

Dynamic Risk Avoidance and the Attraction of Short Memory Volatility Forecasts

Blake LeBaron *

International Business School

Brandeis University

Preliminary: June 20, 2014

Abstract

This paper presents empirical evidence on the feasibility and value of various simple volatility forecasts in the construction of dynamic portfolio rules. Rules using limited past data are compared with rules using longer series in risk estimation. It is shown that in terms of performance, a utility based objective across these rules favors short memory/high gain forecasting rules. This empirical result is connected to recent work in agent-based finance which suggests that short memory volatility sensitive traders may be very important to the dynamics of long swings in asset prices around fundamentals.

Keywords: Learning, Asset Pricing, Financial Time Series, Evolution, Memory

*415 South Street, Mailstop 32, Waltham, MA 02453 - 2728, blebaron@brandeis.edu, www.brandeis.edu/~blebaron. This research is partially supported by funding from the Institute for New Economic Thinking, INET.

1 Introduction

Over thirty years ago Shiller (Shiller 1981) demonstrated that aggregate asset prices moved too much to be explained by their aggregate fundamentals (either earnings or dividends). Since then this result has been rethought, and replicated many times over. It is directly connected to long range return predictability coming from regressions on price-dividend or price-earnings ratios¹. It also is related to cross sectional features in asset pricing, as well as features in the term structure and foreign exchange markets.² Finally, the fact that asset prices may appear riskier than their underlying fundamentals is part of the equity premium puzzle as well. Most of the major puzzles in modern macro finance lead one back to the original excess volatility issue.

Many explanations have been explored involving varying degrees of rationality assumed for economic actors. Because asset prices involve some form of discounted cash flow, all models explaining price volatility must deal with the movements of agents' expectations of future cash flows (the numerator) and/or some measure of risk (the denominator). Explanations involving expected growth rates in cash flows have been tricky given that many of these series show little predictability. One possibility is the long run risk model which assumes that changes in growth rates are highly persistent and obscured by noise (a slow moving trend).³ Learning about trends in cash flows is another possibility. Representative agent examples of this are Adam & Marcet (2010) and Fuster, Hebert & Laibson (2011). A radically different idea is proposed in Boswijk, Hommes & Manzan (2007) and Chiarella, He & Zwinkels (2009) which build expectations from two very different types of heterogeneous agents. In these models the continuing use of trend following beliefs leads to large swings in asset prices.⁴ Another possibility is that aggregate uncertainty is changing, Bansal, Kiku, Shaliastovich & Yaron (forthcoming 2014), or risk aversion is changing either through a persistent habit mechanism as in Campbell & Cochrane (1999) or directly through taste shocks as in Albuquerque, Eichenbaum & Rebelo (2012).

This paper approaches this puzzle from the standpoint of risk perceptions. It is argued that agent expectations may be overly sensitive to changes in recent volatility. This risk perception can lead to very large swings in asset prices.⁵ Further, the price dynamics themselves can make these short run forecast-

¹ See Campbell & Shiller (1988), and then later Cochrane (2008), and Lettau & Ludvigson (2010).

² See Cochrane (2011) for some of the connections between time series and cross sectional features.

³ See Bansal & Yaron (2004).

⁴ This is common to many agent-based financial models. However, relatively few of them are calibrated or estimated to longer range swings in asset prices. See several recent surveys including, Chiarella, Dieci & He (2009), Hommes & Wagener (2009), and Lux (2009).

