

Wealth Dynamics and a Bias Toward Momentum Trading

Blake LeBaron *
Brandeis University

December 2010

Abstract

Evolutionary metaphors have been prominent in both economics and finance. They are often used as basic foundations for rational behavior and efficient markets. Theoretically, a mechanism which selects for rational investors actually requires many caveats, and is far from generic. This paper tests wealth based evolution in a simple, stylized agent-based financial market. The setup borrows extensively from current research in finance that considers optimal behavior with some amount of return predictability. The results confirm that with a homogeneous world of log utility investors wealth will converge onto optimal adaptive forecasting parameters. However, in the case of utility functions which differ from log, wealth selection alone converges to parameters which are economically far from the optimal forecast parameters. This serves as a strong reminder that wealth selection and utility maximization are not the same thing. Therefore, suboptimal financial forecasting strategies may be difficult to drive out of a market, and may even do quite well for some time.

*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 faculty research fellow at the National Bureau of Economic Research.

1 Introduction

Evolution has always played an important background role in both finance and economics. Many researchers have taken comfort in thinking that irrational trading strategies, or less than profitable firms would eventually be removed from the market.¹ The theoretical backing for this strong defense of rationality is not as definitive as its proponents would have us think. Different situations require different restrictions on behavior for convergence to rationality. This paper explores this question in a very standard financial forecasting test case where stock returns have some weak predictability. In this world wealth evolution can select for distorted predictors which are economically far from the optimal true probabilities. This demonstrates that some irrational forecasters may be very difficult to remove from a market, and may even thrive. These forecasters put excessive weight on the recent past relative to a Kalman filter benchmark. Therefore, they would appear to be biased toward favoring momentum strategies.²

The fact that wealth growth and utility maximization are not the same thing is well known in finance. It generated a large debate in the 60's and 70's about the normative case for holding portfolios that maximized the growth rate of wealth.³ This paper, and also the modern literature on growth optimality, looks at the positive question of which strategies survive in a wealth evolutionary dynamical system. In this situation the growth optimal portfolio plays an important role. It is often the strategy which survives in the long run.⁴ In different situations, for different preferences, it may or may not be a utility maximizing strategy. This paper contributes to this question by looking at the case where returns are predictable.

This paper uses a simple example to see if these wealth selection biases might be relevant to actual asset pricing and investor behavior. This is done in the context of a model where asset returns are predictable. It is based on Campbell & Viceira (1999), and more recently Pastor & Stambaugh (2009), where asset returns contain some predictability. The basic result is that wealth growth can select trading strategies far from the optimal predictive rules. Fitness objective functions using utility maximizing objectives will always chose the true forecast parameters, as they should.

After briefly presenting the model, the paper will test the benchmark case of log utility. Then log traders will be replaced by more risk averse traders, and finally some experiments will be

¹ The early comments on this are in Alchian (1950) and Friedman (1953). Another important approach to evolution in economics is Nelson & Winter (1982). Hodgson (1993) is a still a nice general overview with many classical references.

² Momentum strategies can generally be classified as short range trend following strategies. LeBaron (2000) shows that trend following and momentum strategies are very similar. A recent survey of this large area is in Swinkels (2004).

³ This is known as the growth optimal portfolio. See Samuelson (1971) and Hakansson (1971) for the original debate. Also, Kelley (1956) and Breiman (1961) provide the theoretical foundations.

⁴ Various theoretical papers have reached similar conclusions in different frameworks. These include Blume & Easley (1990) and Blume & Easley (2006) which analyze utility maximizing strategies with prices set endogenously. The latter 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 market 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. This is also related to the survival of noise traders as in DeLong, Shleifer, Summers & Waldmann (1991). In a slightly different setup Kogan, Ross & Wang (2006) show that even when irrational traders' wealth goes to zero, their price impact does not. The model considered here differs from this in that trader price impact is ignored.

performed which examine convergence speeds.

2 Results

2.1 Model structure

The economy considered here is a partial equilibrium one where security prices are set exogenously, and are not influenced by changes in wealth. There are two assets in the market. A risk free asset which pays a fixed return, and a risky asset paying a stochastic return with a small predictable component. Returns will be generated at a weekly frequency, and all portfolio rebalancing decisions will be made on a weekly basis.

