and therefore returns, is critical to determining the overall wealth dynamics in a system.

## 5 Conclusions

Realistic heterogeneous agent models of financial markets need to take into account both passive and active learning. Researchers should be aware that they are often using both of these in various modeling frameworks. Each comes with its own set of issues. Passive learning is easy to model, but does not necessarily select for utility maximization, and it may be slow. On the other hand, active learning can move at reasonable speeds, and seems to be an important part of observed behavior. Unfortunately, it is difficult to model, and involves many degrees of freedom. Also, it may often be adapting to noise in financial time series.

In real markets we probably see some combination of these two forms of learning. They may take place at different speeds or time scales, and might generate interesting dynamics as they interact with eachother. Eventually, understanding the impact of both these forms of learning will be important to understanding the dynamics in real financial markets.

## References


Anufriev, M. & Dindo, P. (2010), 'Wealth-driven selection in a financial market with heterogeneous agents', *Journal of Economic Dynamcis and Control* **73**, 327–358.

Arthur, W. B., Holland, J., LeBaron, B., Palmer, R. & Tayler, P. (1997), Asset pricing under endogenous expectations in an artificial stock market, *in* W. B. Arthur, S. Durlauf & D. Lane, eds, 'The Economy as an Evolving Complex System II', Addison-Wesley, Reading, MA, pp. 15–44.

Berrada, T. (2009), 'Bounded rationaity and asset pricing', *Review of Finance* **13**, 693–725.

Black, F. (1986), 'Noise', *Journal of Finance* **41**, 529–543.

Blume, L. & Easley, D. (1990), 'Evolution and market behavior', *Journal of Economic Theory* **58**, 9–40.

Blume, L. & Easley, D. (2006), 'If you're so smart, Why aren't you rich? Belief selection in complete and incomplete markets', *Econometrica* **74**, 929–966.

Boswijk, H. P., Hommes, C. H. & Manzan, S. (2007), 'Behavioral heterogeneity in stock prices', *Journal of Economic Dynamics and Control* **31**(6), 1938–1970.

Breiman, L. (1961), Optimal gambling systems for favorable games, *in* J. Newyman & E. Scott, eds, 'Proceedings of the Fourth Berkeley Symp. of Math Statistics, and Probability', Vol. 1, University of California Berkely Press, Berkely, CA.

Brock, W. A. & Hommes, C. H. (1997), 'A rational route to randomness', *Econometrica* **65**, 1059–1097.

Brock, W. A. & Hommes, C. H. (1998), 'Heterogeneous beliefs and routes to chaos in a simple asset pricing model', *Journal of Economic Dynamics and Control* **22**(8-9), 1235–1274.

Campbell, J. Y. & Viceira, L. M. (2002), *Strategic Asset Allocation*, Oxford University Press, Oxford, UK.

Chen, S.-H. & Yeh, C.-H. (2001), 'Evolving traders and the business school with genetic programming: A new architecture of the agent-based artificial stock market', *Journal of Economic Dynamics and Control* **25**, 363–394.

Chiarella, C., Dieci, R. & He, X.-Z. (2009), Heterogeneity, market mechanisms, and asset price dynamics, *in* T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', Elsevier, USA, pp. 277–344.

Chiarella, C. & He, X.-Z. (2001), ‘Asset pricing and wealth dynamics under heterogeneous expectations’, *Quantitative Finance* **1**, 509–526.

Chiarella, C. & He, X.-Z. (2008), An adaptive model on asset pricing and wealth dynamics with heterogenous trading strategies, *in* D. Seese, C. Weinhardt & F. Schlottmann, eds, ‘Handbook of Information Technology in Finance’, Springer Verlag.

Evstigneev, I. V., Hens, T. & Schenk-Hoppe, K. R. (2006), ‘Evolutionary stable stock markets’, *Economic Theory* **27**, 449–468.

Evstigneev, I. V., Hens, T. & Schenk-Hoppe, K. R. (2009), Evolutionary finance, *in* T. Hens & K. R. Schenk-Hoppe, eds, ‘Handbook of Financial Markets: Dynamics and Evolution’, Handbooks in Finance, North-Holland, Amsterdam, the Netherlands, pp. 509–564.

Friedman, M. (1953), *Essays in Positive Economics*, University of Chicago Press, Chicago, IL.

Goldbaum, D. & Mizrach, B. (2008), ‘Estimating the intensity of choice in a dynamic mutual fund allocation decision’, *Journal of Economic Dynamcis and Control* **32**, 3866–3876.

Hakansson, N. H. (1971), ‘Multi-period mean-vaiance analysis: Toward a general theory of portfolio choice’, *Journal of Finance* **26**.

Hommes, C. H. (2006), Heterogeneous agent models in economics and finance, *in* K. L. Judd & L. Tesfatsion, eds, ‘Handbook of Computational Economics’, Elsevier.

Hommes, C. H. (2010), The heterogeneous expectations hypotheis: Some evidence from the lab, Technical report, CeNDEF, University of Amsterdam.

Kelley, J. L. (1956), ‘A new interpretation of information rate’, *Bell Systems Technical Journal* **35**, 917–926.

LeBaron, B. (2001), ‘Evolution and time horizons in an agent based stock market’, *Macroeconomic Dynamics* **5**(2), 225–254.

LeBaron, B. (2006), Agent-based computational finance, *in* K. L. Judd & L. Tesfatsion, eds, ‘Handbook of Computational Economics’, Elsevier, pp. 1187–1233.

LeBaron, B. (2007), Wealth evolution and distorted financial forecasts, Technical report, International Business School, Brandeis University.