⁵ Another model that emphasizes learning about risk as well as return is Branch & Evans (2011).

ing rules very appealing, and even quite rational for some agents. This has been demonstrated in several agent-based asset pricing models.⁶ The objective of this paper is to take this result to the data to see if the magnitudes line up with those from the simulations. How much are agents gaining or losing from the use of short range forecasting rules, and do the results agree with those from the simulations? Finally, can the empirical features help to calibrate the values in the agent-based models, and help us select between these various competing, but very different, models for excess volatility.⁷

This study does not intend to push the limits of the literature on forecasting returns or volatility, but it obviously has some overlap with this vast body of research.⁸ A critical paper in this area is Goyal & Welch (2003) which questions the out of sample performance of many predictive forecasting rules. This paper takes some of the spirit of their 2003 study by concentrating only on variance forecasting, and ignoring predictions of expected returns.⁹ In the time series sense, the order of magnitude of volatility predictability far exceeds that for returns. There is a huge literature on volatility forecasting, but only a relatively small part of it is concerned with whether volatility forecasting alone can lead to actual economic gains for traders.¹⁰ There is also some connection to recent work on realized volatility since all the volatility estimators will be built from monthly variance estimates using daily data.¹¹ Once the question of volatility forecasting is opened, then one is also concerned with construction of dynamic portfolios. This paper remains relatively simple and agnostic on this issue, following mostly the framework set out in Campbell & Viceira (2002), but many more sophisticated approaches have been proposed.¹² In summary, this paper does not push the technology on building sophisticated forecasting models, but it does show important results for very simple rules, and comparisons across these rules. It is often more interested in gaining insights into the gradients in a simple forecasting space on which learning models will operate than the actual magnitudes of the forecasts per se.

Section 1 details the empirical methodology, including the construction of a new longer horizon data set

⁶See LeBaron (2013a) for examples.

⁷It is tempting to conjecture that an over use of short memory, high gain forecasts is all that you need to generate realistic asset price dynamics. LeBaron (2013b) presents some initial evidence this is not the case, and that it is the interaction between forecasters with different memories that is essential in getting the dynamics right.

⁸An extensive survey is Lettau & Ludvigson (2010).

⁹The debate on predictability is far from over. Recently, Pettenuzzo, Timmerman & Valkanov (2014 forthcoming) have shown that imposing reasonable priors can yield significant out of sample forecasts. This follows on an earlier study, Campbell & Thomson (2008), that suggested reasonable model restrictions might be important.

¹⁰See West, Edison & Cho (1993), Flemming, Kirby & Ostdiek (2001) for early examples, and more recently, Engle & Kelly (2011). Tang & Whitelaw (2011) is also related in its methods for determining the relative contributions of return versus volatility forecasts.

¹¹Andersen, Bollerslev, Diebold & Labys (2003) and Barndorff-Nielsen & Shephard (2004) are two early examples. Barndorff-Nielsen & Shephard (2010) is a recent survey, and Andersen, Bollerslev, Christoffersen & Diebold (2006) surveys many forms of volatility forecasts. However, the technology used in this paper goes back to French, Schwert & Stambaugh (1987). A paper combining realized volatility with variance/covariance portfolio construction is Flemming, Kirby & Ostdiek (2003).

¹²For examples see Cenesizoglu & Timmermann (2012), Johannes, Korteweg & Polson (2014), and Pettenuzzo & Ravazzolo (2014).

using several existing series. Section 2 presents the results, and section 4 concludes.

2 Empirical methodology

The time series used in this study involve several readily available return series from U.S. equity and bond markets. Several series are pasted together to increase the sample size. All series begin in March 1885 and continue through December 2013. For the period from 1926 through 2013 the CRSP (Center for Research in Securities Prices) series are used. Stock returns come from the value weighted series including dividends, at both the daily and monthly frequency. The risk free series is the CRSP short range TBill returns series.

Monthly realized volatility estimates are estimated using daily data, and the methodology from French et al. (1987). An estimate for each month is constructed as,

$$v_t = \sum_{i=1}^m r_{t,i}^2 + 2 \sum_{i=2}^m r_{t,i} r_{t,i-1} - \left(\frac{1}{T} \sum_{i=1}^T r_t \right)^2 \quad (1)$$

where $r_{t,i}$ is the logged daily return for month t , and day of the month i . The variance estimate is adjusted for the small amount of correlation in daily return series, and uses the estimated sample mean from the logged monthly returns, r_t .