The parameters are calibrated to well known results from financial markets to look reasonable.⁵ The risk free return is given as R_f , with $r_f = \log(1 + R_f)$. The return on the single risky asset is given by R_t with $r_t = \log(1 + R_t)$.

The dynamics of r_t are given by

$$r_{t+1} = x_{t+1} + e_{t+1} \quad (1)$$

$$x_{t+1} = \mu + \rho(x_t - \mu) + \eta_{t+1}. \quad (2)$$

This representation follows Campbell & Viceira (1999), and is a reasonable benchmark for financial returns series showing some amount of predictability.

Certain aspects of the stochastic structure of r_t will be important for the framework. Both noise shocks, e_t and η_t , will be normally distributed, and are homoskedastic with variances given by σ_e and σ_η . The annualized values of these are given in table 1. Two other important features will be used in choosing parameters. First, the signal to noise ratio in returns series is small. Predictive regressions run at the annual frequency generally yield very small R^2 values, usually between 0 and 10 percent. Reflecting this, the parameters are set so that the variance of x_t is 2 percent of the total return variance at the weekly frequency. Table 2 reports a monte-carlo simulation of the return process, showing autocorrelations at frequencies of 1 to 4 weeks, and the R^2 for a simulated one year predictive regression. Values in parenthesis are standard deviations across 1000 monte-carlo runs. The relatively short sample lengths are chosen to correspond to those available in many financial time series. The annual prediction experiment assumes the investor knows x_t and regresses the next year's return on the current value. The simulations produce R^2 estimates which are approximately 10 percent with very large dispersion across the simulated cross section. We should expect these numbers to have a slight upward bias due to the fact that in this experiment it is assumed that investors know the value of x_t .

The value of ρ is set to 0.95. This represents the large persistence believed to characterize many predictor variables. For example, Campbell & Viceira (2002) report a value of 0.957 for an estimate of the quarterly impact of lagged dividend price ratios on current ones. The value of 0.95

⁵See Campbell & Viceira (2002) for many examples.

is probably slightly too small for weekly persistence, but there are several reasons for choosing this. First, x_t doesn't exactly represent dividend/price ratios, but is a stand-in for many different predictors. Second, the value of 0.95 is useful in the experiments to see if agents are able to discern between a stationary, and a nonstationary process for x_t . As an initial test, it seems reasonable to move this parameter farther away from 1.

The experiments performed in this paper will concentrate on the evolution of wealth shares across traders. The objective is to find out in a pool of noninteracting strategies with an exogenous returns process, who in the end is left standing through simple compounding of wealth onto successful dynamic portfolios. Agent i 's strategy each period will be to invest $\alpha_{t,i}$ fraction of wealth in the risky asset, and $1 - \alpha_{t,i}$ fraction in the risk free. The portfolio return from t to $t + 1$ is therefore,

$$R_{t+1,i}^p = \alpha_{t,i}R_{t+1} + (1 - \alpha_{t,i})R_f. \quad (3)$$

The wealth share of agent i follows,

$$w_{t+1,i} = \frac{w_{t,i}(1 + R_{t+1,i}^p)}{\sum_{j=1}^N w_{t,j}(1 + R_{t+1,j}^p)}. \quad (4)$$

The dynamics of wealth depends on the realized distribution of returns, wealth shares at period t , and the portfolio strategies at period t , $\alpha_{t,i}$.⁶

Portfolio choice in the model is determined by a simple myopic power utility function in future wealth. Campbell & Viceira (2002) show that for log normal returns the fraction of wealth in the risky asset is approximated by

$$\alpha_{t,i} = \frac{E_t^i(r_{t+1}) - r_f + \sigma_t^2/2}{\gamma\sigma_t^2}. \quad (5)$$

The portfolio weights will be restricted to $-0.5 < \alpha_{t,i} < 2$. This allows for some short selling, and some leverage. This range is designed to replicate a moderate risk taking hedge fund more than the average individual investor.⁷

