Discussion Paper Series

Endogenous time-varying risk aversion and asset returns

By

Michele Berardi

Centre for Growth and Business Cycle Research, Economic Studies,
University of Manchester, Manchester, M13 9PL, UK

May 2012
Number 168

Endogenous time-varying risk aversion and asset returns*

Michele Berardi
The University of Manchester

May 18, 2012

Abstract

Stylized facts about statistical properties for short horizon returns in financial markets have been identified in the literature, but a common cause for their manifestation has yet to be found. We show that a simple asset pricing model with representative agent and rational expectations is able to generate time series of returns that replicate such stylized facts if the risk aversion coefficient is allowed to change endogenously over time in response to unexpected excess returns. The same model, under constant risk aversion, would instead generate returns that are essentially Gaussian. We conclude that an endogenous time-varying risk aversion represents a very parsimonious way to make the model match real data on key statistical properties, and therefore deserves careful consideration from economists and practitioners alike.

Key words: Risk aversion; returns; asset prices; financial markets.

JEL classification: D83, G01, G02, G12.

* Correpsonding author: Dr. Michele Berardi, Economics, School of Social Sciences, The University of Manchester; Tel.: +44 (0)1612754834; Email: michele.berardi@manchester.ac.uk.

Endogenous time-varying risk aversion and asset returns

1 Introduction

The statistical analysis of price variations in financial markets has attracted a lot of attention, both from practitioners and academic economists, in an attempt to find regularities that could help understand and possibly predict the evolution of prices on such markets. Such extensive analysis has led to the identification of a number of statistical properties for financial returns that seem to hold across markets and over time, and that can been summarized in the following set of empirical stylized facts (see, e.g., Cont 2001 and Tseng and Li 2011): i) the distribution of returns is not Gaussian but presents instead fat tails; ii) there is no serial correlation in returns; iii) but there is positive correlation in absolute returns, with slow decay; iv) returns show strong volatility clustering, with large fluctuations that tend to cluster together.

Though there is agreement among researchers on such empirical observations, it has not been possible so far to identify a common origin for them, nor it has been possible to replicate them with standard rational expectations models. We suggest in this work that such common origin could be identified in the time-varying nature of the risk aversion coefficient for investors, and show that an otherwise standard, rational expectations asset pricing model, once enhanced with such feature, can generate time series for returns that replicate very closely the main stylized facts identified in the empirical literature.

The backbone model that we use for our analysis is a simple, standard, present value asset pricing model with stochastic dividends, as presented for example by De Long et al (1990). It is well known that such model, when closed with rational expectations, does a poor job in matching the key stylized facts about financial returns mentioned above, as it implies that returns are normally distributed and independent over time.

A growing body of literature has recently used adaptive learning to try improve the empirical performance of asset pricing models. Examples include Branch and Evans (2010, 2011), Adam, Marcet and Nicolini (2008), Cárceles-Poveda and Giannitsarou (2008), Bullard and Duffy (2001), Brock and Hommes (1998) and Timmermann (1993, 1996).

We suggest instead here a different route and we will show that a simple behavioral modification of an otherwise standard asset pricing model can go a long way in matching statistical properties of historical data for financial returns. The risk aversion coefficient for investors is usually assumed, in the standard economics and finance literature, to be a primitive of the model, a feature that is hard wired into the brain of people when they are born and that does not change. We believe instead that there is scope for modelling the attitude of agents towards risk as a feature that depends on the environment where agents make their decisions, and that it evolves with it.

For example, narrative evidence suggests that many people, who had been very cautious up to that point with their investment decisions and mainly kept their savings into government securities or similar activities, during the stock market bubble of the late nineties and early 2000 abandoned their safe investments and moved their money into more risky assets. Observing high rates of returns on stock markets, those people became more willing to take on risky activities in an attempt to join in and share the high profits that were realized on financial markets at the time. In a sort of herd-like behavior induced by their decreased risk aversion, previously cautious investors entered into the stock market. When prices then started to fall and returns decreased, those same investors became afraid of losses, their risk aversion increased and they flew financial markets, herding into selling their assets and fuelling a sharper decrease in prices.

Alan Greenspan, on this regard, said in a speech at the Federal Reserve Bank of Kansas City (Greenspan, 2005; Italics added):

> "Thus, this vast increase in the market value of asset claims is in part the indirect result of *investors accepting lower compensation for risk*. [...] Any onset of *increased investor caution* elevates risk premiums and, as a consequence, lowers asset values and promotes the liquidation of the debt that supported higher asset prices."

Greenspan suggests in his speech that changes in market values depend, partly, on changes in risk premia required by agents, which in turn depend on the attitude of investors towards risk. In this paper we make formal this argument and show that adding this feature to an otherwise standard model changes completely the statistical properties of simulated asset returns, making them similar to those observed on real markets.

Time-varying risk aversion is not new in economics. In consumption-based asset pricing models, for example, Brandt and Wang (2003) propose a time-varying risk aversion coefficient that responds to news about consumption growth and inflation, while Li (2007) studies asset prices under the assumption of a countercyclical risk aversion. We propose instead a process for risk aversion that depends on unexpected excess returns: agents adapt their attitude towards risk on the basis of the unexpected excess gains that they observe from risky activities.

