Journal of Economic Dynamics & Control 33 (2009) 1106–1122

Contents lists available at ScienceDirect

Journal of Economic Dynamics & Control

journal homepage: www.elsevier.com/locate/jedc

Uncertainty aversion in a heterogeneous agent model of foreign exchange rate formation

Roman Kozhan *, Mark Salmon

FERC, Warwick Business School, The University of Warwick, Scarman Road, Coventry CV4 7AL, UK

A R T I C L E I N F O

Article history:
Received 30 January 2008
Accepted 17 November 2008
Available online 25 February 2009

JEL classification:
C12
C15
C63
D81
F31

Keywords:
Uncertainty aversion
Exchange rate formation
Heterogeneous agents

A B S T R A C T

This paper provides what we believe to be the first empirical test of whether investors in the foreign exchange market are uncertainty averse. We do this using a heterogeneous agents model in which fundamentalist and chartist beliefs of the exchange rate co-exist and are allowed to be either uncertainty neutral or uncertainty averse. Uncertainty aversion is modelled using the maxmin expected utility approach. We find significant evidence of uncertainty aversion in the FX market where in particular fundamentalists are found to be largely uncertainty neutral while chartists are mainly uncertainty averse. Inclusion of uncertainty averse agents significantly improves the empirical performance of the model.

0. Introduction

Foreign exchange theory and modelling has been in a state of turmoil since Meese and Rogoff (1983) demonstrated that the standard theoretical models, of that time, could not outperform a random walk in out-of-sample prediction. The problem is that a random walk can only represent an efficient price if the expected equilibrium exchange rate were to be constant and for many reasons, in a risk averse world, this is an unacceptable position to take. Theory implies that the expected equilibrium rate varies over time to reflect, amongst other factors, time varying risk premia which should imply a degree of predictability. Standard foreign exchange models developed within the rational representative agent framework appear therefore to be dominated by a simple random walk model that is itself seen to be invalid.

Despite considerable research in the intervening 25 years no consensus theoretical paradigm has been developed to resolve the paradox raised by Meese and Rogoff. A recent detailed analysis by Della Corte et al. (2009) concludes, following a Bayesian model averaging exercise, that fundamentals are important but a variety of different models may be relevant to explain behaviour in foreign exchange markets. We believe a new theoretical approach may be needed to explain these results—one which removes the representative agent paradigm and addresses uncertainty as opposed to risk in financial markets. This paper represents a first attempt to bring these two ideas together in order to explain observed movements in foreign exchange rates.

* Corresponding author. Tel.: +44 2476524118.
E-mail addresses: Roman.Kozhan@wbs.ac.uk (R. Kozhan), Mark.Salmon@wbs.ac.uk (M. Salmon).

0165-1889/$ - see front matter
doi:10.1016/j.jedc.2008.11.008

Finance theory has traditionally been based on an analysis of risk, measured using some aspect of a distribution which is itself assumed to be unique and characterises returns. However, a growing body of research has argued that one reason why the standard model, based on expected utility theory, fails to explain many "anomalies" and stylised facts observed across financial markets may be because agents in these markets face uncertainty as opposed to risk, as originally suggested by Knight (1921) and Keynes (1921). Bewley (2002) makes the distinction between risk and uncertainty in the following way: "a random variable is risky if its probability distribution is known, uncertain if its distribution is unknown". Market traders may simply not know which of several potential distributions to apply to evaluate an uncertain prospect or alternatively be faced with a situation in which they have no prior experience and hence no distribution at all to call on to carry out their risk calculus. The required probabilities may therefore either be unknown or unmeasurable. As argued by Keynes and Knight the uncertainty framework may in fact better represent decision making in financial markets than the risk paradigm where agents are assumed not to doubt their models nor, in particular, the associated probability distributions.

A number of authors have developed theories of decision-making under uncertainty including Schmeidler (1989), Gilboa and Schmeidler (1989) and Quiggin (1982). Epstein and Wang (1994), Uppal and Wang (2003), Maenhout (2004) and Hansen and Sargent (2007) among others have developed and applied related ideas of robust decision making to a variety of issues in Economics and Finance. Maccheroni et al. (2006) (see also Maccheroni et al., 2008) have recently proposed variational preferences that unify a preference for robustness and the multiple prior preferences of Gilboa and Schmeidler into a single framework and also developed an uncertainty based CAPM. The approach taken in these papers has been shown, often through simulation, to have the potential to successfully explain a number of anomalies seen in decision-making, such as the Ellsberg and Allais paradoxes as well as to provide explanations for the equity premium and home bias puzzles, for example. While the theoretical basis for decision-making under uncertainty has become relatively well developed, there is a clear lack of empirical evidence supporting the use of this alternative framework as a basis for explaining behaviour in financial markets. In other words, the question of whether people are uncertainty averse (have nonadditive preferences) in reality has not yet been resolved empirically.[1] In this paper we provide what we believe to be the first formal empirical test of uncertainty aversion within the FX market.

One obstacle is that virtually all models of decision-making under uncertainty in finance and foreign exchange rate theory are representative agent models and this is apparently in direct contrast with simple observation that all financial markets are in fact populated by heterogeneous agents. Aggregation arguments under rational expectations are insufficient to reduce markets to a single representative agent (see Kirman, 1992). So first, we develop a model of the foreign exchange market in which we allow fundamentalist and chartist beliefs of the exchange rate to co-exist and where traders are allowed to be either uncertainty neutral or uncertainty averse. As a basis for our model we draw on and extend existing models of the FX market developed in Kirman (1995), Brock and Hommes (1998), De Grauwe and Grimaldi (2006), Chiarella et al. (2006), Kirman et al. (2007) and Boswijk et al. (2007). This approach allows us to form endogenous demand and supply through the interaction of the different types of agent in the market. The realised exchange rate is then determined from the market clearing condition. Heterogeneity within agents' beliefs, whether fundamentalist or chartist, is captured by allowing for different ways the expectations of future prices are formed.

As mentioned above, there are several ways to model decision making under uncertainty and in this paper we use probably the most simple-maxmin expected utility (also known as the multiple priors or worst-case scenario model) of Gilboa and Schmeidler (1989). This approach has been used in several different asset pricing models, see Epstein and Wang (1994), Chen and Epstein (2002), Uppal and Wang (2003), Zhao et al. (2003) and Garlappi et al. (2007). The approach allows for different methods to construct the multiple prior set which reflects the degree of uncertainty or the range of potential models that might explain behaviour in the market. The "extreme" event in this set may in fact be close to some nominal model depending on how the prior set is drawn up. So we can consider both mild uncertainty or considerable ignorance as to the true model in the same framework. We assume that the investor is then faced with forming expectations regarding future exchange rates and considers the worst outcome within this set of models or effectively some interval, where the width of the set or interval is a subjective choice of the investor and hence we are able to capture different degrees of uncertainty aversion. Using the unscented Kalman filter (UKF) and nonlinear least squares (NLS) methods we then estimate our model with both uncertainty averse and uncertainty neutral agents and test whether this critical parameter (the width of the set) is significantly different from zero and hence whether uncertainty aversion exists in the foreign exchange market.

The paper is organised as follows. The next section provides a description of the model. In Section 2 we outline the UKF, used here as a method of state estimation and then parameter estimation by NLS. Section 3 contains the estimation results and model evaluation including specification tests relating to the model. The discussion and interpretation of the results is given in Section 4. Finally we provide some concluding remarks in Section 5.

[1] There have been a number of studies (see for instance Ellsberg, 1961; Mangelsdorf and Weber, 1994; Wakker, 2001) which provide empirical evidence that people are more uncertainty averse than uncertainty loving. However, the data used in these studies were collected either from questionnaires or laboratory experiments. Answering a questionnaire or acting in an artificial experiment may be quite different from taking real decisions in a market which may influence an individual's future well being. As far as we know, the issue of whether the realised prices are in fact influenced by uncertainty aversion has not been addressed in the existing literature.

1. Heterogeneous agent model

We assume that there are two currencies—domestic and foreign which are traded on the foreign exchange market. Denote by s_t the foreign exchange rate at time t—the price of one unit of foreign currency in units of domestic currency. There are N investors competing by trading in the market. Let ρ_t be the interest rate relevant to the foreign currency and r_t be the interest rate for the domestic country over the period t.

An agent's wealth at time t is determined by his trading activity and is equal to

$$W_t = (1 + r_{t-1})d_{t-1} + s_t(1 + \rho_{t-1})f_{t-1},$$

where d_t and f_t denote the trader's demands on domestic and foreign currency, respectively, held at time t. The individual's demands must satisfy the budget constraint $W_t = d_t + s_t f_t$ at each point of time.

There are two types of investor: fundamentalist and what we will call chartist but really these are defined by the tools they use to form expectations. The former believes that there exists an equilibrium price (fundamental value) $\bar{s}_t$ towards which the exchange rate will always move. More precisely, their expectation of the change in the exchange rate is proportional to the observed difference between the fundamental value and the previous level of the exchange rate and is expressed by the formula

$$E_t(s_{t+1}|F) = s_{t-1} + \nu(\bar{s}_t - s_{t-1}) \quad \text{with } 0 \leqslant \nu \leqslant 1, \tag{1}$$

where expectations are calculated conditional on the information available at time t. We discuss below in Section 1.4 how $\bar{s}$ is determined.

We assume chartists use a simple long–short moving average rule in order to predict a future deviation from its past level. Their exchange rate forecast is then given by

$$E_t(s_{t+1}|C) = s_{t-1} + h\left(\frac{1}{M_s}\sum_{i=1}^{M_s} s_{t-i} - \frac{1}{M_l}\sum_{i=1}^{M_l} s_{t-i}\right) \quad \text{with } h > 0, \tag{2}$$

where M_s and M_l are the lengths of the short and long moving average windows, respectively.