Optimal return forecasts in this state space world are given by the usual Kalman filter that gives a forecast of

$$\hat{x}_{t+1|t} = \mu + \rho_i(\hat{x}_{t|t-1} - \mu) + \omega_i(r_t - \hat{x}_{t|t-1}), \quad (6)$$

where $\hat{x}_{t+1|t}$ is the forecast of x_{t+1} given time t information. ρ_i is the memory parameter, and ω_i is closely related to the Kalman gain parameter.⁸ The basic experiments will race investment strategies in terms of relative wealth across different values of the (gain, memory) pair, (ω_i, ρ_i) . Given the parameters from table 1 the optimal forecast parameter pair is given by is $(\omega_*, \rho^*) =$

⁶If agents were assumed to consume a constant fraction of wealth, λ , each period, the above relative wealth relations would not change.

⁷LeBaron (2007) reports results which are robust across several other portfolio bounds.

⁸ See LeBaron (2007) for a detailed derivation.

(0.0164, 0.95). This is a target that will be used when examining wealth distributions in the model simulations. Finally, σ_t^2 will be set to σ_r^2 . This is a reasonable approximation given that the signal to noise ratio in the model is small.

2.2 Log utility

This section focusses on the case where all agents have $\gamma = 1$ or log preferences. Agents follow adaptive forecasting rules as in equation 6. They will be heterogeneous both in terms of the gain and memory parameters. The objective is to see how well wealth is drawn to the optimal forecasting parameters determined by the Kalman filter.

The two panels of figure 1 show the long run properties for different forecasting parameters. The parameter pairs are organized on a grid by the memory and gain parameters for the adaptive forecasts. Memory parameters vary from 0.9 to 1.0 incremented by 0.01. Gain parameters vary from 0 to 0.1 incremented by 0.01.⁹ This gives a 11 by 13 grid for a total of 143 rules. The dark lines mark the optimal forecast parameters, or strategies formed by using the true conditional expectation. The gain levels at zero are an important value to keep track of since this corresponds to ignoring any new return information in the forecast, and using a constant forecast set to the unconditional mean return.

The top panel of figure 1 shows the cross sectional mean of wealth fractions over 100 runs recorded after 500 years. The contour height (displayed on the right legend) is in units of density divided by uniform density ($1/143 = 0.007$). For example, a contour level of 2 indicates that the corresponding rule has twice the wealth density it would have under a uniform wealth distribution. The figure shows a clear long run concentration on the optimal forecasting rule. This indicates that for the $\gamma = 1$ investor wealth will concentrate on the true conditional expectation based trading rule in the long run.

The lower panel of figure 1 shows the estimated expected utility for the different strategies. This is estimated with the time series mean taken over the 500 years and over the 100 cross section. It is reported as an annual certainty equivalent return. This is estimated as,

$$\log(1 + r_p^*) = E(\log(1 + r_{p,t})), \quad (7)$$

where the above expectation is estimated by taking both time and cross sectional means. It shows a region around the optimal rule with a annual certainty equivalent of about 10 percent. This drops off as the distance to the optimum increases. It is interesting to note that the drop off is very steep as one moves to the constant forecasting rules on the left side of the panel where the gain is zero. This indicates the economic usefulness of the adaptive strategy for the $\gamma = 1$ investor. There is also an asymmetry in the shape in that utility levels are not all that sensitive to reductions in the memory parameter. Dropping the memory parameter to 0.9 has only a small impact on the

⁹ Extra points are added at 0.005, and 0.0164. Both give a finer grid near zero, and the later makes sure that the true optimal forecast parameters are in the grid.

certainty equivalence level.

2.3 Risk aversion ($\gamma = 3$)

The previous section demonstrated that wealth does converge to the optimal forecasting parameters for the $\gamma = 1$ (log) investor as predicted by theory. Unfortunately, it is not at all clear this is a good level of risk aversion for actual investors.¹⁰ It is generally felt that $\gamma = 1$ is somewhat low. Economists and financial advisors often use a wide range of values of $\gamma > 1$. All these higher levels of risk aversion are probably a much better approximation to the actual population than $\gamma = 1$.