While in standard economics risk aversion is a feature that depends solely on the curvature of the utility function being maximized by agents, whose form is assumed constant over time, in behavioral economics Kahneman and Tversky (1979)'s prospect theory argues that expected utility maximization is a poor representation of how people make choices under risk, and suggests instead an alternative framework where people's attitude towards risk is situation dependent. We will continue to use here the expected utility maximization framework, but modify it to allow the curvature of the utility function, and therefore the risk attitude of agents, to evolve over time endogenously. An evolutionary justification for changes in attitudes towards risk is provided by Netzer (2009), who shows that from an evolutionary perspective the utility function of agents, and therefore their risk aversion, should depend on the probability distribution of alternatives about which agents need to make decisions. In our contest, such alternatives are represented by returns from risky vs. risk-free activities, and agents adapt their perceptions about the distribution of such alternatives using observations about unexpected excess returns on the stock market.

The plan of the paper is as follows: in Section 1.1 we discuss the literature related to our work, while in Section 1.2 we provide some evidence on the statistical properties of financial returns by looking at the S&P500 index as a representative case. Section 2 introduces the basic model, Section 2.1 discusses endogenous time-varying risk aversion and Section 2.2 derives the equilibrium solution for the model. Section 3 presents results from simulations of the model with constant and with endogenous time-varying risk aversion, and compares the resulting series for returns with those from real data. Section 4 concludes.

1.1 Related literature

We now discuss further in detail some of the literature cited in the previous section, and how it relates to our work. Branch and Evans (2010) show that real time learning dynamics, in an otherwise standard consumption based asset pricing model, calibrated to U.S. stock data, is capable of reproducing regime-switching returns and volatilities. Branch and Evans (2011) introduce learning about risk and returns in the De Long et al (1990) framework and show that escape dynamics emerge which look like stock market crashes, even though the escape route is not from a bubble high but from the equilibrium fundamental value. Hommes and Zhu (2011) use the concept of stochastic consistent expectations equilibrium to explain excess volatility in a standard present value asset pricing model with stochastic dividends similar to the one we consider here. Adam et al (2008) show how adaptive learning can generate excess volatility in a consumption based asset pricing model and present an estimated version of the model to US data which can replicate some asset price puzzles such as stock price volatility, the persistence of the price-dividend ratio and the predictability of long-horizon returns. All these works focus on the long-horizon properties of asset prices returns, while we will focus our attention on trying to explain and replicate statistical properties of returns in the short run.

In terms of time-varying risk aversion, Brandt and Wang (2003) present a model where the coefficient of risk aversion changes in response to news about consumption growth and inflation, and find empirical support for the hypothesis that aggregate risk aversion varies in response to news about inflation. Li (2007), instead, assumes a countercyclical risk aversion which drives a pro-cyclical risk premium in asset prices, but finds that such feature may not help explain important facts such as the predictability of long-horizon stock returns or the univariate mean-reversion of stock prices. Finally, Smith and Whitelaw (2009) find empirically evidence in support of the hypothesis that risk aversion moves countercyclically.

An important theoretical paper for our modelling choice of risk aversion is Netzer (2009), who proposes a model of the evolution and adaptation of hedonic utility and provides an evolutionary explanation for risk attitudes that adapt according to changes in the perceived distribution of possible alternatives. We will discuss this point at length in Section 2.1.1.

In terms of evidence about short-horizon returns on financial markets, stylized facts are presented in a number of papers in the literature, such as Cont (2001) and Tseng and Li (2011). Both works show that the same stylized facts discussed in Section 1.2 below hold for a large number of financial time series, including Standard & Poor's 500 Index, NASDAQ Composite Index and Hang Seng Index, series for individual stock prices such as IBM, Microsoft and BMW, and even series for exchange rates.

1.2 Empirical evidence on returns

As we mentioned before, the main stylized facts identified for short-horizon returns on financial markets are: i) the distribution of returns in not Gaussian, and presents instead fat tails; ii) there is no correlation in returns; iii) but there is positive correlation (with slow decay) in absolute returns; iv) returns show volatility clustering, i.e., large fluctuations tend to cluster together.

We will present here evidence in support of such statistical regularities from daily returns from the S&P500 index, for the period 02/01/1957 until 12/04/2012,[1] but we want to stress, as said before, that returns computed from many other asset price series present the same key features (see, e.g., Tseng and Li, 2011).

Returns are computed as

$$r_t = \frac{p_t - p_{t-1}}{p_{t-1}},$$

where p_t is the price of the asset or index at time t. It is also common practice to normalize returns as follows

$$nr_t = \frac{r_t - \mu_t}{\sigma_r},$$

where μ_t and σ_t are the mean and standard deviation of returns. For the S&P500 index, $\mu_t = 0.000294$ and $\sigma_r = 0.0100$ and returns and normalized returns are plotted in Figure 1. It is clearly evident the volatility clustering of returns, with large movements that tend to cluster together at particular times.