There is a substantial evidence that foreign exchange markets are populated by both these beliefs in one form or another. The surveys by Taylor and Allen (1992) and Cheung et al. (2004) show that chartist methods are widely used in financial institutions[2] along with fundamentals models, invariably based on interest rate parity conditions. Our particular choice of models for forming expectations just represent simple examples and more complex and potentially realistic models could be considered, in particular to extend the simple long–short moving average rule we have employed to represent chartist beliefs.

The information available to both types of trader at time t includes past levels of the exchange rate and past and present values of the fundamental variables, interest rates in our case. We do not allow agents to observe the contemporaneous equilibrium exchange rate since this is naturally not yet known in the market. A similar assumption is also used by Hellwig (1982), Blume et al. (1994), Boswijk et al. (2007), etc.

For analytic tractability, following Brock and Hommes (1998) and Boswijk et al. (2007) we assume that all investors have homogeneous expectations about the conditional second moment of the exchange rate $E_t(s_t^2|I) = E_t(s_t^2)$, $I = F, C$. The variable I indicates the individual's beliefs as being chartist or fundamentalist at time t based on past information. (Hereafter we omit the indicator I for notational convenience in statements which are true for both types of agent and where it does not cause a misunderstanding).

1.1. Demand functions

We identify four different individual demand functions which determine the market clearing exchange rate. In particular, $f_t^n(I)$ and $f_t^u(I)$, $I = F, C$ denote individual demands for the foreign currency by investors (uncertainty neutral and uncertainty averse, denoted by superscripts n and u, respectively).

1.1.1. Uncertainty neutral agents

Agents' risk preferences are characterised by a quadratic utility function $U(x) = x - \gamma x^2$, where we assume that the risk-aversion coefficient γ is the same for all traders. We make this assumption simply to focus the impact of uncertainty but clearly a more complex model with differing degrees of risk aversion could be developed.[3] Let an uncertainty neutral agent decide to hold f_t^n foreign currency units at time t. Then, any uncertainty neutral investor maximises the quadratic expected utility function of the next-period's wealth

$$E_t(U(W_{t+1}^n)|I) = E(W_{t+1}^n|I) - \gamma E((W_{t+1}^n)^2|I) \underset{f_t^n}{\to} \max, \tag{3}$$

[2] Including Central Banks as personal conversation within the Bank of England has shown us while preparing this paper.

[3] A referee has suggested our results would be explained by introducing heterogeneity in risk aversion among traders. We do not explore this question in order to keep analysis simple, but we note that there will be differential effects of uncertainty aversion and risk aversion on the model. Risk aversion affects demand in a linear manner while uncertainty aversion introduces a nonlinear transformation into the demand function. Introducing both heterogeneity in both uncertainty and risk aversion would call for eight types of agent making the analysis of uncertainty aversion much more difficult.

where the agent's wealth at $t+1$ is given by

$$W^n_{t+1} = (1+r_t)(W^n_t - s_t f^n_t) + s_{t+1}(1+\rho_t) f^n_t. \tag{4}$$

Maximising the expected utility with respect to f^n_t, domestic agents are able to determine their optimal trade, which is given in the following theorem.

Lemma 1. *Given the exchange rate level s_t the optimal trade for an uncertainty neutral agent is to hold f^n_t units of foreign currency, where*

$$f^n_t = \frac{E_t(s_{t+1}|I)(1+\rho_t) - s_t(1+r_t)}{2\gamma E_t(s^2_{t+1}|I)(1+\rho_t)^2}.$$

See Appendix for the proof.

1.1.2. Uncertainty averse agents

Uncertainty averse agents maximise their maxmin quadratic expected utility function of future wealth (see Gilboa and Schmeidler, 1989; Garlappi et al., 2007). Their preferences are expressed by the set of possible future expected v alues of the exchange rate determined by a (symmetric) bandwidth δ_I around the base or uncertainty neutral expectation. That is, an uncertainty averse domestic agent assumes that the future exchange rate takes its value in the interval $[E_t(s_{t+1}|I) - \delta_I, E_t(s_{t+1}|I) + \delta_I]$, $I = F, C$. The maximisation problem of such an agent can be written as follows:

$$E_t(U(W^u_{t+1})|I) = \min_{s \in [E_t(s_{t+1}|I)-\delta_I, E_t(s_{t+1}|I)+\delta_I]} E_t(W^u_{t+1}(s)|I) - \gamma E_t((W^u_{t+1}(s))^2|I) \underset{f^u_t}{\rightarrow} \max \tag{5}$$

with respect to budget constraint

$$W^u_{t+1}(s_{t+1}) = (1+r_t)(W^u_t - s_t f^u_t) + s_{t+1}(1+\rho_t) f^u_t.$$

Let us denote

$$C(I) = \frac{E_t(s_{t+1}|I)(1+\rho_t) - s_t(1+r_t)}{2\gamma E_t(s^2_{t+1}|I)(1+\rho_t)^2},$$

$$C_{\max}(I) = \frac{(E_t(s_{t+1}|I) + \delta_I)(1+\rho_t) - s_t(1+r_t)}{2\gamma E_t(s^2_{t+1}|I)(1+\rho_t)^2},$$

$$C_{\min}(I) = \frac{(E_t(s_{t+1}|I) - \delta_I)(1+\rho_t) - s_t(1+r_t)}{2\gamma E_t(s^2_{t+1}|I)(1+\rho_t)^2}.$$

Lemma 2. *Given the level of exchange rate s_t the optimal strategy for an uncertainty averse agent is to hold f^u_t units of foreign currency, where*

$$f^u_t = \begin{cases} C_{\min}(I) & \text{if } s_t < E_t(s_{t+1}|I) - \delta_I, \\ 0 & \text{if } E_t(s_{t+1}|I) - \delta_I \leqslant s_t \leqslant E_t(s_{t+1}|I) + \delta_I, \\ C_{\max}(I) & \text{if } E_t(s_{t+1}|I) + \delta_I < s_t \end{cases}$$

for $I = F, C$

See the Appendix for the proof.

1.2. Learning through social interaction

Consistent with observed behaviour in the markets investors may change the way they make their decisions at every period of time, as discussed in Menkhoff and Taylor (2007). They may switch the way they form expectations about future exchange rates (become fundamentalists or chartists) and also their reaction to uncertainty in the market can change. The learning mechanism of agents we employ is similar to case-based reasoning and is based on the cumulative gain of particular groups of agents and a comparison with the past experience of other investors. This sort of updating is implemented in heterogeneous agents models by, *inter alia*, Kirman (1993), Kirman et al. (2007), De Grauwe and Grimaldi (2006) and Boswijk et al. (2007).

At the beginning of period t, agents compare the realised utilities of the different strategies and invest into those which perform better. To be more precise, each trader chooses a strategy in a probabilistic manner where the probability of choosing a strategy depends on its past performance. According to this model the probability of an investor becoming a fundamentalist at time t can be calculated as

$$P_{t+1}(F) = \frac{e^{\beta G^n_t(F)}}{e^{\beta G^n_t(F)} + e^{\beta G^n_t(C)}},$$

where $G_t^n(F)$ and $G_t^n(C)$ are discounted sums of the one-period utilities of the uncertainty neutral fundamentalists and chartists, respectively, given that both types of agent had the same initial level of wealth. That is,

$$G_t^n(I) = \sum_{j=1}^{m} \omega^{j-1} U(g_{t-j+1}^n(I))$$

with $g_t^n(I) = (1 + r_{t-1})(W_{t-1} - s_{t-1} f_{t-1}^n(I)) + s_t(1 + \rho_{t-1}) f_{t-1}^n(I)$ and $U(x) = x - \gamma x^2$. The parameter β here is called the intensity of choice (see Boswijk et al., 2007) and ω plays the role of a discount factor.

Simultaneously with changing the method of expectation formation, a trader can change his reaction to the level of uncertainty present in the market. Sentiment indicators are frequently found to be significant in explaining asset returns which reflects the time varying nature of uncertainty in markets. If based on the past performance an uncertainty averse strategy appears better in terms of utility than being uncertainty neutral, traders will become more careful and less aggressive as specified by the maxmin model. If the information available in the market is treated as "certain", then agents will be more willing to choose the expected utility strategy. Under severe uncertainty the maxmin strategy can be as bad as a "do nothing" strategy while under mild uncertainty it will earn some positive utility and will be less sensitive to bad outcomes. In the same way as described above, the probability of an investor becoming uncertainty neutral is obtained from the following formula:

$$P_{t+1}(n, I) = \frac{e^{\beta G_t^n(I)}}{e^{\beta G_t^n(I)} + e^{\beta G_t^u(I)}},$$

where $G_t^n(I)$ and $G_t^u(I)$ are discounted sums of the one-period utilities of corresponding uncertainty neutral and uncertainty averse strategies, respectively.

More formally,

$$G_t^j(I) = \sum_{j=1}^{m} \omega^{j-1} U(g_{t-j+1}^j(I))$$

with $g_t^j(I) = (1 + r_{t-1})(W_{t-1} - s_{t-1} f_{t-1}^j(I)) + s_t(1 + \rho_{t-1}) f_{t-1}^j(I),\ j = n, u.$

Note that uncertainty averse traders use the utility function U in the same way as their uncertainty neutral opponents to evaluate the past performance of different strategies, since all uncertainty about the past prices has been resolved. Since all types of traders have the same risk preferences, the difference in the performance of strategies is reflected only in the level of demand from uncertainty neutral and uncertainty averse traders.