The period from 1885-1915 is trickier for data construction. The data sets from Schwert (1990) are used here. They consist of value weighted indices (including dividends) measured at both the daily and monthly frequencies. Monthly volatility estimates are constructed exactly as done for the CRSP series. The risk free rate is set to the New York commercial paper series from the FRED historical archive. This monthly yield series goes back to 1857. It is not a government bond series, but this series matches well the short end of the term structure captured by the CRSP Tbill series. For all time periods inflation is estimated using the CPI series from the Shiller monthly data set available at his website.¹³

This is then used to construct a family of smoothed volatility estimators. A forecast for volatility at time $t + 1$ is constructed from constant and decreasing gain estimators applied to the monthly volatility series. The constant gain estimate is given by,

$$E_t(v_{t+1}) = \sigma_{t,j}^2 = (1 - \lambda_j) \sigma_{t-1,j}^2 + \lambda_j v_t, \quad (2)$$

where the parameter λ_j controls how far back into the past the estimator reaches, or the filter gain. The

¹³For information on Shiller's series see Shiller (2000).

functional form is similar to the forecast generated by a GARCH(1,1), and corresponds directly to the Risk-metrics methods for estimating conditional variances used in various risk management systems. These later estimators are based on the same simple exponential filter. Decreasing gain estimates follow,

$$\sigma_t^2 = +\frac{t-1}{t}\sigma_{t-1}^2 + \frac{1}{t}v_t, \quad (3)$$

which corresponds to an estimate of the mean volatility over the period $[1, t]$.

These estimates are then used in a standard portfolio construction framework based on linearized Constant Relative Risk Aversion (CRRA) preferences,¹⁴

$$\alpha_{t,j} = \frac{E(r_{e,t+1} - r_{f,t+1}) + \frac{1}{2}\sigma_{t,j}^2}{\gamma\sigma_{t,j}^2}, \quad (4)$$

where $\alpha_{t,j}$ is the fraction of wealth to invest in the risky asset for volatility rule j , and γ is the coefficient of relatively risk aversion. The excess return of equity over the risk free, $r_{e,t+1} - r_{f,t+1}$ is estimated unconditionally over the entire sample. All portfolio variability comes entirely from changes in conditional variance forecasts. Portfolio returns are generated using the weights as

$$R_{p,j,t+1} = \alpha_{t,j}R_{e,t+1} + (1 - \alpha_{t,j})R_{f,t+1} - \pi_{t+1} \quad (5)$$

where π_{t+1} is the expost inflation over the period $t, t+1$.

Various moments can be estimated for the portfolio returns. The most important of these is the certainty equivalent return given by

$$(1 + R_{ce})^{1-\gamma} = E((1 + R_p)^{1-\gamma}), \quad (6)$$

and

$$\log(1 + R_{ce}) = E(\log(1 + R_p)), \quad (7)$$

for $\gamma = 1$. This corresponds to the optimal portfolio choice of the agent and is a standard risk adjusted measure for the utility value of the strategy. Many agent based models will use something close to this as a rule fitness value, so it makes sense here as a reasonable representation of a standard agent-based adaptive learning mechanism.

¹⁴See Campbell & Viceira (2002) for derivations.

3 Empirical results

Results are presented using the sample of monthly returns, volatility, and risk free rates starting in March 1885 and ending in December 2013. For tables reporting trading strategy results the sample starts in March 1905 to give the strategy some initial data to start the learning rules. Table 1 presents summary statistics on the monthly return series, both nominal and real. There are few surprising results. Real returns in stocks are 7.9 percent with a standard deviation of 18 percent. The real risk free TBill returns are 1.2 percent. Monthly stock returns show the expected strong evidence for excess kurtosis. Figure 1 displays the entire monthly return series graphed along with the monthly realized volatility. It is obvious that increases in monthly volatility are reflected in large, usually persistent, increases in volatility.