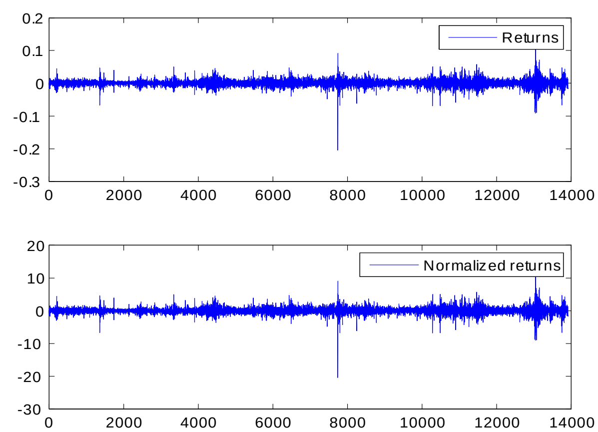

Figure 1: Returns and normalized returns for the S&P500 index.

We then bin normalized returns into 1000 equally spaced groups according to their value, and obtain an histogram (adjusted for the number of observations, in order to show relative frequencies) that approximates the empirical distribution of data (Figure 2). We then compare it with the histogram for a series of simulated Gaussian returns with same mean and variance (Figure 3): it is possible to see that the distribution of empirical returns presents fat tails compared to the Gaussian one, with significant probability mass on returns with

[1]Data are freely available at: http://research.stlouisfed.org/fred2/series/SP500?cid=32255

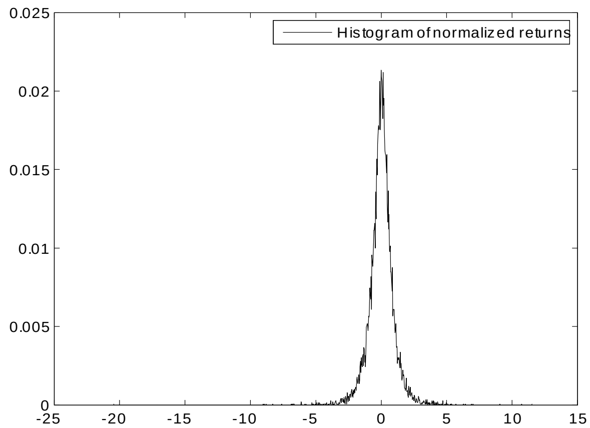

Figure 2: Histogram of S&P500 normalized returns.

high deviation from the mean. Non-normality is also confirmed by the 4th moment of the

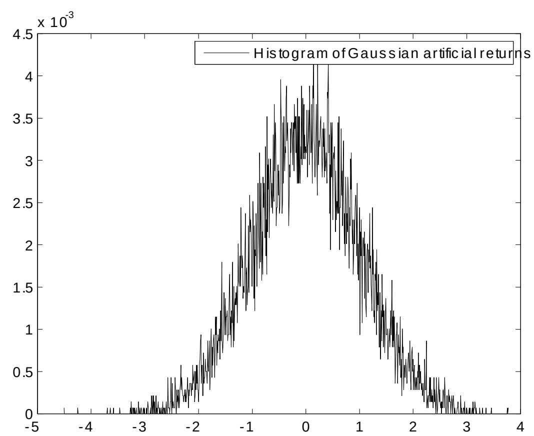

Figure 3: Histogram of artificial Gaussian returns.

distribution, kurtosis, which in the data is 24.378, while the value for a Gaussian distribution is 3, and by the Jarque-Bera test, which rejects at 5% significance level the null hypothesis that the sample comes from a normal distribution with unknown mean and variance.

Another key stylized fact concerns the correlation of returns and absolute returns over time. We therefore compute and plot the autocorrelation functions (ACFs) for both series in Figure 4. This clearly shows that returns are uncorrelated, but absolute returns show

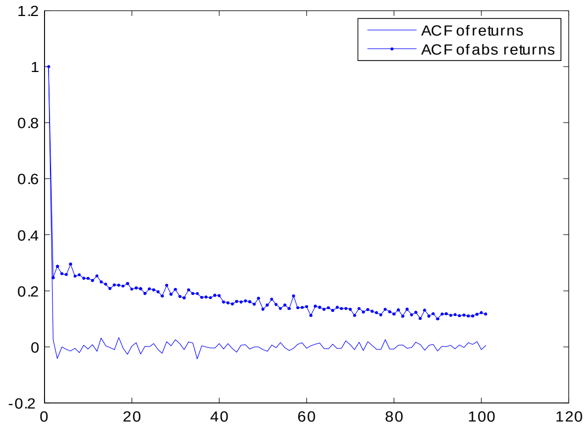

Figure 4: ACFs for returns and absolute returns for the S&P500 index.

positive correlation with slow decay over time.

We have presented in this section evidence about key stylized facts concerning statistical regularities in financial returns. We now develop a simple model that will be able to replicate closely all such facts.

2 The model

We start with a standard, present value asset pricing model with stochastic dividends. There are two types of assets: a risk free one, elastically supplied, with a gross rate of return $R = \beta^{-1}$, where β is the discount factor; and a risky asset, whose price is p_t and pays a stochastic dividend d_t. The supply of the risky asset is exogenous and stochastic.

The dynamic equation for wealth (W_t) is then represented by

$$W_{t+1} = RW_t + (p_{t+1} + d_{t+1} - Rp_t)z_t^d$$

where z_t^d is the demand for the risky asset.