This learning mechanism ensures that strategies with higher realised utility in the recent past become more attractive to agents. In this way, traders do not systematically make mistakes but learn about the strategy with the highest current performance.

1.3. *The market clearing exchange rate*

In order to be able to define the aggregate demand functions we denote the proportion of fundamentalists in the market by x_t and let y_t^F and y_t^C define the proportions of uncertainty neutral investors among fundamentalists and chartists, respectively. These proportions will change with time according to the probabilities specified above.

As we can see each individual demand function is a function of the expected value of the level of the foreign exchange rate s_t. Depending on past information an investor decides how to build his expectation: based on fundamental variables or by inferring structure from historical exchange rate patterns (or in other words—be fundamentalist or chartist). Let us denote the aggregate demand function at time t by $\Phi_t(s_t)$.

These demand functions can be presented in the form

$$\Phi_t(s_t) = N(x_t y_t^F f_t^{i,n}(F) + x_t(1 - y_t^F) f_t^{i,u}(F) + (1 - x_t) y_t^C f_t^{i,n}(C) + (1 - x_t)(1 - y_t^C) f_t^{i,u}(C)).$$

The market clearing exchange rate s_t^* is then a solution to the market clearing equation

$$\Phi_t(s_t^*) = 0, \tag{6}$$

that is, aggregate demand must be equal to aggregate supply which is held at 0 for simplicity. We denote the pricing function by $\tilde{\mathbf{f}}(\mathscr{F}_t) = \Phi_t^{-1}(0)$, where $\mathscr{F}_t$ is the information available at time t.[4]

1.4. *The fundamental exchange rate*

Fundamentalists are assumed to form their beliefs about the latent fundamental equilibrium exchange rate, $\bar{s}_t$, based on the uncovered interest parity (UIP) condition. They believe that $\log(\bar{s}_t) = \log(\bar{s}_{t-1}) + \log(r_{t-1}) - \log(\rho_{t-1})$. Notice that UIP is

[4] If the function Φ is not invertible, we can treat $\tilde{\mathbf{f}}$ as a generalised inverse considering the solution that minimises Φ (the distance minimising solution).

only used to determine the latent equilibrium exchange rate since it is well known that market rates fail to satisfy the UIP condition. In fact, our model could provide one explanation for the existence of the carry trade through the presence of uncertainty averse heterogeneous agents in the market. Since the fundamental exchange rate is unobservable, fundamentalists are assumed to use a state space estimation method to form their estimate of $\bar{s}$. As in the standard Kalman Filter they form prior beliefs and use a Bayesian rule to update these beliefs but because our model is nonlinear we have to move beyond the simple Kalman filter and use what is known as the UKF. Fundamentalists are rational in the way they take into account both fundamental information and the presence of chartists in the market. We assume they use this information to derive the posterior distribution of the fundamental price from which the optimal state estimate can be derived. This estimate is then used together with the realised exchange rate when forming their expectations of the market rate through Eq. (1) and their demands which in turn to determine the market rate along with chartist demands. Assuming a Gaussian error term, the expression for the fundamental price $\bar{s}_t$ will be

$$\log(\bar{s}_t) = \log(\bar{s}_{t-1}) + \log(r_{t-1}) - \log(\rho_{t-1}) + \sigma_{\bar{s}}\xi_t, \tag{7}$$

with $\xi_t \sim N(0,1)$. In the absence of chartists, fundamentalist traders would drive the exchange rate to the fundamental price and it would coincide with the rational expectations equilibrium. Hereafter we use the notation $\bar{s}_t$ to also denote the estimate of the fundamental rate where it does not lead to confusion.

Chartists are boundedly rational since while recognising the presence of fundamentalists they choose not to take any fundamental information into account.[5] They simply believe in their chartist rule and use it to extrapolate an exchange rate forecast which therefore systematically deviates from the fundamental exchange rate.

The timing of the model is as follows. Traders, based on the recent performance of different strategies decide upon their type and their attitude to uncertainty. Given these choices the proportion of agent types is determined. At that time, those traders who decided to use the chartist strategy form their expectations of the future exchange rate by extrapolating past prices and compute their demand throughout the maximisation problem (3). Fundamentalists update their beliefs about the fundamental rate taking into account current interest rates and form their expectations and demand functions. The market then clears at the market clearing price level at the end of the day which provides the model's output for the level of the exchange rate from the solution to Eq. (6).

2. Estimation

We now describe the estimation of the latent fundamental rate and the parameters of the model. Given that the model is highly nonlinear, the latent fundamental price is estimated using an UKF, which is described in the following subsection. Having then estimated the fundamental price we use NLS to estimate the parameters of the model.

2.1. Unscented Kalman filter

The UKF allows us to deal directly with the nonlinearities present in the model without approximation or linearisation using a form of particle filter and hence provides a much more accurate estimate of the evolution of the nonlinear stochastic process than the standard Kalman filter or the extended Kalman filter. This algorithm was proposed by Julier and Uhlmann (2004) and allows us to solve the problem of nonlinear filtering using nonlinear transformations of Gaussian distributions. We give a brief sketch of the UKF algorithm below; a more extensive description can be found in Julier and Uhlmann (2004) and Van der Merwe (1998).

We are interested in estimating a nonlinear model of the form

$$\mathbf{y}_t = \mathbf{f}(\mathbf{u}_t, \mathbf{x}_t, \mathbf{n}_t), \tag{8}$$

$$\mathbf{x}_t = \mathbf{h}(\mathbf{x}_{t-1}, \mathbf{u}_t, \mathbf{w}_t), \tag{9}$$

where $\mathbf{y}_t \in \mathbb{R}^{n_y}$ denotes the observable output time series, in our case the level of the market clearing exchange rate s_t, $\mathbf{u}_t \in \mathbb{R}^{n_u}$ the input data (interest rates and past exchange rates $[\rho_{t-1}, r_{t-1}, s_{t-1}]^T$), $\mathbf{x}_t \in \mathbb{R}^{n_x}$ the unobserved state of the system, in our case the log of fundamental rate $\log(\bar{s}_t)$, $\mathbf{w}_t \in \mathbb{R}^{n_w}$ the process noise and $\mathbf{n}_t \in \mathbb{R}^{n_n}$ the measurement noise. The functions $\mathbf{f}$ and $\mathbf{h}$ are said to represent the output measurement and state transition models, respectively; (8) represents the solution to (6) for the observed market clearing exchange rate and (9) corresponds to the UIP condition determining the fundamental exchange rate (7).

The unscented transformation method directly evaluates the nonlinearities in $\mathbf{f}$ and $\mathbf{h}$ through a sequential Monte Carlo simulation and directly calculates the first two moments of $\mathbf{x}$ and $\mathbf{y}$ from the nonlinear system itself. Let us denote by $\mathbf{x}_t^a = [\mathbf{x}_t^T \mathbf{w}_t^T \mathbf{n}_t^T]^T$ and by $\mathbf{x}_{t|t}^a$ and $\mathbf{P}_{t|t}^a$ the mean and covariance of $\mathbf{x}^a$ at time t. In order to provide a transformation, a set of $2n_a + 1$ weighted samples of "sigma" points $\mathcal{S}_i = \{\mathbf{W}_i, \mathcal{X}_i\}$, $n_a = n_x + n_w + n_v$ are chosen so that they completely capture

[5] This again reflects reality since chartist methods are typically used in short run decision making (i.e., intraday) when no new fundamental information may be available.

the true mean and covariance of the prior random variable $\mathbf{x}_t^a$. This may be carried out as follows:

$$\begin{aligned}
\mathcal{X}_0^a &= \mathbf{x}_{t|t}^a,\\
\mathcal{X}_{i,t|t}^a &= \mathbf{x}_{t|t}^a + \left(\sqrt{(n_a+\lambda)\mathbf{P}_{t|t}^a}\right)_i, \quad i = 1,\dots,n_x,\\
\mathcal{X}_{i,t|t}^a &= \mathbf{x}_{t|t}^a - \left(\sqrt{(n_a+\lambda)\mathbf{P}_{t|t}^a}\right)_i, \quad i = n_x+1,\dots,2n_x,\\
W_0^{(m)} &= \frac{\lambda}{n_a+\lambda},\\
W_0^{(c)} &= \frac{\lambda}{n_a+\lambda} + (1-\alpha^2+\beta),\\
W_i^{(m)} &= W_i^{(c)} = \frac{1}{2(n_x+\lambda)}
\end{aligned}$$

with $\lambda = \alpha^2(n_a+\kappa) - n_a$, $\kappa \geqslant 0$, $0 \leqslant \alpha \leqslant 1$ and $\beta \geqslant 0$. Here $(\sqrt{(n_a+\lambda)\mathbf{P}_{t|t}^a})_i$ denotes the ith column of the matrix square root of $(n_a+\lambda)\mathbf{P}_{t|t}^a$. In our implementation we set $\kappa = 2$, $\alpha = 0.9$, $\beta = 2$.