The first set of certainty equivalent return estimates are given in table 2. These correspond to a maximum α of 1 with no leverage allowed. The table gives results for several different levels of risk aversion, γ , and 5 different forecasting rules. All values are presented as annualized returns. “Decreasing” refers to the decreasing gain rule. This is equivalent to estimating the volatility from the monthly mean over the first t sets of months for use in the $t + 1$ portfolio. It serves as a useful benchmark rule for an agent not trying to adaptively adjust their volatility estimate.¹⁵

The first row of the table corresponds to log utility ($\gamma = 1$) and displays little difference across any of the forecasting rules. Risk aversion for these agents is so low that most of the agents are at the constraint and hold 100 percent equity in their portfolios, yielding little difference across any of the strategies. The row below this with values in parenthesis implements a bootstrap comparison of each of the constant gain forecasts with the decreasing gain forecast. Pairs of the portfolio returns are drawn at random, but maintaining the temporal dependence for each time pair. The value reports the fraction of bootstrap runs where the constant gain forecast generates a larger certainty equivalent return. This value is important since this is what one would be interested in an adaptive learning environment. It is essentially the probability of switching to the constant gain forecasting rule from a rule using a very long memory forecast. For $\gamma = 1$ these values are all relatively small. Only the rule corresponding to a 1 year half life, shows any major chance of change 0.34.

Once γ is increased to 2 the situation changes. Now certainty equivalent returns show larger values for relatively short gains (less than 10 years). The estimated probability of getting a larger adjusted return for a 1 year half life rule has increased to 0.87. A reasonable value for γ in terms of generating realistic asset

¹⁵In other experiments the unconditional mean of the volatility from the entire sample is used in place of the decreasing gain rule. This gives an obvious out of sample advantage for this rule, but had little impact on the results.

pricing dynamics is often in the range of 3 to 3.5.¹⁶ Now the improvement from moving to the one year constant gain volatility forecast is much more dramatic. The certainty equivalent return increases from 2.7 percent to 3.9 percent in annual risk adjusted terms. The probability that the high gain forecast improves on the constant gain rule is now 0.99. Certainty equivalent values fall as one moves from the 1 year rule out to the 25 year rule, but several of the higher gain (shorter memory) rules still show large improvements over the decreasing gain forecast. The results for $\gamma = 5$ repeat those for $\gamma = 3$. The potential utility draw of large gain volatility forecasting rules remains, with over 96 percent of bootstrapped paired sample showing a gain for the 1 year half life rule.

The certainty equivalent measure is probably the best measure to be checked in terms of rule usefulness, but it is not the end of the story in terms of performance. Since these rules are estimated using variance based forecasts it might be interesting to explore how much of the performance gain is coming from the mean versus the variance of the dynamic portfolio. Table 3 compares the expected returns for the various dynamic strategies. For $\gamma = 1$ and $\gamma = 2$ the table shows little difference which is what should be expected. However, for $\gamma = 3, 5$ there are increases in conditional means. They are really only large for the shortest memory, 1 year half life rule. These need to be directly compared with the results for the portfolio standard deviation in table 4. This table shows a reduction in variance for the more risk loving agents, but an increase in variance for $\gamma = 5$. In this unusual case the gain in performance appears to be coming more from the change in mean than variance, which is interesting given that this is a variance forecasting rule.

Table 5 presents the annualized Sharpe ratios for the various strategies. These again show improvements in moving from the decreasing gain strategy to the 1 year half life strategy for all levels of $\gamma > 1$. In the case of Sharpe ratios the strategies clearly begin to converge back to the decreasing gain levels for longer half life rules. The column for the 25 year rule clearly shows this. Of course, Sharpe ratios are not really a very good measure in any of these cases since the strategy is not designed to maximize this, but it reiterates the general attractiveness of a high gain variance forecasting rule.