LeBaron, B. (2010), Heterogenous gain learning and the dynamics of asset prices, Technical report, International Business School, Brandeis University, Waltham, MA.

Levy, M., Levy, H. & Solomon, S. (1994), 'A microscopic model of the stock market: cycles, booms, and crashes', *Economics Letters* **45**, 103–111.

Lux, T. (2009), Stochastic behavioral asset pricing stochastic behavioral asset pricing models and the stylized facts, *in* T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', North-Holland.

Markowitz, H. (1976), 'Investment for the long run: New evidence for an old rule', *Journal of Finance* **31**, 1273–1286.

Pastor, L. & Stambaugh, R. F. (2009), 'Predictive systems: Living with imperfect predictors', *Journal of Finance* **64**, 1583–1628.

Pastor, L. & Veronesi, P. (2009), 'Learning in financial markets', *Annual Review of Financial Economics* **1**, 361–381.

Radner, R. (1998), Economic survival, *in* D. P. Jacobs, E. Kalai & M. I. Kamien, eds, 'Frontiers of Research in Economic Theory', Econometric Society Monographs, Cambridge University Press, pp. 183–209.

Samuelson, P. (1971), 'The "fallacy" of maximizing the geometric mean in long sequences of investing or gambling', *Proceedings of the National Academy of Science* **68**, 2493–2496.

Sims, C. A. (1980), 'Macroeconomics and reality', *Econometrica* **48**, 1–48.

Winter, S. (1982), Competiton and selection, *in* J. Eaton & J. Milgate, eds, 'The New Palgrave', Stockton Press, pp. 545–548.

Yan, H. (2008), 'Natural selection in financial markets: Does it work?', *Management Science* **54**(11), 1935–1950.

Table 1: *Return Parameter Values*

| Parameter | Value |
|---|---|
| $r_f$ | 0.02 |
| $E(r_t)$ | 0.07 |
| $\sigma_r$ | 0.20 |
| $\sigma_x^2/\sigma_r^2$ | 0.02 |
| $\rho$ | 0.95 |

**Description:** Parameters for return time series. All values are annualized, but simulations are done at the weekly frequency. $r_f$ is the risk free interest rate. $E(r_t)$ is the unconditional expected real return on the risky asset. $\sigma_r$ is the corresponding annual standard deviation. $\sigma_x^2/\sigma_r^2$ is the signal to noise ratio in the returns series. $\rho$ is the AR(1) persistence parameter for the expected return process.

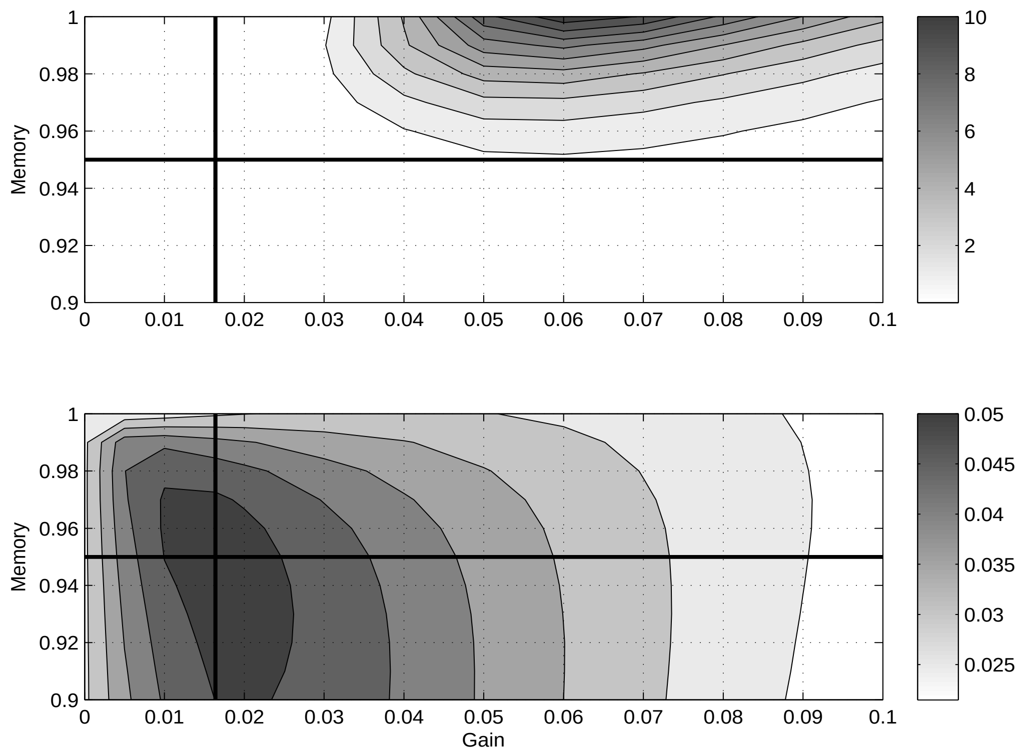


Figure 1: **Wealth and utility surfaces for $\gamma = 3$.**
This upper panel in this figure shows the wealth distribution after 500 years estimated as a mean over a 100 run cross section. The figure shows the density over the different strategies indexed by the memory and Kalman gain parameters. The height measures the density at each grid point relative to a uniform density. The lower panel measures the expected utility of each rule reported in units of annual certainty equivalent returns. The maximum of the wealth density is at the (gain, memory) pair of $(0.06, 1.00)$. The annual certainty equivalent return at this point is 2.91 percent which compares to an annual certainty equivalent return of 5.25 at the optimal forecast parameters.