Speculation as Explanation

Modeling cacao prices by means of heterogeneous agents' behavior

Abstract

In this paper a Heterogeneous Agent Model for the cacao market price is extended and analyzed. It explicitly incorporates the recent financialization of commodity markets. It takes into account non-linear speculator behavior as an explanatory factor. The speculators are assumed to follow either a mean-reversion strategy or a trend following strategy and are allowed to switch strategies based on past performance. The fundamental price that guides the behavior of the speculator following a mean-reversion strategy is explicitly modeled with fundamental factors. The model is estimated on daily frequency data to provide detailed insight into price developments. The model estimates stay close to their theoretical expectations, though are not all significant. Forecast performance is similar to benchmark models amongst which a modified model that uses a moving average as fundamental price. The results further show a decrease of the fraction of speculators following a mean-reversion strategy after periods of overpricing. Based on these results, the conclusion is drawn that the cacao market price model evaluated here reveals significant dynamic speculator behavior. The behavior is difficult to reconcile with expectations raised by the literature. Further, it is concluded that HAM's can be modeled to lay closer to economic reality by using an explicitly modeled fundamental price. Further improvements along the lines directed by the estimation results, might lead to increased statistical performance as well as modeled price dynamics more in line with expectations.

Thesis by: Rogier Maarten Hanselaar

E-mail: rogier.hanselaar@gmail.com

Supervisor: Prof. Dr. Dick van Dijk

Co-reader: Dr. Michel van der Wel

Erasmus School of Economics

Erasmus University Rotterdam

August 13, 2013, Rotterdam

Contents

1. Introduction

In recent years, prices of commodities have become quite volatile. In 2001, the price of London Brent Crude oil was below 20 dollars per barrel, rose to 140 dollars per barrel in mid-2008 and dropped again to about 40 dollars per barrel near the end of 2008 (Ellen & Zwinkels, 2010). Agricultural commodity prices rose by nearly 40 percent in 2007 and declined almost as strongly again in 2008 (Rosegrant, 2008). In 2010 grain prices saw a sharp rise and in 2011 grain prices spiked again (Lagi, Bar-yam, Bertrand, & Bar-yam, 2011).

Volatile commodity prices impair proper market functioning. Traditionally, commodity market functions have been to facilitate the price discovery process, to allow the transfer of price risk, and to provide an optimal allocation of goods (Flassbeck, Bicchetti, Mayer, & Rietzler, 2011; Lagi et al., 2011). Increased price volatility implies that prices become unreliable. As a consequence, the resulting uncertainty discourages hedging by commodity users and producers, and makes hedging more expensive due to the increased number of margin calls (Flassbeck et al., 2011). Next, the mispricing implies that the equilibrium price is not attained, where equilibrium price refers to the price dictated by demand and supply stemming from the real economy that produces an optimal allocation of resources (Lagi et al., 2011).

While the sub-optimal allocation of goods implies an economic loss, it also raises ethical questions when the goods are food commodities. As a result of the sudden high food prices in 2007/2008, many low-income households in developing countries had trouble sustaining themselves. Food riots broke out, as food became too expensive for a large fraction of the population. At the macro-level, the high volatility of commodity prices led to much uncertainty, which discouraged investments and slowed long-term growth especially in developing countries (Flassbeck et al., 2011; Lagi et al., 2011). To be able to reduce price volatility by means of policies, it is necessary to analyze its causes by modeling price development.

In this paper, a commodity price model is advanced and analyzed with the aim of providing additional insight into the causes underlying price movements. The model is novel in that it incorporates factors that drive supply and demand stemming from the real economy and combines this with non-linear speculator behavior based on these fundamental factors. In this paper further new terrain is explored, as the model is estimated on relatively high frequency data.

Incorporated elements commodity price dynamics

Fundamental factors are included in the model under scrutiny, as commodity price movement has been attributed to certain supply and demand factors in line with the efficient market hypothesis, which states that actual prices reflect fundamental values determined by supply and demand (Barberis & Thaler, 2002). For instance, supply of food commodities is notoriously depend on climatic conditions, while supply of oil can be affected by political instability in the oil producing countries. On the demand side, rapidly expanding economies require more energy, contributing to an already high demand for fossil fuels. Next, these economies produce an emerging middle class that demands more meat, the production of which requires more maize or related food crops. Further, the execution of governmental biofuel programs aimed at producing bioethanol and biodiesel from food crops causes a higher demand for food crops as well (Flassbeck et al., 2011; Rosegrant, 2008).

While models that only incorporate these supply and demand factors most likely offer some explanation for the price evolution in commodity markets, they are restricted in their capabilities of simulating the bubbles and crashes present in historical commodity prices (He & Westerhoff, 2005). Hence, the commodity price model elaborated upon here encompasses more than fundamentals alone by including the behavior of speculators.

There is abundant motivation for financial investors, and with them speculators, to have been active in commodity markets. Commodities have been identified as an attractive asset class. They offer the same average return as equity, are negatively correlated with both equity and bonds and are less volatile than equity due to the low correlation between commodities. Next, commodities form a good hedge against inflation, as they are used as inputs in the production process and their prices thus partially determine consumer prices; moreover, they form an interesting hedge against dollar currency fluctuations, as many commodities are paid for in dollars (Gorton & Rouwenhorst, 2005). Finally, there are arguments stating that commodity prices are in a bullish period due to structurally increasing demand (Radetzki, 2006). The burst of the dot-com bubble, spurring investors to look for alternatives for equity, combined with deregulation of the commodity markets in the beginning of the 21st century, has increased consciousness of the benefits of investing in commodities. Accordingly, the amount of commodity-related assets under management by investors has risen sharply in recent years (Flassbeck et al., 2011).