The previous tables have all limited α to be less than or equal to 1. In table 6 this is relaxed, and the upper bound on α_t is increased to 2. Results in the table show that the improvements in certainty equivalent returns are not driven by the portfolio restriction. Large improvements in moving from decreasing to constant gain learning are observed for all levels of risk aversion. For example, at $\gamma = 3$ the annualized certainty equivalent return increases from 2.7 to 3.9 percent. Even for $\gamma = 1$ the improvement is still present,

¹⁶ See LeBaron (2012) for examples.

and the probability of drawing a constant gain rule that improves on the decreasing gain rule is 0.86.¹⁷

Results have up to now supported the argument that agents would actively select a simple, but high gain, active variance forecasting rule over constant gain, or other longer range forecasts. However, the constant gain rule may be far from the optimal forecasting rule for volatility. A complete search over all possible forecasts would be infeasible, but a very simple comparison is performed. A low order autoregressive (AR) model is fit to the volatility series, and it will be used to forecast future variances. Using the Ljung-Box residual Q-test identifies an AR(3) model as in,

$$\sigma_t^2 = \alpha + \sum_{j=1}^3 \beta_j \sigma_{t-j}^2 + \beta_4 r_{t-1}.$$

This model is also augmented with the well known “leverage” term to account for volatility in month $t + 1$ increasing when prices fall in month t .¹⁸ Also, a simple AR(1) model is used for comparison. All parameters are estimated over the full sample, with the intention of giving them an unfair advantage over the other models. The certainty equivalent values are presented for the 3 AR models in table 7. They show some small improvements over the decreasing gain rule, but little difference across the 3 rules themselves. These should be compared to results in 2. For $\gamma = 3$ the 1 year half life rule generated a certainty equivalent return of 3.9 percent, which can be compared to a value of 3.5 percent for the AR(3) rules. For $\gamma = 5$ the same comparison shows a certainty equivalent return of 2.6 percent in table 2 which is exactly the same as the AR(3) with leverage in table 7. The initial evidence suggests that there will be no dramatic improvements in certainty equivalent returns from moving to more sophisticated models.¹⁹

Table 8 compares the in sample mean squared errors for the various forecasts. There should be no question about which forecast should win in this race.²⁰ The models perform as expected with smallest MSE’s going to the more complex and sophisticated AR models. However, the gains over the 1 year half life forecast do not appear large in MSE improvement. The MSE falls from 0.22 to 0.19 in going to the AR model with leverage. This repeats the result that the constant gain rules may be relatively simple and robust rules which work well in terms of a dynamic trading strategy.

As a final casual assessment of the high gain forecasting rule its dynamics are compared to the level

¹⁷It is important to note that the results in table 6 are not relevant to agent-based models such as LeBaron (2012) and LeBaron (2013a) since both of these papers impose a no leverage restriction.

¹⁸See Hasanhodzic & Lo (2011) and originally Christie (1982).

¹⁹Obviously, this is operating in the family of models which keep the expected excess return forecasts constant. Also, these are all trivial models given what is possible in terms of forecasting. Recent examples of much more sophisticated models are in Johannes et al. (2014) and (Pettenuzzo et al. 2014 forthcoming).

²⁰Obviously, one should question the potential in sample bias in the AR MSE’s. A fair test would estimate these values out of sample, but this increase in MSE would only strengthen the result that the performance comparison is very close across rules.

of the VIX index. As is well known the VIX is an option based predictor of next month's volatility on the aggregate S&P500 index. Figure 2 compares the dynamics of the 1 year constant gain forecast to the VIX in units of annual standard deviation. The general level of variability from lows of near 10, to highs around 30 are similar except for the unusual period in 2008 when the VIX get close to 60. Obviously, the constant gain forecast is much smoother than the monthly observations of the VIX, but the general patterns are similar. If overall perceptions of risk change by nearly a factor of 4 over time (16 for the variance), then it would seem likely that this dynamics may play a big role in the changes in aggregate asset prices.

4 Summary

This paper tested a simple trading strategy for monthly equity returns. The objective is not so much to demonstrate that this rule, which puts a heavy weight on recent volatility, works well, but to show that its performance is at least comparable to rules using much longer memory. Empirically, the evidence suggests that these simple high gain forecast rules do well when compared with lower gain, longer memory, rules. Therefore, adaptive learning mechanisms used in real markets would end up with rules that are sensitive to short term changes in volatility. This feature is a reliable component of heterogeneous gain learning models as in LeBaron (2012). In these simulated models much of wealth has an extreme short range sensitivity in volatility, and this sensitivity to endogenously generated changing variances plays a key role in generating long range persistent changes in price/dividend ratios.