Investors have a CARA utility function of the form

$$U(W_{t+1}) = -e^{-\alpha_t W_{t+1}}, \tag{1}$$

where α_t is the coefficient of absolute risk aversion, but agents are myopic mean variance maximizers, and therefore maximize[2]

$$E_t\left[\alpha_t E_t W_{t+1} - \alpha_t^2/2Var_t(W_{t+1})\right] \tag{2}$$

where E_t and Var_t are the conditional expectation and variance of wealth based on the subjective probability distribution of agents.

The coefficient of risk aversion in (1) and (2) has a time t subscript, to make it explicit that we will allow such parameter to evolve over time: the specification of its endogenous dynamics will be given in the next section. Though risk aversion evolves over time, we are going to assume that at each time t agents take such coefficient as given in their maximization problem, only to revise it in the following period on the basis of new evidence. In this respect, agents in our setting implement an anticipated utility model in the spirit of Kreps (1998). Such modeling strategy has been largely adopted in the macroeconomics literature on bounded rationality and learning (e.g., Sargent 1993, 1999 and Evans and Honkopohja 2001) and we believe it can represent a good approach also in the present contest where behavioral parameters rather than beliefs evolve over time.

From the above setting, it follows that the optimal demand for the risky asset is given by

$$z_t^d = \frac{E_t\left(p_{t+1}+d_{t+1}\right) - \beta^{-1}p_t}{\alpha_t \sigma_t^2}, \tag{3}$$

where σ_t^2 is agents' conditional variance of excess returns $p_{t+1} + d_{t+1} - Rp_t$ and is given by

$$\sigma_t^2 = E_t\left[\left(p_{t+1}+d_{t+1}\right) - E_t\left(p_{t+1}+d_{t+1}\right)\right]^2. \tag{4}$$

Equating demand and supply, denoted by z_t^s, we obtain the pricing equation

$$p_t = \beta E_t\left(p_{t+1}+d_{t+1}\right) - \beta\alpha_t\sigma_t^2 z_t^s. \tag{5}$$

The exogenous process for dividends is assumed to be given by

$$d_t = d_0 + u_t, \tag{6}$$

where d_0 is a constant and u_t is an i.i.d., zero mean, normally distributed disturbance.

[2] Equivalently, we can assume that agents believe wealth to be normally distributed, in which case maximizing the expected value of (1) is equivalent to maximizing the expected value of $-\exp\{-\alpha_t E_t W_{t+1} + \alpha_t^2/2Var_t\left(W_{t+1}\right)\}$, and we obtain the same pricing equation.

We also assume that supply follows the exogenous random process

$$z_t^s = s_0 + v_t, \tag{7}$$

where s_0 is a constant and v_t is an i.i.d., zero mean, normally distributed disturbance uncorrelated with u_t.

2.1 Endogenous time-varying risk aversion (ETVRA)

The coefficient of absolute risk aversion, α_t, is modelled as time-varying and endogenous, depending on unexpected (excess) returns in the stock market: higher returns than expected make agents more willing to take on the risk involved in investing in the stock market. Specifically, we postulate a process of the form

$$\alpha_t = \alpha_{t-1} + \gamma_\alpha \pi_t, \tag{8}$$

where $\gamma_\alpha \leq 0$ and

$$\pi_t = p_t + d_t - E_{t-1}(p_t + d_t), \tag{9}$$

i.e., risk aversion decreases when excess profits in the financial market are higher than expected.[3] Parameter γ_α represents the sensitivity of risk aversion to unexpected returns: with $\gamma_\alpha = 0$ we have the standard case of constant risk aversion over time, while negative values of γ_α mean that agents are willing to take on more risk when they see excess returns in the stock market compared to what they expected (and vice-versa). Such endogenous dynamics for the coefficient of risk aversion will make the reduced form parameters in the solution for prices be time-varying.

2.1.1 An evolutionary justification to ETVRA

Netzer (2009) considers a model of the evolution and adaptation of hedonic utility, and shows that the utility function that we can expect to emerge under evolutionary considerations is strictly related to the cumulative density function (CDF) of the choice variable for agents. In particular, he shows that the utility function that minimizes the probability of choosing the wrong alternative when individuals have only a finite ability to distinguish between choices is exactly equal to the CDF of such choices, a result also found by Robson (2001).[4]

[3] Note that $E_{t-1}\beta p_{t-1} = \beta p_{t-1}$, so these terms drop out of equation (9).

[4] Netzer (2009) also shows that the utility function that maximizes the expected fitness can be obtained by integrating a concave transformation of the distribution function of the alternatives, so that the strong

It follows that the Arrow-Pratt coefficient of risk aversion equals the (negative of) the ratio between the derivative of the probability density function (PDF) and the PDF itself. Assuming that alternatives arrive according to a truncated exponential distribution with rate $\lambda > 0$, the coefficient of absolute risk aversion turns out to be equal to λ. If agents don't know the exact distribution of such alternatives and have to estimate it, a shift in their estimates results in a shift in the coefficient of risk aversion.