The prediction step of the UKF can be sketched in the following way:

$$\begin{aligned}
\mathcal{X}_{t+1|t}^x &= \mathbf{h}(\mathcal{X}_{t|t}^x, \mathcal{X}_{t|t}^w),\\
\mathbf{x}_{t+1|t} &= \sum_{i=1}^{2n_a+1} W_i^{(m)} \mathcal{X}_{i,t+1|t}^x,\\
\mathbf{P}_{t+1|t} &= \sum_{i=1}^{2n_a+1} W_i^{(c)} [\mathcal{X}_{i,t|t+1}^x - \mathbf{x}_{t+1|t}][\mathcal{X}_{i,t|t+1}^x - \mathbf{x}_{t+1|t}]^T,\\
\mathcal{Y}_{t+1|t} &= \mathbf{f}(\mathcal{X}_{t+1|t}^x, \mathcal{X}_{t|t}^n),\\
\mathbf{y}_{t+1|t} &= \sum_{i=1}^{2n_a+1} W_i^{(m)} \mathcal{Y}_{i,t+1|t}.
\end{aligned}$$

The state update equations are as follows:

$$\begin{aligned}
\mathbf{P}_{\mathbf{xy}} &= \sum_{i=1}^{2n_a+1} W_i^{(c)} [\mathcal{X}_{i,t|t+1}^x - \mathbf{x}_{t+1|t}][\mathcal{Y}_{i,t|t+1} - \mathbf{y}_{t+1|t}]^T\\
\mathbf{P}_{\mathbf{yy}} &= \sum_{i=1}^{2n_a+1} W_i^{(c)} [\mathcal{Y}_{i,t|t+1} - \mathbf{y}_{t+1|t}][\mathcal{Y}_{i,t|t+1} - \mathbf{y}_{t+1|t}]^T,\\
\mathbf{K}_{t+1} &= \mathbf{P}_{\mathbf{xy}}\mathbf{P}_{\mathbf{yy}}^{-1},\\
\mathbf{x}_{t+1|t+1} &= \mathbf{x}_{t+1|t} + \mathbf{K}_{t+1}(\mathbf{y}_t - \mathbf{y}_{t+1|t}),\\
\mathbf{P}_{t+1|t+1} &= \mathbf{P}_{t+1|t} - \mathbf{K}_{t+1}\mathbf{P}_{\mathbf{yy}}\mathbf{K}_{t+1}^T.
\end{aligned}$$

2.2. NLS estimation

The UKF method allows us to write the model in the form

$$\mathbf{y}_t = \mathbf{y}_{t+1|t}(\theta) + \sigma\varepsilon_t,$$

where $\mathbf{y}_t$ is the observed time series (exchange rate levels in our case) and ε_t are independent identically distributed random variables. The parameters $\theta \in \Theta$ of the model are then estimated (in a distribution free manner) by minimising the sum of squared errors

$$SE = \frac{1}{N}\sum_{i=0}^{N} (\mathbf{y}_t - \mathbf{y}_{t+1|t}(\theta))^2 \underset{\theta\in\Theta}{\rightarrow} \min.$$

The generalised NLS estimates are asymptotically normally distributed

$$\sqrt{N}(\hat{\theta} - \theta) \overset{a}{\sim} N(\mathbf{0}, \Omega),$$

where $\Omega = A^{-1}BA^{-1}$ is the White–Domowitz robust covariance matrix with

$$A = 2N^{-1}\sum_{t=1}^{N} E(\nabla\mathbf{y}_t'\nabla\mathbf{y}_t) \quad \text{and}$$

$$B = 4N^{-1}\sum_{t=1}^{N} E(\varepsilon_t^2 \nabla \mathbf{y}_t' \nabla \mathbf{y}_t) + 4N^{-1}\sum_{\tau=1}^{N-1}\sum_{t=\tau+1}^{N} E(\varepsilon_t \varepsilon_{t-\tau}[\nabla \mathbf{y}_t' \nabla \mathbf{y}_{t-\tau} + \nabla \mathbf{y}_{t-\tau}' \nabla \mathbf{y}_t])$$

(see White and Domowitz, 1984 for details).

3. Empirical results

As we have said the function $\mathbf{f}$ in the measurement equation (8) solves the market clearing condition (6); that is, the measurement equation is $s_t = \tilde{\mathbf{f}}(\mathbf{u}_t, \bar{s}_t) + \tilde{\varepsilon}_t$. The function $\mathbf{h}$ in the state transition equation is in the form of (7) with $\sigma_{\bar{s}}\xi_t$ being a noise. Inputs to the system are the interest rates r_t and ρ_t as well as past realised prices $s_{t-j}, j = 1, \ldots, t-1$.

The proxies we use for the domestic and foreign risk-free rates r and ρ are the UK and US Interbank LIBOR overnight rates, respectively, and we use daily closing rates on the GBP/USD exchange rate over the period from 2 January 1997 till 30 June 2008, as shown in Fig. 1.

For identification we need to tie down several parameters and since we are more interested in the degree of uncertainty aversion in the market we fix the risk aversion coefficient to be $\gamma = 2$. Under the assumption of identical risk aversion for both types of trader, this coefficient in fact has no influence on the equilibrium price in our model. The set of parameters remaining for estimation then consists of $\theta = \{\delta_F, \delta_C, \sigma_{\bar{s}}, v, h, \sigma, \beta\}$. As our main research question is to test for uncertainty aversion in traders, we are in particular interested in estimating the parameters δ_F and δ_C and testing if they are significantly different from zero. If these parameters are jointly insignificant this would imply that observed spot exchange rates are generated from a model where only uncertainty neutral agents interact. The coefficients v and h are adjustment parameters in the expectation formation process for both types of agents. Heterogeneity in expectations is a main driver of the model and these parameters along with the intensity of choice parameter β are the most important parameters for the benchmark model without uncertainty averse agents. The two remaining parameters σ and $\sigma_{\bar{s}}$ need to be estimated as they are input values to the UKF.

The converged NLS estimates of the parameters, their standard deviations and the p-values of tests of significance are given in Table 1.

Based on these estimation results we conclude that six of the parameter estimates are significantly different from 0 while δ_F is not. These results indicate strong uncertainty aversion in chartists but not fundamentalists since the estimated parameter δ_C is significantly different from zero and hence the null hypothesis is rejected. Uncertainty aversion therefore appears to be an important feature of at least some traders' behaviour in the foreign exchange market. One intuitive interpretation for this result could be as follows; fundamentalists have a strong belief in one *economic* model and their expectations are tied down by the exogenous interest rates through $\bar{s}$ but chartists use an "*ad hoc*" time series model driven by past realised rates and hence are naturally more uncertain as to the true model driving the fundamental exchange rate. The average proportion of uncertainty averse agents in the market is 21.56% which indicates a significant impact on the realised exchange rate.

In order to test the hypothesis of no uncertainty aversion in the market we specify the joint null hypothesis that all traders in the market are uncertainty neutral (i.e. they are simply quadratic utility maximisers) by

$$H_0\colon \delta_F = \delta_C = 0$$

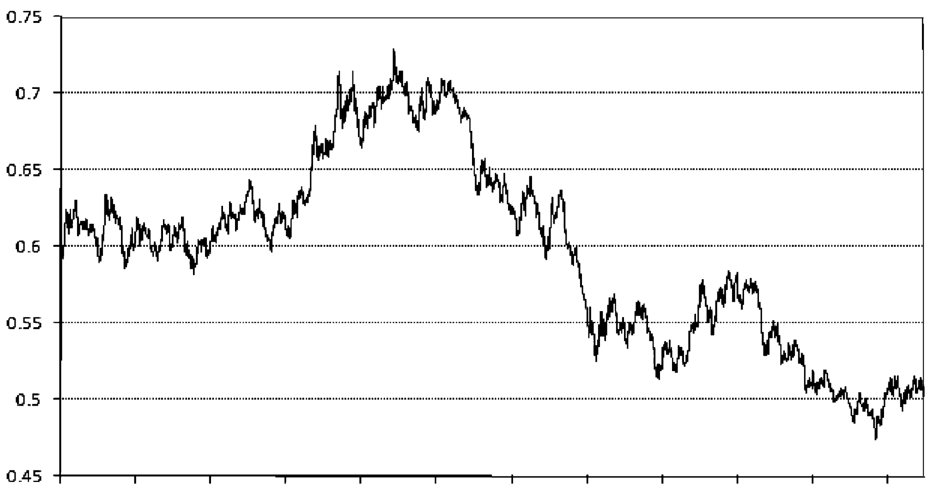

Fig. 1. Daily GBP/USD prices during 2 January 1997–30 June 2008 (top) and the estimated model residuals (bottom).

Table 1
Parameters estimates: unrestricted model.

Parameters	Estimates	Standard deviation	p-Values
δ_F	1.35677×10^{-9}	3.98765×10^{-3}	>0.5000
δ_C	0.01652	1.56621×10^{-3}	<0.0001
$\sigma_{\bar{s}}$	0.00527	2.12940×10^{-4}	<0.0001
ν	0.99868	8.98731×10^{-2}	<0.0001
h	0.07832	1.33536×10^{-2}	<0.0001
σ	0.00120	2.45013×10^{-4}	<0.0001
β	0.01374	5.59801×10^{-3}	<0.0001

Parameter estimates of the model based on the nonlinear least squares estimation method (column 2). Columns 3 and 4 contain standard deviations of estimates and their p-values of the test for significance.

Table 2
Parameters estimates: restricted model.

Parameters	Estimates	Standard deviation	p-Value
$\sigma_{\bar{s}}$	0.01311	4.78994×10^{-4}	<0.0001
ν	0.99982	1.59371×10^{-2}	<0.0001
h	0.11988	6.33276×10^{-3}	<0.0001
σ	0.00443	1.65087×10^{-4}	<0.0001
β	0.03721	1.35288×10^{-3}	<0.0001

Parameter estimates of the model based on the nonlinear least squares estimation method (column 2) under the assumption $\delta_F=0$ and $\delta_C=0$. Columns 3 and 4 contain standard deviations of estimates and their p-values of the test for significance.

with the alternative

$$H_1: \delta_F > 0, \quad \delta_C > 0.$$

We test this hypothesis using a Wald test on the restriction $\delta_F=\delta_C=0$. The corresponding F statistic is computed by

$$F = \frac{(RSS_{\delta_C=\delta_F=0}(\hat{\theta}_{3:7}) - RSS_U(\hat{\theta}))/2}{\hat{\sigma}^2} \overset{a}{\sim} \chi^2_2,$$

where $RSS_{\delta_C=\delta_F=0}(\hat{\theta}_{3:7})$ and $RSS_U(\hat{\theta})$ are residual sum of squares of the restricted and unrestricted models and $\hat{\sigma}^2$ is the estimate of the variance of the residuals from the unrestricted model. Estimates of the parameters of the restricted model are given in Table 2. The value of the test statistic is $F=31.08$ which is substantially larger than 5% critical value $\chi^2_{0.05,2}=5.99$. Thus, we reject the null hypothesis of no uncertainty aversion within traders in the FX market.