Figure 2 repeats the two panel plot of long run wealth and expected utility levels for $\gamma = 3$. The upper panel has changed dramatically. Wealth is no longer converging to the optimal forecast parameters. The wealth density is maximized at a (gain, memory) pair of (0.06, 1.00), well off the optimum values which are again marked by the dark lines. The lower panel reports the utility levels as annual certainty equivalents which are given by,

$$(1 + r_p^*)^{1-\gamma} = E((1 + r_{p,t})^{1-\gamma}) \quad (8)$$

This shows clearly that in utility terms, the optimal forecast parameters are utility maximizing for investors. The maximum certainty equivalent return is estimated as 5.25 percent per year at the optimal forecast parameters. It is only 2.91 at the wealth maximizing parameters, which is a loss of almost 2.5 percent, or 50 percent of the certainty equivalent return. The utility surface is again relatively flat in the the memory parameter with little change coming from reducing this to 0.90. The surface is steep in the direction of reducing the gain parameter to zero which again corresponds to moving to a constant weight portfolio.

Figure 3 shows the dynamics of the wealth distributions over time. As in the previous figures these are again cross sectional means over 100 different runs. The pattern is interesting in that the movement away from the optimal values begins early. Even at 5 years, the wealth density is drifting to the northeast corner. By 20 years there is a pronounced large bias in the gain parameter. This continues at 50 and 200 years.

3 Conclusions

This paper has shown that wealth evolution alone can converge to forecasting strategies with distorted beliefs. These results were predicted by theoretical work on strategy evolution, but the results here are performed using a familiar forecasting setup calibrated to known patterns in financial time series.

The most important qualitative aspect of the forecast distortion is the selection of gain parameters which are well above the optimal value. This translates into wealth concentrating on

¹⁰ For example, see Bliss & Panigirtzoglou (2004), and references therein.

strategies which put too much weight on current returns in their forecasts. Such strategies could correspond to the presence of momentum and trend following strategies in actual markets.

The experiments also confirm that the true forecast parameters would be selected under a utility maximization objective. In financial markets subject to learning, gradients carrying wealth to both the these maxima may exist. The wealth maximizing point is obvious, as it is driven by the relative success of certain strategies in terms of wealth growth. Almost any sensible model of wealth dynamics should have some aspect of this wealth dynamic built in. The utility maximization point would attract wealth under some form of active learning, where agents actively shift strategies onto those estimated to maximize utility.¹¹ In real markets there may be a tension in terms of the accumulation of wealth between these two forms of learning.

Several extensions to this model would appear to be important. First, experiments will need to follow the theoretical literature, and much of the computational literature, and endogenize prices.¹² One can always question whether the convergence results given here would be affected by price changes as wealth moves around. This paper deliberately eliminated this effect, but in the future it has to be part of the analysis of evolution and financial markets. A second simplification which may have a large impact was the assumption that the innovations to the expected return process and the return noise process are independent. This deviates from some of the forecasting evidence and many of the prediction models in use. It also has a nontrivial impact on the structure of the Kalman filter forecasting system.¹³ If there is a large enough negative correlation then a large recent returns could have a negative impact on predicted future expected returns, and therefore would warrant a negative gain parameter in this system. Implementing a richer evolutionary system which better addresses this issue is another important extension.

The basic point of this paper is simple. Wealth evolution, on its own, will not reliably perform the function that it is often assumed to do. This obviously forces some difficult choices for researchers building heterogeneous adaptive models. If optimal forecasts hold any sway as a target markets might be tending toward, then this movement must be coming from the active learning side. Worse, the passive learning side will be slowly, and steadily working against this drift. Unfortunately, modeling active learning is much more difficult than passive learning. These results suggest that there will always remain some amount of wealth concentrated on strategies that would be difficult to explain from the standpoint of optimal forecasting. Understanding exactly what one would expect this wealth distribution to look like as it moves through time, and its impact on prices remains an interesting question.

¹¹ See LeBaron (2011 forthcoming) for further discussions.

¹² See LeBaron (2010) for an example with some similar strategies, and a rich trading environment with endogenous pricing, and agent survival.

¹³ See Pastor & Stambaugh (2009) for discussion.