Formally, given a probability density function $f(x)$ for the alternatives x, drawn from a set $X = [a, b]$, Netzer shows that in the contest of hedonic utility under uncertainty, the utility function $U(x)$ that minimizes the probability of choosing the wrong alternative for an agent is equal in the limit to $F(x)$, the CDF for $f(x)$, and therefore the Arrow-Pratt coefficient of absolute risk aversion (RA) equals

$$RA(x) = -\frac{f'(x)}{f(x)}.$$

It follows that if $f(x) = \frac{\lambda e^{-\lambda x}}{1-e^{-\lambda b}}$, a truncated exponential distribution with rate parameter λ, then

$$RA(x) = \lambda.$$

$U(x)$ exhibits constant absolute risk aversion, but risk aversion can still vary over time as agents update their beliefs about the distribution f of alternatives, which results in a change in the estimate for λ.

In our setting, f would be the perceived distribution of excess returns from the risky activity, and agents adapt their beliefs about λ over time by moving in the direction indicated by unexpected returns. Equation (8) can therefore be interpreted as a simple adaptive learning rule with constant gain, through which agents keep track of the evolution of returns on the stock market by adjusting their beliefs in the direction suggested by forecast errors. A shift in probability is signalled by π_t, a realization of unexpected excess returns on the stock market. Parameter γ_α represents therefore a behavioral feature that captures the degree of adjustment of agents' perceived distribution of excess returns to evidence about unexpected realizations. As beliefs about the distribution of returns adapt, so does risk aversion. In particular, a decrease in the estimate of λ, which corresponds to a shift in probability mass to alternatives with larger payoffs, reduces risk aversion.

connection between utility function and distribution of choices still holds under this alternative criterion. In particular, both utility functions imply the same relationship (up to a multiplicative transformation) between the coefficient of absolute risk aversion and the probability density function of alternatives.

2.2 Equilibria

In this Section we present first the rational expectations (RE) equilibria under constant risk aversion, and then introduce time-varying risk aversion dynamics.

2.2.1 Constant risk aversion

Under RE, with constant risk aversion ($\alpha_t \equiv \alpha$), it is well known that equations (5)-(7) admit two possible equilibria, the fundamental solution and the bubble one.

The price equation in the fundamental equilibrium is represented by

$$p_t = \beta (1-\beta)^{-1} \left(d_0 - \alpha s_0 \sigma^2\right) - \beta\alpha\sigma^2 v_t, \tag{10}$$

while in the "bubble" equilibrium it is given by

$$p_t = \alpha\sigma^2 s_0 - d_0 + \beta^{-1} p_{t-1} - \alpha\sigma^2 v_{t-1} + \xi_t, \tag{11}$$

where ξ_t is a martingale difference sequence that opens the door to sunspot variables affecting prices. It is also known (see, e.g., Branch and Evans, 2011) that the fundamental equilibrium is adaptively learnable by agents, while the "bubble" one is not. Moreover, since $\beta < 1$, the process for p_t in the bubble equilibrium is not stable.

2.2.2 ETVRA

With a time-varying risk aversion, the solution to the model must be enriched to include an equation for the evolution of risk aversion and for the variance of returns. We stress again here that agents in our setting, in the spirit of the anticipated utility model of Kreps (1998), solve at each point in time their maximization problem (and we, as modeler, find the equilibrium) as if risk aversion was constant, but then modify their attitude towards risk as new evidence on unexpected excess returns comes available next period.

For the fundamental equilibrium, we therefore have that the evolution of the system is represented by

$$\begin{aligned} p_t &= \beta (1-\beta)^{-1} \left(d_0 - \alpha_t s_0 \sigma_t^2\right) - \beta\alpha_t\sigma_t^2 v_t && (12) \\ \alpha_{t+1} &= \alpha_t + \gamma_\alpha \pi_t, && (13) \\ \pi_t &= -\beta\alpha_t\sigma_t^2 v_t + u_t && (14) \\ \sigma_t^2 &= \frac{1 \pm \sqrt{1 - 4\alpha_t^2\beta^2\sigma_v^2\sigma_u^2}}{2\alpha_t^2\beta^2\sigma_v^2} && (15) \end{aligned}$$

where equation (14) for unexpected excess returns comes from computing (9) in the fundamental equilibrium. Expression (15) for the conditional variance of excess returns is obtained from (4) by substituting in the equations for prices and dividends.

By substituting into expression (13) for α_{t+1} those for π_t and σ_t^2, it is possible to see that the dynamics for the risk aversion coefficient are represented by a non linear difference equation that depends on parameters β and γ_α and on the variances and realizations of the two processes u_t and v_t:

$$\alpha_{t+1} = \alpha_t + \gamma_\alpha \left[u_t - \beta \alpha_t \frac{1 \pm \sqrt{1 - 4\alpha_t^2 \beta^2 \sigma_v^2 \sigma_u^2}}{2\alpha_t^2 \beta^2 \sigma_v^2} v_t \right].$$

Simulations will be done with the (-) root, as this is the one that Branch and Evans (2011) show to be stable under learning. It is also the one that ensures that simulated prices follow a smooth process without unrealistic and sudden large jumps.