3.1. Model evaluation

In order for us to have confidence in the test results reported above we need to be sure that the model provides a reasonably good representation of reality (although we have used robust standard errors throughout the inference). We proceed to evaluate the model in three ways; first we compute standard residual diagnostic tests, then we compare the model with the random walk model in an out-of-sample prediction exercise and then finally we carry out two further prediction based tests to examine if there is any exploitable structure in the model's output based on testing directional change predictability and economic value based on using forecasts from the model in a simple trading rule.

Before discussing these results we would like to stress that it is logically impossible to formally validate this model using just aggregate exchange rate data.[6] The model we have built has a complex "micro" structure which leads to an aggregate or macro output in terms of the level of observed exchange rate. The question of model validation rests on whether we can uniquely identify the underlying "micro" model or whether in fact there may be a number of micro models that have the same implications for the observed aggregate exchange rate. The issue is one of model identification as opposed to structure identification as originally discussed by Preston (1978) and more recently by Hendry et al. (2002). Standard parametric identification ensures there is a unique parametrisation of a structure and we are confident in that since we have a nonsingular parameter variance covariance matrix (see Rothenberg, 1971). However, parameter identification conditions do not rule out that there may be several different models which have the same observable

[6] We would like to thank a referee for emphasising this point to us.

implication. This is the issue we face in trying to validate our model and in fact it affects much econometric research although it is generally unacknowledged.

So while we use prediction tests below we note that a model without predictive ability could be correct—if the correct model is a random walk and secondly, in this particular case we cannot claim that predictive ability in the aggregate exchange rate necessarily implies that our heterogeneous agent model is the unique correctly specified data generation process. We interpret the following test results therefore as simply providing some indication that our model is not inconsistent with the data.

We first examine the usual residual specification tests. Standard goodness of fit measures indicate an R^2 value of 0.98. The residuals from the regression do not show significant evidence of autocorrelation (the p-value of Breusch–Godfrey serial correlation LM test is 0.1319). Homoscedasticity is not rejected by Breush–Pagan's test at 5% significance level (p-value is 0.0881) and the augmented Dickey–Fuller tests for unit root rejects the nonstationarity of residuals (p-value is <0.0001). So we have some confidence the model structure is capturing much of the structure seen in the level of the exchange rate.

We now compare the ability of our model to predict out-of-sample compared to a random walk, following the classic paper by Meese and Rogoff (1983). We provide several tests on the predictive power of the model in a recursive one step ahead prediction mode.

First, the root mean-square error of our model's out-of-sample forecasts is 0.003147 while that for the random walk model is 0.005093. The resulting Diebold–Mariano statistic is -21.58 (p-value <0.0001) clearly indicating superior predictive performance from our model compared to a random walk.

We next test the ability of the model to make correct one step-ahead predictions of the directional change of the exchange rate compared to random selection as well as a test of the significance of the economic value of these predictions. We use the Pesaran–Timmerman test (see Pesaran and Timmermann, 1992, 1994) to examine directional change and the Anatolyev–Gerko test (see Anatolyev and Gerko, 2005) to assess the economic significance of the predicted returns. The Pesaran–Timmermann statistic is used to test the null of no market timing or that the proportion of correct sign predictions equals the proportion which can be expected under the null of independence. The Anatolyev and Gerko test is one of no mean predictability (i.e. independence from past exchange rate returns) and is based on both market timing and the generation of significant profit using a trading rule. Essentially this is a Hausman test that compares two estimates of mean returns from a simple trading rule, both of which will be consistent under the null of no predictability but will differ under the alternative (of dependence on past returns). The formal testing procedures are described in the Appendix.

The null hypothesis underlying both Pesaran–Timmermann and the Anatolyev–Gerko tests is essentially that the predictor from the model again cannot beat a random walk strategy. Rejection means that the model can predict significantly better than the baseline random walk model and hence we would expect it to dominate virtually all the standard macro based models of the exchange rate that were considered by Meese and Rogoff.

Following Kozhan and Salmon (2007), we consider different trading strategies which take into account the presence of uncertainty. We examine predictability of the exchange rate by measuring when the predicted value of the exchange rate deviates from the previous value by more than a given value k. The trading strategy based on the forecast is then to trade only if the forecast value for the next day's exchange rate is larger than the current price level by more than k and not to trade otherwise. The rationale behind this is reflected in a so-called no-trade condition (see Dow and Werlang, 1992). If a forecast value does not differ from the current level of the exchange rate by more than k, then the trader does not believe that there is a clear enough signal on which to trade, in other words there is too much uncertainty in the market.

Table 3 presents the results of the above tests for different values of k. We see that the Pesaran–Timmerman test rejects the null of no directional predictability at the 10% significance level for various values of the threshold variable and the Antolyev–Gerko test rejects the null hypothesis for all of k's at the 5% significance level. Moreover, as k increases, the model's predictive power increases in terms of economic value. This supports our contention that small predictive signals

Table 3
Tests for model's predictive power.

k	Dir. ch.	*Av. returns*	P–T stat.	p-Value	A–G stat.	p-Value	n
0	51.00	0.0191	1.661	0.0483	2.075	0.0189	2898
0.0001	51.42	0.0223	1.348	0.0887	2.422	0.0077	2001
0.0002	52.03	0.0315	1.530	0.0628	3.405	0.0003	1278
0.0003	52.36	0.0520	1.356	0.0875	5.365	0.0000	846
0.0004	54.06	0.0722	1.958	0.0250	7.394	0.0000	566
0.0005	54.20	0.0693	1.713	0.0433	7.462	0.0000	404
0.0006	53.45	0.0680	1.144	0.1262	7.275	0.0000	275
0.0007	52.15	0.0780	0.621	0.2670	7.275	0.0000	209

The trading strategy is based on, to trade only if the forecast value for the next day's exchange rate is larger than the current price level by more than k and not to trade otherwise. The percentage of correct directional predictions and average daily returns based on the strategy are given in columns 2 and 3 for the corresponding values of k. Columns 4 and 5 present values of statistic and p-values for Pesaran–Timmermann test, columns 6 and 7 provide results of Anatolyev–Gerko test. Column 8 shows numbers of transactions during the horizon.

are noisy because of uncertainty in the market while large signals are informative and have predictive power. This indication of predictability, taking account of the k uncertainty band, provides further empirical validation of the model over the random walk. Once again this suggests the model is better supported by the data than the standard macro fundamental models considered by Messe and Rogoff and is also consistent with the finding of uncertainty aversion in the market.

4. Discussion and interpretation

The most important result we seem to have found is that at least some traders in the FX market appear to be uncertainty averse. The estimation results show that the inclusion of uncertainty averse agents improves the performance of the model (through the F test) and the uncertainty aversion parameter for chartists is significantly different from zero. However, traders do not remain equally uncertainty averse throughout the sample. As we can see in Fig. 2 there are periods of higher and lower uncertainty aversion in the market, but since fundamentalists are uncertainty neutral (as δ_F is insignificantly different from zero) these periods are highly correlated with periods of chartist activity.

Fig. 2 plots the fraction of the different types of agent in the market. The majority of the time fundamentalists dominate chartists—the average fraction of fundamentalists over the whole period is 62.4%. This means that exchange rate forecasts based on fundamentals (interest rates) are more precise and more profitable than the trend following approximation we have considered the chartists use. Moreover, the precision of these forecasts is quite tight which allows fundamentalists to be uncertainty neutral and more confident about their predictions. At the same time, it is chartists who are found to be mainly uncertainty averse.

Let us look at behaviour of traders more closely and focus the analysis on two subperiods (from January 1999 to December 1999 and from January 2007 to December 2007). If the fundamental price does not show clear trends then chartist expectations become less precise which leads to the increase in a number of uncertainty averse chartists in the market. As the demand of uncertainty averse traders is smaller than the demand of uncertainty neutral traders, chartists trade less actively during this period of time. However, given unexpected changes in the fundamental rate, fundamentalists make errors in their predictions and lose money. Evolutionary pressure will then cause fundamentalists to switch towards using chartists strategies, which increases the proportion of chartists in the market. In addition, in trying to push the price back to its fundamental value, fundamentalists create short-term trends in the exchange rate. Chartists pick up on these trends, make money on them and their weight in the market expands. Once the exchange rate approaches the fundamental level and uncertainty in the market vanishes, the fundamentalists' forecast becomes more precise once again and the fraction of chartists immediately drops. This behaviour can be clearly observed in Fig. 3. Spikes in fundamentalist activity mainly correspond to reversals in exchange rate trends when the chartist strategy becomes unprofitable. At that time chartists are mainly uncertainty averse and that helps them survive in the market. As uncertainty averse traders consider the worst-case scenario, their demands are relatively small and they do not suffer from big losses. As long as trends become clearer and stronger the fraction of chartists increases. As uncertainty about the behaviour of the exchange rate in the short-term decreases, more and more chartists switch to an uncertainty neutral strategy. If changes in the exchange rate happen unexpectedly but in the same direction as the existing trend, such as in October 2007, we observe a double effect—chartists benefit as the trend becomes stronger and fundamentalists lose because a gap exists between their beliefs