References

- Alchian, A. (1950), 'Uncertainty, evolution, and economic theory', *Journal of Political Economy* **58**, 211–221.
- Bliss, R. R. & Panigirtzoglou, N. (2004), 'Option-implied risk aversion estimates', *Journal of Finance* **59**(1), 407–446.
- 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.
- 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.
- Campbell, J. Y. & Viceira, L. M. (1999), 'Consumption and portfolio decisions when expected returns are time varying', *Quarterly Journal of Economics* **114**, 433–495.
- Campbell, J. Y. & Viceira, L. M. (2002), *Strategic Asset Allocation*, Oxford University Press, Oxford, UK.
- DeLong, J. B., Shleifer, A., Summers, L. H. & Waldmann, R. (1991), 'The survival of noise traders in financial markets', *Journal of Business* **64**, 1–19.
- Evstigneev, I. V., Hens, T. & Schenk-Hoppe, K. R. (2006), 'Evolutionary stable stock markets', *Economic Theory* **27**, 449–468.
- Friedman, M. (1953), *Essays in Positive Economics*, University of Chicago Press, Chicago, IL.
- Hakansson, N. H. (1971), 'Multi-period mean-variance analysis: Toward a general theory of portfolio choice', *Journal of Finance* **26**(4), 857–884.
- Hodgson, G. M. (1993), *Economics and Evolution: Bringing life back into Economics*, University of Michigan Press, Ann Arbor, MI.
- Kelley, J. L. (1956), 'A new interpretation of information rate', *Bell Systems Technical Journal* **35**, 917–926.
- Kogan, L., Ross, S. A. & Wang, J. (2006), 'The price impact and survival of irrational traders', *Journal of Finance* **61**, 195–229.
- LeBaron, B. (2000), 'The stability of moving average technical trading rules on the Dow Jones index', *Derivatives Use, Trading, and Regulation* **5**(4), 324–338.

- LeBaron, B. (2007), Wealth evolution and distorted financial forecasts, Technical report, International Business School, Brandeis University.
- LeBaron, B. (2010), Heterogeneous gain learning and the dynamics of asset prices, Technical report, International Business School, Brandeis University.
- LeBaron, B. (2011 forthcoming), 'Active and passive learning in agent-based financial markets', *Eastern Economic Journal* .
- Nelson, R. & Winter, S. (1982), *An evolutionary theory of economic change*, Harvard University Press, Cambridge, MA.
- Pastor, L. & Stambaugh, R. F. (2009), 'Predictive systems: Living with imperfect predictors', *Journal of Finance* **64**, 1583–1628.
- 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.
- Swinkels, L. A. P. (2004), 'Momentum investing: Survey', *Journal of Asset Management* **5**(2), 120–143.

Table 1: *Return Parameter Values*

Parameter	Value
r_f	0.02
$E(r_t)$	0.07
σ_r	0.20
σ_x^2/σ_r^2	0.02
ρ	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. σ_r is the corresponding annual standard deviation. σ_x^2/σ_r^2 is the signal to noise ratio in the returns series. ρ is the AR(1) persistence parameter for the expected return process.

Table 2: *Return Time Series Simulations*

Sample	ρ_1	ρ_2	ρ_3	ρ_4	R^2
25 Years	0.018 (0.028)	0.013 (0.028)	0.016 (0.029)	0.015 (0.028)	0.110 (0.106)
50 Years	0.019 (0.020)	0.017 (0.020)	0.017 (0.020)	0.015 (0.020)	0.100 (0.075)

Description: Mean values from 1000 monte-carlo return simulations corresponding to 25 and 50 years. ρ_j is the return autocorrelation at j week lag. R^2 is the R^2 of a annual regression of year $t+1$ returns on x_t , the expected return, at the end of year t . Numbers in parenthesis are the standard deviations of these estimated values from the 1000 length cross section.




Figure 1: **Wealth and utility surfaces for $\gamma = 1$**

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. Both maximums correspond to the optimal Kalman forecast parameters of (0.016, 0.950). The maximum certainty equivalent return is 0.100, or 10 per cent per year.




Figure 2: **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.




Figure 3: **Wealth surfaces for $\gamma = 3$: Time evolution**

This set of 4 figures displays the time evolution of the cross sectional averages of wealth distributions. Distributions are sampled at the 5, 20, 50, and 200 year time periods.