In the bubble equilibrium, instead, the evolution of the system under ETVRA is described by

$$\begin{aligned}
p_t &= \left(\alpha_{t-1}\sigma_{t-1}^2 s_0 - d_0\right) + \beta^{-1} p_{t-1} - \alpha_{t-1}\sigma_{t-1}^2 v_{t-1} + \xi_t, \\
\alpha_{t+1} &= \alpha_t + \gamma_\alpha \pi_t, \\
\pi_t &= u_t \\
\sigma_t^2 &= \sigma_u^2 + \sigma_\xi^2.
\end{aligned}$$

It is easy to see that in this case α_t simply follows a random walk process

$$\alpha_{t+1} = \alpha_t + \gamma_\alpha u_t.$$

In a bubble equilibrium, therefore, asset prices diverge, unexpected excess returns from the risky activity are white noise and the ETVRA follows a random walk: for all these reasons we do not focus on such equilibrium in this work. In a fundamental equilibrium, instead, asset prices follow a stationary process, but unexpected excess returns depend on risk aversion of agents and the dynamics for the ETVRA coefficient are highly nonlinear. These features, we will show, impact significantly on the implied nature of returns from the risky asset, and are able to generate dynamics that in many dimensions closely resemble those observed in real financial markets.

3 Simulations and statistical analysis

In order to compare the statistical properties of observed returns on financial markets with those that come from our model, we simulate price dynamics in the fundamental RE equilibrium, first with constant risk aversion and then with ETVRA. From such series we compute returns and analyze their properties in comparison with the stylized facts reported in Section 1.2.

We calibrate the model as follows: $\beta = .95$, $s_0 = 1$, $d_0 = 5$, $\sigma_v^2 = .5$, $\sigma_u^2 = .9$. These parameter values are taken from Branch and Evans (2011), apart from d_0 which we have set to a higher value in order to avoid prices hitting the zero lower bound.[5] We also set the initial value for risk aversion, α_1, equal to 0.75, which is the value used by Branch and Evans (2011) for their constant risk aversion. Sensitivity analysis regarding these parameters has shown that our results are robust to different (but realistic) parameterizations. The key new parameter γ_α is instead calibrated to 0.025: extensive investigations have shown that the higher is γ_α, the further away from normality is the distribution of returns. We will discuss further this point later on when reporting results from simulations.

3.1 Returns dynamics with constant risk aversion

We start by simulating our model with constant risk aversion and compute returns from the resulting time series for asset prices. Mean and standard deviation of simulated returns, $\tilde{\mu}_t$ and $\tilde{\sigma}_t$, are respectively 0.000119 and 0.0156, in line with those found in real data.

Returns from a representative run of simulated data are plotted in Figure 5, together with their normalized counterpart. It is already clear to the eye that the statistical properties deeply differ from those of real returns, plotted in Figure 1, as there seems to be no volatility clustering in simulated returns under constant risk aversion.

We then plot the histogram for normalized returns in Figure 6, again with data binned in 1000 intervals and highs adjusted for the number of observations. Comparing it with Figures 2 and 3, it is immediate to see the resemblance with the second one, which suggests that simulated returns under constant risk aversion are normally distributed. Such evidence is supported by the Jarque-Bera test, which can not reject at 5% significance level the null hypothesis that the sample comes from a normal distribution with unknown mean and

[5] Branch and Evans (2011) avoid such problem by modelling supply slightly differently from us and allowing it to become endogenous when prices sharply decline: this is meant to capture the drying up of asset float in financial markets that perform poorly. We chose not to use such mechanism here, but simulations showed that it would not affect our results.

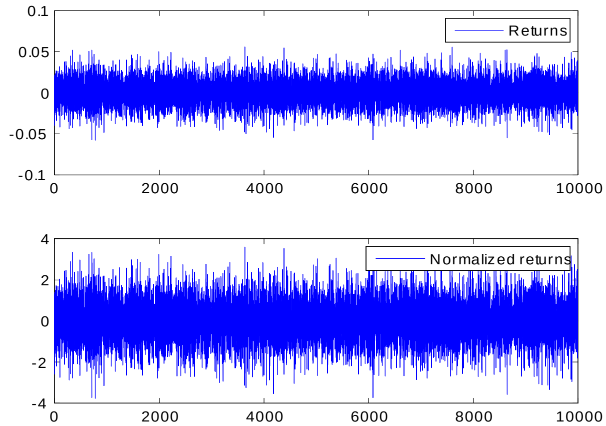

Figure 5: Returns and normalized returns for the simulated series with constant risk aversion.

variance. Normality of returns in this case is also confirmed by the estimated kurtosis, which is 3.0049. Figure 7 presents the ACFs for simulated returns and their absolute value: it is

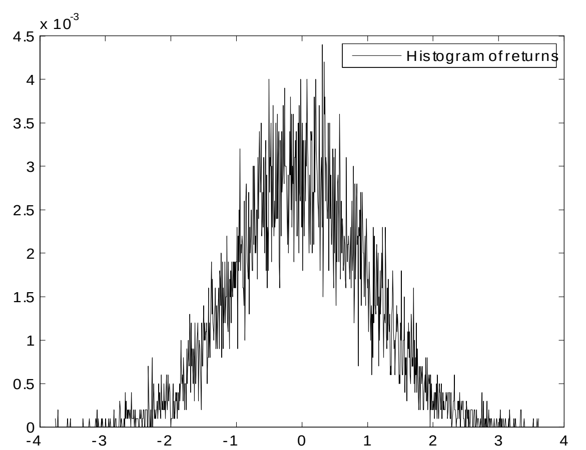

Figure 6: Hisogram of normalized returns for the simulated series with constant risk aversion.

immediate to see that there is no serial correlation in either of the two series, contrary to what real returns show.