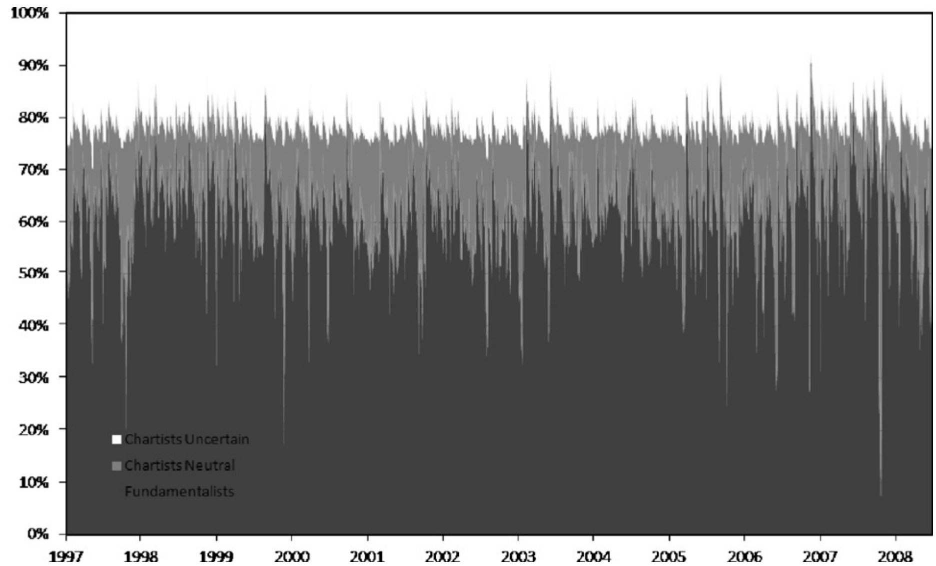

Fig. 2. Proportions of traders in the market—fundamentalists, chartists uncertainty averse and chartists uncertainty neutral.

about the fundamental price and the exchange rate level. As a consequence of these events, we see a large switch towards the chartist strategy and a temporary decrease in fundamentalists' activity Figs. 3 and 4.

The second example (see Figs. 5 and 6) shows the coexistence of fundamentalists and chartists in the market. The relative fractions of agent types do not change much if there are trends in the fundamental price (from the middle of June 1999 to the middle of August 1999). Both fundamentalist and chartist strategies are profitable as they can simultaneously exploit trends in the exchange rate and the movement of the exchange rate towards the fundamental price. This leads to approximately equal revenues for both chartists and fundamentalists: the former predict a trend and the latter follow the fundamental price. During this trend period in the fundamental price we observe the persistent presence of chartists in the market. The difference in wealth during this period is less volatile than usual which brings some stability to the proportion of traders in the market.

We can also observe in Figs. 5 and 6 that even in the absence of clear trends in the exchange rate at the end of November 1999, a large difference between the fundamental rate and the realised exchange rate immediately causes a big drop in the proportion of fundamentalists in the market.

The important observation is that the majority of chartists have been found to be uncertainty averse while the degree of uncertainty aversion among fundamentalists is almost zero. Periods when chartists are more confident in the market when the degree of uncertainty is low are usually at the end of active chartist periods. The intuition behind this behaviour is that once technical traders become more powerful (their proportion increases) they create trends effectively throughout their own herding behaviour. At the same time, as this trend becomes clear chartists become more confident in their predictions. Hence they use point predictors to make money rather than interval-based forecasts which would reflect uncertainty.

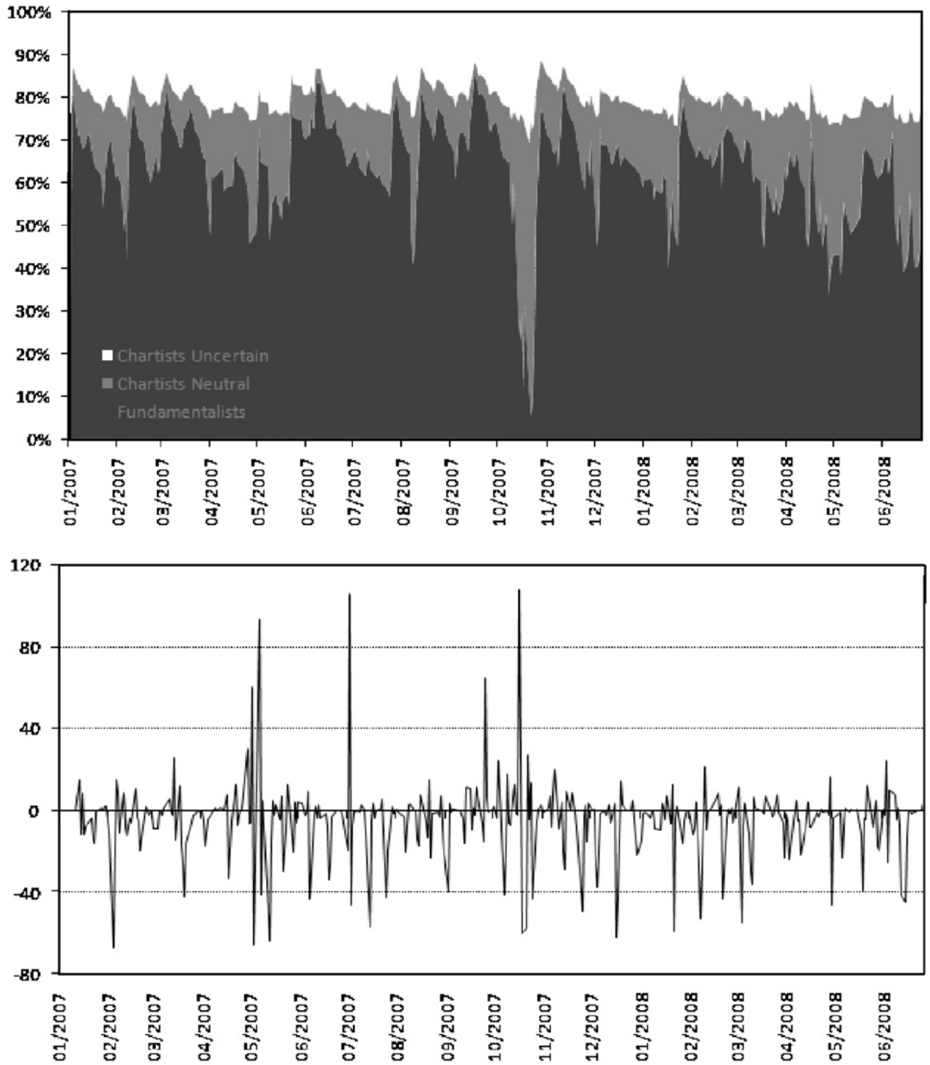

Fig. 3. Proportion of trader types, exchange rate levels vs. fundamental price, interest rate differential and difference in utilities between fundamentalist and chartist traders over period from 2 January 2007 to 30 June 2008.

Fig. 4. Exchange rate levels versus fundamental price over period from 2 January 2007 to 30 June 2008.

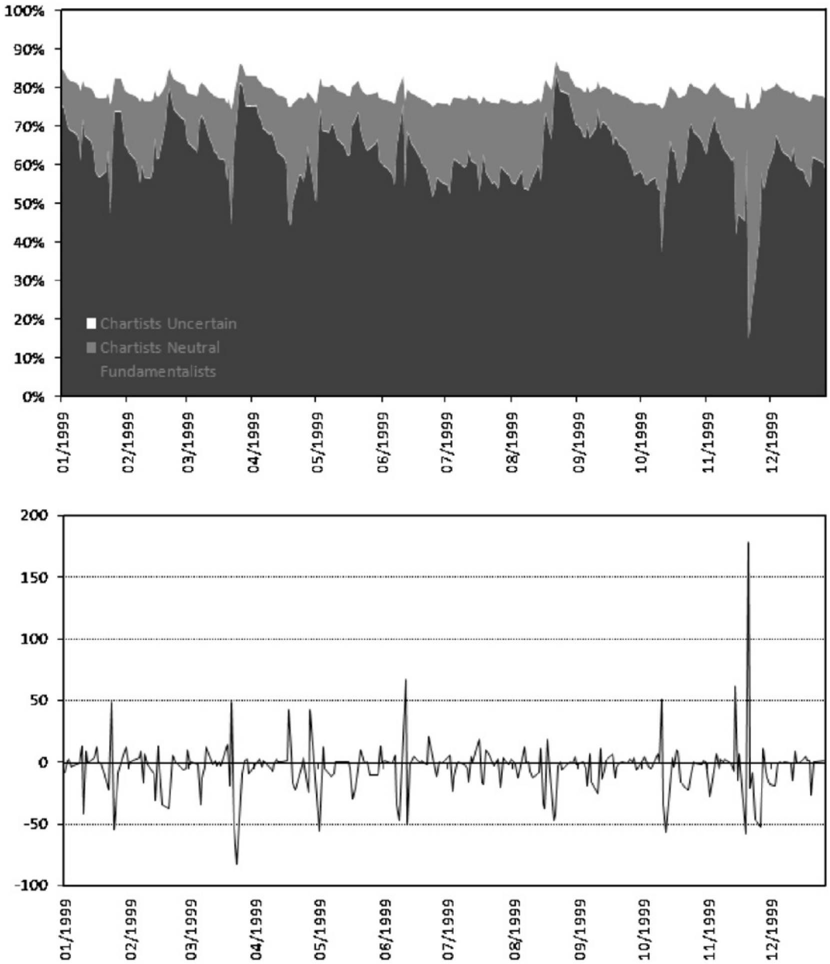

Fig. 5. Proportion of trader types, exchange rate levels vs. fundamental price, interest rate differential and difference in utilities between fundamentalist and chartist traders over period from 2 January 1999 to 29 December 1999.