As a mechanism, weighting recent volatility too heavily seems to pass a common sense test that agents would feel unusually scared about recent volatility in the market, and have a hard time in driving themselves to a longer range perspective for volatility estimation. It also contributes to potentially stabilizing strategies, using price/dividend ratios, being weakened as the traders get scared about moving into an undervalued, but volatile market. If the results had suggested a more balanced performance across all the rules, or an advantage to the decreasing gain rules, then this agent-based mechanism would need to be rethought. Given that the simple high gain rules did well, this type of volatility forecasting dynamics should be considered along with the many other mechanisms that have been proposed to explain long range swings of asset prices from their fundamental values.

References

- Adam, K. & Marcet, A. (2010), Booms and busts in asset prices, Technical Report 2010-E-2, Bank of Japan, Tokyo, Japan.
- Albuquerque, R., Eichenbaum, M. S. & Rebelo, S. (2012), Valuation risk and asset pricing, Technical Report 18617, National Bureau of Economic Research.
- Andersen, T. G., Bollerslev, T., Christoffersen, P. F. & Diebold, F. X. (2006), Volatility and correlation forecasting, in G. Elliott, C. W. J. Granger & A. Timmermann, eds, 'Handbook of Economic Forecasting', North Holland, pp. 778–878.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), 'Modeling and forecasting realized volatility', *Econometrica* **96**, 579–625.
- Bansal, R., Kiku, D., Shaliastovich, I. & Yaron, A. (forthcoming 2014), 'Volatility, the macroeconomy and asset prices', *Journal of Finance*.
- Bansal, R. & Yaron, A. (2004), 'Risks for the long run: A potential resolution of asset pricing puzzles', *Journal of Finance* **59**, 1481–1509.
- Barndorff-Nielsen, O. E. & Shephard, N. (2004), 'Power and bipower variation with stochastic volatility and jumps', *Journal of Financial Econometrics* **2**, 1–37.
- Barndorff-Nielsen, O. E. & Shephard, N. (2010), Measuring and modeling volatility, in R. Cont, ed., 'Encyclopedia of Quantitative Finance', John Wiley, pp. 1898–1901.
- Boswijk, H. P., Hommes, C. H. & Manzan, S. (2007), 'Behavioral heterogeneity in stock prices', *Journal of Economic Dynamics and Control* **31**(6), 1938–1970.
- Branch, W. A. & Evans, G. W. (2011), 'Learning about risk and return: A simple model of bubbles and crashes', *American Economic Journal: Macroeconomics* **3**, 159–191.
- Campbell, J. & Shiller, R. (1988), 'Stock prices, earnings, and expected dividends', *Journal of Finance* **63**, 661–676.
- Campbell, J. Y. & Cochrane, J. H. (1999), 'By force of habit: A consumption-based explanation of aggregate stock market behavior', *Journal of Political Economy* **107**, 205–251.