All the above evidence on simulated returns under constant risk aversion clearly indicates that the model with constant risk aversion is not able to replicate the key stylized facts that have been consistently observed in series for returns on real financial markets.

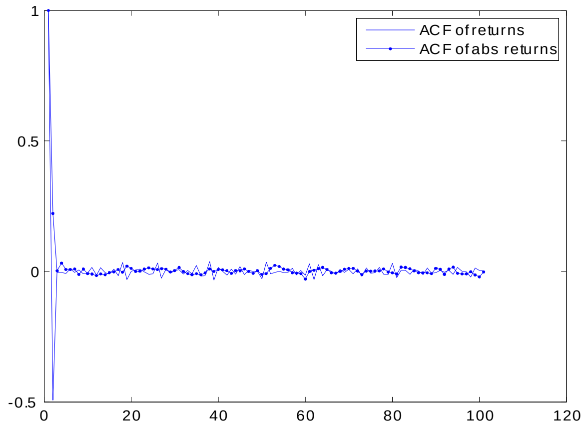

Figure 7: ACFs for returns and absolute returns for the simulated series with constant risk aversion.

3.2 Returns dynamics with ETVRA

We now simulate the model with endogenous time-varying risk aversion and compute returns from the resulting time series for asset prices. Mean and standard deviation of simulated returns, $\tilde{\mu}_t$ and $\tilde{\sigma}_t$, are respectively 0.000090 and 0.0110, again in line with those found in real data.

Returns from a representative run of simulated data are plotted in Figure 8, together with their normalized counterpart. It is evident now that ETVRA generates volatility clustering in returns.

In Figure 9 we then present the histogram for simulated normalized returns, again binned in 1000 intervals and with highs adjusted for the number of observations. Comparing this picture with Figures 2 and 3, it seems evident now that the distribution of returns under ETVRA differs significantly from a Gaussian and is instead similar to the one observed for returns on real financial markets. This impression is confirmed again by the Jarque-Bera test, which this time rejects at 5% significance level the null hypothesis that the sample comes from a normal distribution with unknown mean and variance. Also in terms of kurtosis, we find a value of 21.48, in line with the one estimated for S&P500 returns and much higher than the one, close to 3, obtained under constant risk aversion and consistent with a Gaussian distribution.

We note here that kurtosis for the distribution of simulated returns is sensitive to the

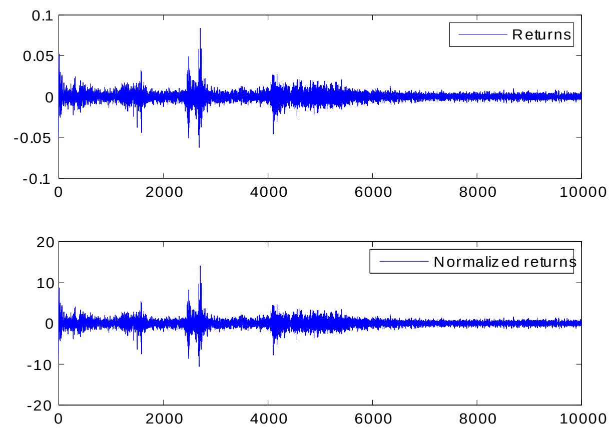

Figure 8: Returns and normalised returns for the simulated series with ETVRA.

value of γ_α: higher values of γ_α imply higher kurtosis, as the distribution of returns drifts further away from normality.[6] As risk aversion becomes more sensitive to unexpected excess returns from the risky asset, the feedback effect from returns to asset demand and therefore prices is strengthened, making the distribution of returns more heavy tailed. In Figure 10

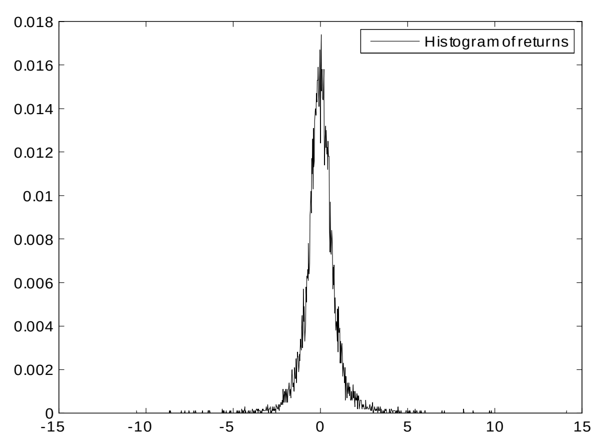

Figure 9: Histogram of normalized returns for the simulated series with ETVRA.

we then present the ACFs for simulated returns under ETVRA and for their absolute value:

[6] The case with constant risk aversion corresponds to $\gamma_\alpha = 0$, and we have seen in Section 3.1 that it implies normality of returns and a kurtosis close to 3.

we can see that while there is no serial correlation in returns, absolute returns show positive correlation with slow decay, very similar to the one observed in S&P500 returns. All the above