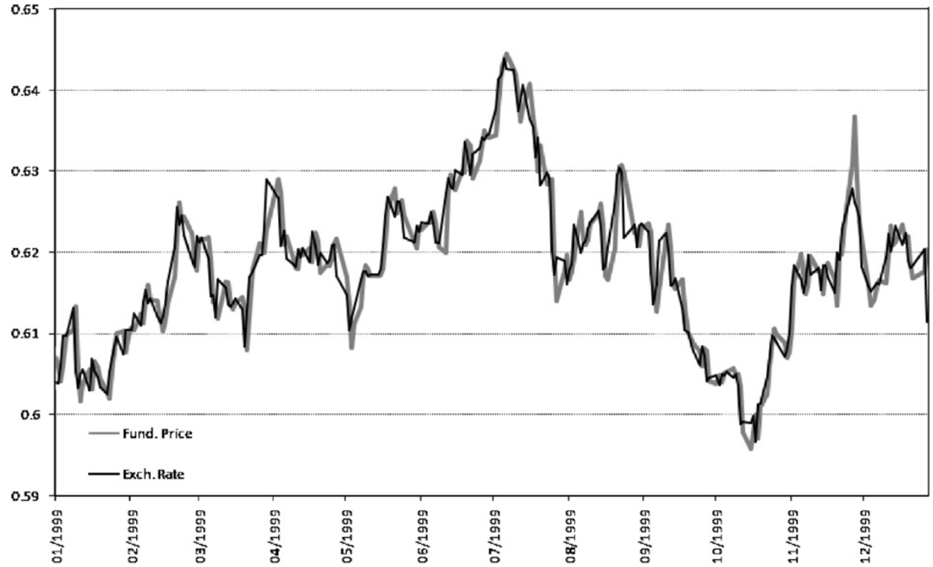

Fig. 6. Exchange rate levels versus fundamental price over period from 2 January 1999 to 29 December 1999.

The proportion of uncertainty averse chartists is relatively stable over the whole period of time. This is explained by the fact that uncertainty averse traders are prepared for the worst possible outcome and are less sensitive to large negative shocks in the exchange rate. The proportion of uncertainty neutral traders varies over time and depends highly on market conditions.

From the estimation results we see that the parameter δ_F, which reflects uncertainty aversion of fundamentalist traders, is almost zero so we are not able to distinguish the performance of uncertainty neutral fundamentalists from the performance of uncertainty averse fundamentalists; they are essentially the same. So in fact their proportions in the market are identical and for this reason we present all fundamentalists as being uncertainty neutral. If fundamentalists predict the future rate incorrectly and face losses it does not make any difference to them to switch to an uncertainty averse strategy as it will not improve their performance.

5. Conclusions

In this paper we provide what we believe to be the first formal test of uncertainty aversion within traders in an FX market using observed daily GBP/USD data over a 10 year period 2 January 1997 to 30 June 2008. We have developed a model of exchange rate formation with uncertainty averse investors that can either hold fundamental or chartist beliefs which has been estimated using NLS and UKF techniques. The estimation results indicate the statistical significance of uncertainty aversion within the market and in particular we find that fundamentalists are largely uncertainty neutral while chartists are uncertainty averse. We have also shown through a range of statistical tests that the model is not inconsistent with observed data and dominates the random walk model for exchange rates. The activity of chartists increases during periods showing clear trends in the level of the exchange rate and they become more confident (uncertainty neutral) when these trends are long. As soon as any trend breaks down the majority of traders switch back to the fundamentalist strategy.

The approach proposed in the paper has several limitations. Chartists are usually more sophisticated than simply trend followers and therefore their forecasts might be more precise than suggested by the model. This is especially true due to the dramatic increase in algorithmic trading among investors and the use of such techniques as genetic programming. Also, our model has been estimated using daily data while in real market prices are obviously determined on an tick by tick basis.

Despite these limitations the paper provides a first step to detecting and testing the behavioural attitude of FX traders to uncertainty by providing what may be the first empirical test. It therefore extends the existing literature of decision making under uncertainty and provides an approach to rigorously examining further models on real data.

We have not discussed the implications of the model for the carry trade but it is clear from the structure of the model that the UIP condition will be violated by the realised market clearing exchange rates for considerable periods of due to the presence of uncertainty averse heterogeneous agents in the market.

Acknowledgements

We would like to thank the three anonymous referees, associate editors of the special issue William Branch and Mikhail Anufriev, Cars Hommes, Reiner Franke, Thomas Lux and all participants at the EU NEST Workshop on "Complexity in

Financial Markets" held at Marseille in April 2007 for their comments on an earlier version of this paper. This paper was supported by an EU NEST STREP, FP6 Grant on "Complex Markets", no. 516446.

Appendix A

Proof of Lemma 1. This lemma is an easy consequence of Lemma 2 letting $\delta_I = 0$. □

Proof of Lemma 2. Given $\tilde{s}$ one can rewrite the expected utility of the terminal wealth as

$$I(\tilde{s}) = ((1+r_t)(W_t - s_t f_t^u) + \tilde{s}(1+\rho_t)f_t^u - \gamma(1+\rho_t)^2 f_t^u \sigma^2,$$

where $\sigma^2 = E(s_{t+1}^2)$.

The explicit form of the preference functional

$$V(W_{t+1}(f_t^u)) = \min_{\tilde{s}\in[E_t(s_{t+1}|I)-\delta_I, E_t(s_{t+1}|I)+\delta_I]} I(\tilde{s})$$

can be found through the minimisation problem $I(\tilde{s}) \underset{\tilde{s}\in[E_t(s_{t+1}|I)-\delta_I, E_t(s_{t+1}|I)+\delta_I]}{\rightarrow} \min$.

The derivative of the functional $I(\tilde{s})$ is

$$\frac{\partial I(\tilde{s})}{\partial \tilde{s}} = (1+\rho_t)f_t^u$$

hence, $\partial I(\tilde{s})/\partial\tilde{s} \geqslant 0$ if $f_t^u \geqslant 0$ and $\partial I(\tilde{s})/\partial\tilde{s} < 0$ if $f_t^u < 0$. Let us denote

$$\tilde{s}(f_t^u) = \begin{cases} E_t(s_{t+1}|I) - \delta_I & \text{if } f_t^u \geqslant 0 \\ E_t(s_{t+1}|I) + \delta_I & \text{if } f_t^u < 0 \end{cases} = \underset{\tilde{s}\in[E_t(s_{t+1}|I)-\delta_I, E_t(s_{t+1}|I)+\delta_I]}{\arg\min} I(\tilde{s}). \tag{10}$$

The expected utility can be rewritten as

$$V(f_t^u) = (1+r_t)(W_t - s_t f_t^u) + \tilde{s}(f_t^u)(1+\rho_t)f_t^n - \gamma(1+\rho_t)^2 f_t^{u2}\sigma^2.$$

At the point $f_t^u = 0$ the preference functional $V(f_t^u) = (1+r_t)W_t$ does not depend on $\tilde{s}(f_t^u)$ and therefore is continuous function on $\mathbb{R}$.

The derivative of the preference functional is given by the expression

$$\frac{\partial V}{\partial f_t^u} = \begin{cases} 2\gamma\sigma^2(1+\rho_t)^2(C_{\min}(I) - f_t^u), & f_t^n > 0, \\ 2\gamma\sigma^2(1+\rho_t)^2(C_{\max}(I) - f_t^u), & f_t^u < 0, \end{cases}$$

where

$$C_{\min}(I) = \frac{(E_t(s_{t+1}|I) - \delta_I)(1+\rho_t) - s_t(1+r_t)}{2\gamma E_t(s_{t+1}^2|I)(1+\rho_t)^2},$$

$$C_{\max}(I) = \frac{(E_t(s_{t+1}|I) + \delta_I)(1+\rho_t) - s_t(1+r_t)}{2\gamma E_t(s_{t+1}^2|I)(1+\rho_t)^2}.$$

Thus,

$$f_t^u = \begin{cases} C_{\min}(I) & \text{if } C_{\min}(I) > 0, \\ C_{\max}(I) & \text{if } C_{\max}(I) < 0, \\ 0 & \text{if } C_{\min}(I) \leqslant 0 \leqslant C_{\max}(I). \end{cases}$$

The statement of the lemma can be easily obtained from the previous equation. □

Pesaran–Timmermann test. Let s_t be the realised value of the exchange rate and let $s_{t|t-1}$ denote its forecast. Define the probabilities

$$P_{11} = P(s_{t|t-1} < 0, s_t < 0), \quad P_{12} = P(s_{t|t-1} < 0, s_t \geqslant 0),$$
$$P_{21} = P(s_{t|t-1} \geqslant 0, s_t < 0), \quad P_{22} = P(s_{t|t-1} \geqslant 0, s_t \geqslant 0).$$

The diagonal elements of this contingency table provide the proportion of correct predictions. P_{ij} denotes the probability of a realisation in the cell of the i'th row and j'th column of the contingency table. In general, the Pesaran–Timmermann test considers a number of categories $i,j \in \{1,\ldots,m\}$; we only need to consider $m = 2$. Denote by $P_{i0} = \sum_{j=1}^m P_{ij}$ the probability of cells in the i'th row and $P_{0j} = \sum_{j=1}^m P_{ij}$ the probability of cells in the j'th column. The null hypothesis is expressed as

$$H_0: \sum_{i=1}^m (\hat{P}_{ii} - \hat{P}_{i0}\hat{P}_{0i}) = 0,$$

which says that the predictor cannot predict significantly more correct directional changes that a random walk predictor (i.e. 50%). Here the probabilities estimates $\hat{P}_{ij}$ are frequencies of the corresponding events observed in the data.