- Campbell, J. Y. & Thomson, S. B. (2008), 'Predicting excess stocks returns out of sample: Can anything beat the historical average', *Review of Financial Studies* **21**(1509-1531).
- Campbell, J. Y. & Viceira, L. M. (2002), *Strategic Asset Allocation*, Oxford University Press, Oxford, UK.
- Cenesizoglu, T. & Timmermann, A. (2012), Do return prediction models add economic value?, Technical report, HEC Montreal.
- 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. & Zwinkels, R. C. J. (2009), Heterogeneous expectations in asset pricing: Empirical evidence from the s&p 500, Technical report, University of Technology, Sydney, Australia.
- Christie, A. A. (1982), 'The stochastic behavior of common stock variances: Value, leverage and interest rate effects', *Journal of Financial Economics* **10**, 407–432.
- Cochrane, J. (2008), 'The dog that did not bark: A defense of return predictability', *Review of Financial Studies* **21**(4), 1533–1575.
- Cochrane, J. H. (2011), 'Discount rates', *Journal of Finance* **66**(4), 1047–1109.
- Engle, R. & Kelly, B. (2011), Dynamic equicorrelation, Technical report, University of Chicago, Booth School of Business.
- Flemming, J., Kirby, C. & Ostdiek, B. (2001), 'The economic value of volatility timing', *Journal of Finance* **56**(1), 329–352.
- Flemming, J., Kirby, C. & Ostdiek, B. (2003), 'The economic value of volatility timing using "realized" volatility', *Journal of Financial Economics* **67**, 473–509.
- French, K. R., Schwert, G. W. & Stambaugh, R. F. (1987), 'Expected stock returns and volatility', *Journal of Financial Economics* **19**, 3–29.
- Fuster, A., Hebert, B. & Laibson, D. (2011), Natural expectations, macroeconomic dynamics, and asset pricing, in D. Acemoglu & M. Woodford, eds, 'NBER Macroeconomics Annual', Vol. 26, University of Chicago Press, pp. 1–48.

- Goyal, A. & Welch, I. (2003), 'Predicting the equity premium with dividend ratios', *Management Science* **49**(5), 639–654.
- Hasanhodzic, J. & Lo, A. W. (2011), Black's leverage effect is not due to leverage, Technical report, Boston University.
- Hommes, C. H. & Wagener, F. (2009), Complex evolutionary systems in behavioral finance, in T. Hens & K. R. Schenk-Hoppe, eds, 'Handbook of Financial Markets: Dynamics and Evolution', North-Holland, pp. 217–276.
- Johannes, M., Korteweg, A. & Polson, N. (2014), 'Sequential learning, predictability, and optimal portfolio returns', *Journal of Finance* **69**(2), 611–643.
- LeBaron, B. (2012), 'Heterogeneous gain learning and the dynamics of asset prices', *Journal of Economic Behavior and Organization* **83**(3), 424–445.
- LeBaron, B. (2013a), Heterogeneous gain learning and long swings in asset prices, in R. Frydman & E. S. Phelps, eds, 'Rethinking Expectations: The Way Forward For Macroeconomics', Princeton University Press, pp. 169–206.
- LeBaron, B. (2013b), The impact of heterogeneous gain learning in an agent-based financial market, Technical report, International Business School, Brandeis University.
- Lettau, M. & Ludvigson, S. (2010), Measuring and modeling variation in the risk-return trade-off, in Y. Ait-Sahalia & L. P. Hansen, eds, 'Handbook of Financial Econometrics', Vol. 1, Elsevier, pp. 617–690.
- Lux, T. (2009), 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, pp. 161–215.
- Pettenuzzo, D. & Ravazzolo, F. (2014), Optimal portfolio choice under decision-based model combinations, Technical report, Department of Economics, Brandeis University.
- Pettenuzzo, D., Timmerman, A. & Valkanov, R. (2014 forthcoming), 'Forecasting stock returns with economic constraints', *Journal of Financial Economics*.
- Schwert, G. W. (1990), 'Indexes of United States stock prices from 1802-1987', *Journal of Business* **63**, 399–426.
- Shiller, R. J. (1981), 'Do stock prices move too much to be justified by subsequent changes in dividends?', *American Economic Review* **71**, 421–36.

Shiller, R. J. (2000), *Irrational Exuberance*, Princeton University Press, Princeton, NJ.

Tang, Y. & Whitelaw, R. F. (2011), 'Time varying Sharpe ratios and market timing', *Quarterly Journal of Finance* **1**(3), 465–493.

West, K. D., Edison, H. J. & Cho, D. (1993), 'A utility-based comparison of some models of exchange rate volatility', *Journal of International Economics* **35**(1-2), 23–45.