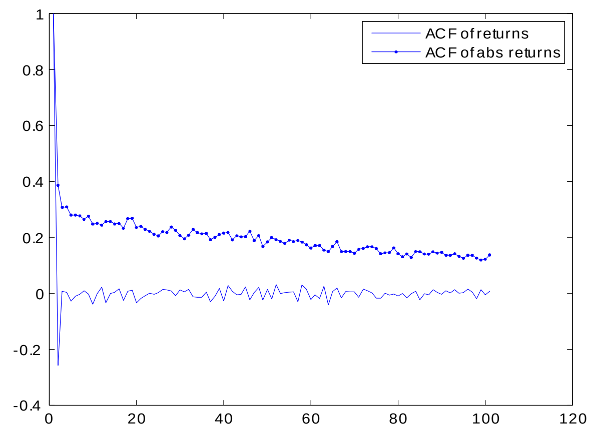

Figure 10: ACFs for returns and absolute returns for the simulated series with ETVRA.

evidence on simulated returns with ETVRA shows that simply introducing an endogenous risk aversion coefficient that responds to unexpected excess returns from the risky activity in an otherwise standard asset pricing model fundamentally changes the statistical properties of simulated returns, and makes them strikingly resemble those observed on real financial markets.

4 Conclusions

We have proposed in this paper a possible unifying explanation for many stylized facts about returns on financial markets that had so far eluded a common understanding. Our suggestion is that features like volatility clustering, non-normality of distribution and heavy tails, absence of correlation in returns but positive correlation (with slow decay) in their absolute value, can all emerge in an otherwise standard, rational expectations model of asset pricing if an endogenous time-varying risk aversion is introduced. Simulations of the model with a coefficient of risk aversion that is allowed to change endogenously in response to unexpected excess returns in the risky activity produced returns that are surprisingly close, in all these dimensions, to real data, while simulated returns from the same model under constant risk aversion display normality of distribution, absence of volatility clustering and lack of correlation both in returns and in their absolute value. We conclude that endogenous time-varying risk aversion represents a very parsimonious way to explain key stylized facts in financial markets and deserves careful consideration from economists and practitioners alike.

References

[1] Adam, K., Marcet, A., Nicolini, J.P., 2008. Stock market volatility and learning. European Central Bank Working Paper Series n. 862.

[2] Branch, W.A., Evans, G.W., 2010. Asset return dynamics and learning. Review of Financial Studies 23, 1651-1680.

[3] 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.

[4] Brandts, M.W., Wang, K.Q., 2003. Time-varying risk aversion and unexpected inflation. Journal of Monetary Economics 50, 1457-1498.

[5] Brock, W., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. Journal of Economic Dynamics & Control 22, 1235-1274.

[6] Bullard, J., Duffy, J., 2001. Learning and excess volatility. Macroeconomic Dynamics 5, 272–302.

[7] Carceles-Poveda, E., Giannitsarou, C., 2008. Asset pricing with adaptive learning. Review of Economic Dynamics 11, 629–651.

[8] Cont, R., 2001. Empirical properties of asset returns: stylized facts and statistical issues. Quantitative Finance 1, 223-236.

[9] De Long, J.B., Shleifer, A., Summers, L.H., Waldmann, J., 1990. Noise trader risk in financial markets. Journal of Political Economy 98, 703-38.

[10] Evans, G.W., Honkapohja, S., 2001. Learning and Expectations in Macroeconomics. Princeton University Press: Princeton.

[11] Greenspan, A., 2005. Reflections on Central Banking. Speech, Federal Reserve Bank of Kansas City Jackson Hole Symposium, Jackson Hole, WY, August 26, 2005. http://www.federalreserve.gov/boarddocs/speeches/2005/20050826/default.htm

[12] Hommes, C., Zhu, M., 2011. Learning under misspecification: a behavioral explanation of excess volatility in stock prices and persistence in inflation. CeNDEF Working Papers 11-04.

[13] Kahneman, D., Tversky, A., 1979. Prospect theory: an analysis of decision under risk. Econometrica 47, 263-291.

[14] Kreps, D., 1998. Anticipated Utility and Dynamic Choice. In D.P. Jacobs, E. Kalai, and M. Kamien, eds., Frontiers of Research in Economic Theory, Cambridge: Cambridge University Press, 242–74.

[15] Li, G., 2007. Time-varying risk aversion and asset prices. Journal of Banking and Finance 31, 243-257.

[16] Netzer, N., 2009. Evolution of time preferences and attitudes towards risk. American Economic Review 2009, 937–955.

[17] Robson, A.J., 2001. The biological basis of economic behavior. Journal of Economic Literature 39, 11-33.

[18] Sargent, T.J., 1993. Bounded Rationality in Macroeconomics. Clarendon Press: Oxford.

[19] Sargent, T.J., 1999. The Conquest of American Inflation. Princeton University Press: Princeton.

[20] Smith, D.R., Whitelaw, R.F., 2009. Does risk aversion change over time? Mimeo.

[21] Timmermann, A., 1993. How learning in financial markets generates excess volatility and predictability in stock prices. Quarterly Journal of Economics 108, 1135–1145.

[22] Timmermann, A., 1996. Excess volatility and predictability of stock prices in autoregressive dividend models with learning. Review of Economic Studies 63, 523–557.

[23] Tseng, Jie-Jun, Li, Sai-Ping, 2011. Asset returns and volatility clustering in financial time series. Physica A 390, 1300–1314.