The test is based on the standardised statistic

$$z_n = \sqrt{n} V_n^{-1/2} Z_n \overset{a}{\sim} N(0,1),$$

where n is the number of observations, and

$$Z_n = \sum_{i=1}^{m} (\hat{P_{ii}} - \hat{P_{i0}}\hat{P_{0i}}),$$

$$V_n = \left(\frac{\partial f(\mathbf{P})}{\partial \mathbf{P}}\right)'_{\mathbf{P}=\hat{\mathbf{P}}} (\hat{\Psi} - \hat{\mathbf{P}}\hat{\mathbf{P}}') \left(\frac{\partial f(\mathbf{P})}{\partial \mathbf{P}}\right)_{\mathbf{P}=\hat{\mathbf{P}}}.$$

$\hat{\Psi}$ is an $m^2 \times m^2$ diagonal matrix with $\hat{\mathbf{P}}$ as its diagonal elements,

$$\left(\frac{\partial f(\mathbf{P})}{\partial \mathbf{P}}\right)_{\mathbf{P}=\hat{\mathbf{P}}} = \begin{cases} 1 - P_{0i} - P_{i0} & \text{for } i = j, \\ -P_{j0} - P_{0i} & \text{for } i \neq j. \end{cases}$$

Anatolyev–Gerko test. Let r_t be the observed log-returns of the exchange rate and $r_{t|t-1}$ be their forecasts for $t = 1, \ldots, n$. The forecasts depend on the past information $\mathscr{F}_{t-1} = \{r_{t-1}, r_{t-2}, \ldots\}$. Let the trading rule of the investor be based on the forecast variable $r_{t|t-1}$, in particular, the investor takes a long position in USD if $r_{t|t-1} \geqslant 0$ and a short position in dollars if $r_{t|t-1} < 0$. Then the one-period return from using the trading strategy is $R_t = \text{sign}(r_{t|-1}) \cdot r_t$. The null hypothesis is conditional mean independence so that

$$H_0\colon\ E(r_t|\mathscr{F}_{t-1}) = \text{const}$$

or that $r_{t|t-1}$ and r_t are independent. The expected one-period return $E(R_t)$ can be consistently estimated under the null by two estimators:

$$A_n = \frac{1}{n}\sum_t R_t$$

and

$$B_n = \left(\frac{1}{n}\sum_t \text{sign}(r_{t|t-1})\right)\left(\frac{1}{n}\sum_t r_t\right).$$

A_n estimates the average return from using the trading strategy whereas B_n estimates the average return from using the benchmark strategy that issues buy/sell signals randomly with probabilities corresponding to the proportion of buys and sells implied ex post by the trading strategy. When r_t is predictable investing in the trading strategy will generate higher returns than the benchmark and the difference between A_n and B_n will be sizable. The variance of the difference $A_n - B_n$ is

$$V = Var(A_n - B_n) = \frac{4(n-1)}{n^2} p_r(1-p_r) Var(r_t),$$

where $p_r = Pr\{\text{sign}(r_{t|t-1}) = 1\}$. The estimator for the variance is $\hat{V} = 4/n^2 \hat{p}_r(1-\hat{p}_r)\sum_t (r_t - r_{t|t-1})^2$ with $\hat{p}_r = 1/2(1 + 1/n\sum_t \text{sign}(r_{t|t-1}))$. The excess profitability statistic is then given by

$$EP = \frac{A_n - B_n}{\sqrt{\hat{V}}} \overset{d}{\to} N(0,1)$$

under the null hypothesis.

References

Anatolyev, S., Gerko, A., 2005. A trading approach to testing for predictability. Journal of Business & Economic Statistics 23 (4), 455–462.
Bewley, T., 2002. Knightian decision theory. Part I. Decisions in Economics and Finance 25, 79–100.
Blume, L., Easley, D., O'Hara, M., 1994. Market statistics and technical analysis: the role of volume. Journal of Finance 49, 153–181.
Boswijk, H., Hommes, C., Manzan, S., 2007. Behavioral heterogeneity in stock prices. Journal of Economic Dynamics & Control 31 (6), 1938–1970.
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.
Chen, Z., Epstein, L., 2002. Ambiquity, risk and asset returns in continuous time. Econometrica 70, 1403–1443.
Cheung, Y., Chinn, M., Marsh, I., 2004. How do uk-based foreign exchange dealers think their market operates? International Journal of Finance & Economics 9 (4), 289–306.
Chiarella, C., Dieci, R., Gardini, L., 2006. Asset price and wealth dynamics in a financial market with heterogeneous agents. Journal of Economic Dynamics & Control 30, 1755–1786.
De Grauwe, P., Grimaldi, M., 2006. The Exchange Rate in a Behavioral Finance Framework. Princeton University Press, Princeton.
Della Corte, P., Sarno, L., Tsiakas, I., 2009. An economic evaluation of empirical exchange rate models. Review of Financial Studies, forthcoming, doi:10.1093/rfs/hhn058.
Dow, J., Werlang, S., 1992. Uncertainty aversion, risk aversion, and the optimal choice of portfolio. Econometrica 60 (1), 197–204.
Ellsberg, D., 1961. Risk, ambiguity and Savage axioms. Quarterly Journal of Economics 75, 643–669.

Epstein, L., Wang, T., 1994. Intertemporal asset pricing under Knightian uncertainty. Econometrica 62 (2), 283–322.
Garlappi, L., Uppal, R., Wang, T., 2007. Portfolio selection with parameter and model uncertainty: a multi-prior approach. Review of Financial Studies 20 (1), 41–81.
Gilboa, I., Schmeidler, D., 1989. Maxmin expected utility with non-unique prior. Journal of Mathematical Economics 18, 141–153.
Hansen, L., Sargent, T., 2007. Robustness. Princeton University Press, Princeton.
Hellwig, M., 1982. Rational expectations equilibrium with conditioning on past prices: a mean-variance example. Journal of Economic Theory 26, 279–312.
Hendry, D., Lu, M., Mizon, G., 2002. Model identification and non-unique structure. Nuffield Economics Research via WWW 2002-W10.
Julier, S., Uhlmann, J., 2004. Unscented filtering and nonlinear transformation. IEEE Review 92 (3).
Keynes, J., 1921. A Treatise on Probability. McMillan, London.
Kirman, A., 1992. Whom of what does representative agent represent? Journal of Economic Perspectives 6, 117–136.
Kirman, A., 1993. Ants, rationality and recruitment. Quarterly Journal of Economics 108, 137–156.
Kirman, A., 1995. The behaviour of the foreign exchange market. Bank of England Quarterly Bulletin 35, 286–293.
Kirman, A., Ricciotti, R.F., Topol, R.L., 2007. Bubbles in foreign exchange markets: it takes two to tango. Macroeconomics Dynamics 11 (1), 102–123.
Knight, F., 1921. Risk, Uncertainty and Profit. Houghton Mifflin, Boston.
Kozhan, R., Salmon, M., 2007. On uncertainty, market timing and the predictability of tick by tick exchange rates. FERC Working Paper.
Maccheroni, F., Marinacci, M., Rustichini, A., 2006. Ambiguity aversion, robustness, and the variational representation of preferences. Econometrica 74 (6), 1447–1498.
Maccheroni, F., Marinacci, M., Rustichini, A., Taboga, M., 2008. Portfolio selection with monotone mean-variance preferences. Economic Working Paper Series 664, Bank of Italy, Economic Research Department.
Maenhout, P., 2004. Robust portfolio rule and asset pricing. The Review of Financial Studies 17 (4), 951–983.
Mangelsdorf, L., Weber, M., 1994. Testing Choquet expected utility. Journal of Economic Behavior and Organization 25, 437–457.
Meese, R., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? Journal of International Economics 14, 3–24.
Menkhoff, L., Taylor, M., 2007. The obstinate passion of foreign exchange professionals: Technical analysis. Journal of Economic Literature 45 (4), 936–972.
Pesaran, H., Timmermann, A., 1992. A simple nonparametric test of predictive performance. Journal of Business & Economic Statistics 10 (4), 461–465.
Pesaran, H., Timmermann, A., 1994. A generalization of the non-parametric Henriksson–Merton test of market timing. Economic Letters 44, 1–7.
Preston, A., 1978. Concepts of structure and model identifiability for econometric systems. In: Stability and Inflation. Wiley, New York (Chapter 16).
Quiggin, J., 1982. A theory of anticipated utility. Journal of Economic Behavior and Organization 3, 323–343.
Rothenberg, T., 1971. Identification in parametric models. Econometrica 39, 577–592.
Schmeidler, D., 1989. Subjective probability and expected utility without additivity. Econometrica 57, 571–587.
Taylor, M., Allen, H., 1992. The use of technical analysis in the foreign exchange market. Journal of International Money and Finance 11 (3), 304–314.
Uppal, R., Wang, T., 2003. Model misspecification and underdiversification. Journal of Finance 58, 2465–2486.
Van der Merwe, R., 1998. Sigma-point Kalman filters for probabilistic inference in dynamic state-space models, Ph.D. Thesis, University of Stellenbosch.
Wakker, P., 2001. Testing and characterizing properties of nonadditive measures through violations of the sure-thing principle. Econometrica 69 (4), 1039–1059.
White, H., Domowitz, I., 1984. Nonlinear regression with dependant observations. Econometrica 52 (1), 143–161.
Zhao, Y., Haussmann, U., Ziemba, W., 2003. A dynamic investment model with control on the portfolio's worst case outcome. Mathematical Finance 13 (4), 481–501.