Table 1: Summary statistics

Series	Mean	Std.	Skewness	Kurtosis
VW index	0.105	0.179	0.275	11.7
VW index (real)	0.079	0.181	0.332	11.3
TBill	0.038	0.008	0.842	5.0
TBill (real)	0.012	0.033	−0.063	11.7

Summary statistics reported in annualized return units. Monthly CRSP VW returns + spliced with Schwert series from March 1885 through December 2013.

Table 2: Certainty equivalent

	Gain: half life years				
CRRA	Decreasing	1	5	10	25
$\gamma = 1$	0.060	0.059	0.060	0.060	0.060
		(0.344)	(0.000)	(0.000)	(0.000)
$\gamma = 2$	0.042	0.048	0.044	0.041	0.040
		(0.868)	(0.638)	(0.370)	(0.216)
$\gamma = 3$	0.027	0.039	0.033	0.030	0.028
		(0.988)	(0.922)	(0.820)	(0.664)
$\gamma = 5$	0.018	0.026	0.020	0.018	0.017
		(0.968)	(0.814)	(0.640)	(0.456)

Certainty equivalent returns (annualized). Estimated 1905-2013. Volatility forecasting rules use gain levels with given half lives measured in years. Decreasing corresponds to a decreasing gain forecast. Numbers in parenthesis are bootstrap estimates for the probability of the given rule generating a higher certainty equivalent than the decreasing gain rule.

Table 3: Annualized returns

	Gain: half life years				
CRRA	Decreasing	1	5	10	25
$\gamma = 1$	0.077	0.075	0.077	0.077	0.077
$\gamma = 2$	0.076	0.071	0.069	0.068	0.068
$\gamma = 3$	0.057	0.064	0.060	0.058	0.057
$\gamma = 5$	0.037	0.048	0.040	0.038	0.038

Table 4: Annualized standard deviation

	Gain: half life years				
CRRA	Decreasing	1	5	10	25
$\gamma = 1$	0.184	0.178	0.184	0.184	0.184
$\gamma = 2$	0.182	0.149	0.157	0.161	0.165
$\gamma = 3$	0.142	0.129	0.133	0.136	0.138
$\gamma = 5$	0.087	0.093	0.088	0.088	0.090

Table 5: Annualized Sharpe ratios

	Gain: half life years				
CRRA	Decreasing	1	5	10	25
$\gamma = 1$	0.375	0.380	0.375	0.375	0.375
$\gamma = 2$	0.374	0.426	0.390	0.373	0.364
$\gamma = 3$	0.349	0.440	0.393	0.373	0.360
$\gamma = 5$	0.333	0.432	0.366	0.347	0.335

Table 6: Certainty equivalent (max = 2)

	Gain: half life years				
CRRA	Decreasing	1	5	10	25
$\gamma = 1$	0.078	0.091 (0.860)	0.082 (0.664)	0.077 (0.428)	0.074 (0.252)
$\gamma = 2$	0.037	0.060 (0.986)	0.046 (0.847)	0.041 (0.693)	0.037 (0.498)
$\gamma = 3$	0.027	0.039 (0.952)	0.031 (0.793)	0.028 (0.631)	0.026 (0.442)
$\gamma = 5$	0.018	0.025 (0.940)	0.020 (0.790)	0.018 (0.598)	0.017 (0.413)

Table 7: Autoregressive comparison

	Model			
CRRRA	Decreasing gain	AR(1)	AR(3)	AR(3)+Leverage
$\gamma = 1$	0.060	0.059	0.059	0.060
$\gamma = 2$	0.042	0.043	0.046	0.046
$\gamma = 3$	0.027	0.033	0.035	0.036
$\gamma = 5$	0.018	0.022	0.023	0.026

Table 8: MSE forecasts $\times 10^4$

	Model		
Decreasing gain	AR(1)	AR(3)	AR(3)+Leverage
0.261	0.197	0.192	0.190
	Gain: half life years		
Decreasing gain	1	5	10
0.261	0.222	0.249	0.258

Figure 1: Returns and Volatility




Figure 2: Constant Gain Forecast and VIX