Speculation is able to offer additional explanation if agents are not assumed to instantly exploit mispricing through arbitrage, contrary to the efficient market hypothesis. Compelling reasons to relax this assumption are the exposure to certain risks, the incurring of certain costs, and the direction of trading that arbitraging implies. Firstly, there is fundamental risk, which refers to possibilities that the fundamentals change while the arbitraging agent tries to exploit the price difference. Secondly there is noise trader risk, which refers to the possibility that noise traders push the price further away from its true value on the short-term, while the arbitrageur is trading in the other direction trying to exploit a mispricing. This can force the arbitrageur to liquidate his position, as his clients whose money he is investing might decide he is incompetent and withdraw their money, resulting in large losses. Next to these risks, an arbitrageur is hindered by implementation costs, such as commissions, bid-ask spreads, and short-sales constraints which reduce the profitability of the exploitation of mispricing. Lastly, if the noise traders are trend followers, the arbitraging agent might prefer to trade in the same direction as the noise traders – rather than against it – knowing that the price change attracts more trend followers pushing the price further up (Barberis & Thaler, 2002).

There have been multiple mechanisms through which financial investment and speculation in commodity markets could have increased price volatility and contributed to bubble forming. One such way is through Exchange Traded Products (ETPs), which mostly track a single commodity index. Traditionally ETPs used futures contracts as collateral. However, ETPs have recently increasingly been using physical commodities to reduce counterparty risk, which reduces commodity supply as much inventory is earmarked as collateral and is not available for delivery. This drives up spot prices and pushes the commodity price curve into backwardation, which makes index investments more attractive and causes more demand for ETPs, again strengthening this process and potentially igniting bubble formation (Flassbeck et al., 2011). Another way in which price volatility and bubble formation could have been increased is through herding behavior of speculators, in particular through herding caused by speculators believing they can extract information from other agents' behavior. As Banerjee (1992) points out, this leads to inefficient outcomes and very volatile equilibrium behaviour. Moreover, it implies that a position taken by an actor is often not based on his or her own information, but on other agents' past actions. Accordingly, price changes may occur without any new information entering the market, and prices may change greatly when new information does present itself (Flassbeck et al., 2011). Combined with the limits to arbitrage and the increase in financial investor activity in commodity markets mentioned above, these trend extrapolating mechanisms make it probable that financialisation has induced increased price volatility and bubble formation in commodity markets. As such, the model advanced here explicitly includes the effects of financialisation through speculator behavior.

Existing literature on modeling speculator behavior

Recent academic work has explored the possibility of including speculator behavior as explanatory factors in price models. Hommes (2006) provides an early overview of possible ways to implement such an addition, and calls them Heterogeneous Agent Models (HAMs). Central to each approach is the positioning of a stabilizing group of speculators (fundamentalists) against a destabilizing group of speculators (chartists). The former group believes there is some true price, the fundamental value, to which prices converge in the long run. When these speculators note a deviation from the fundamental value, they trade against this deviation pushing prices back towards the fundamental value. The latter group believes in price trends. When a shock pushes prices up, these speculators read a price trend in it. Consequently, they will buy to exploit this perceived trend, pushing prices further away from the fundamental value. The price trend becomes a self-fulfilling prophecy.

HAMs combined with non-linear interaction between both speculator types allows bubbles to form and bull and bear markets to be created irregularly, while in the long-run the prices mean-revert. HAMs can reproduce various stylized facts of asset returns, such as persistence in asset prices, unpredictability of returns on the short run and predictability on the long run, excess and clustered return volatility, and fat tails of the return distribution. Further they generate high and persistent trading volumes caused by differences in beliefs, and produce persistent deviations from fundamental values. While there is a plethora of different speculators (and just as many mechanisms through which they influence prices) aggregation over them can be expected to leave a stabilizing group and a destabilizing group central (Hommes, 2006; Lagi et al., 2011).

A few authors have applied these HAMs to commodity prices. Lagi et al. (2011) use HAMs to determine what factors have caused the price spikes in food commodities in recent years, and conclude that speculation and ethanol conversion are the dominant causes. Reitz & Westerhoff (2007) apply HAMs to a wide variety of commodities, and make the amount of fundamentalist traders dependent on the deviation of the actual price from the fundamental price; they model the fundamental price as the long run average of the actual prices. Westerhoff & Reitz (2005) model maize prices with a HAM in which the amount of technical traders depends on the distance of the actual price from the fundamental price, with a higher proportion for a larger distance. Again the fundamental price is modeled as the long run average of the actual prices. Both assume a latent pool of speculators ready to step in. Ellen & Zwinkels (2010) improve upon this framework by assuming a fixed amount of speculators that are allowed to choose either a fundamentalist or chartist strategy. The speculators base their choice on past performance of this strategy, and may switch often. The fundamental price is calculated as the 2 year moving average of actual prices.

Model set-up

In this paper, the Heterogeneous Agent framework is utilized to model speculator behavior. The model of Ellen & Zwinkels (2010) is adopted and extended. The two types of speculator strategies mentioned above are used: the fundamentalist strategy in which speculators form their price expectations on fundamental prices and the chartist strategy in which speculators form their expectations on perceived price trends. The first strategy causes market prices to revert to the fundamental prices, while the second strategy chases market prices away. Speculators are modeled to choose their strategy based on the past performance of those strategies. The fundamental price is determined by real supply and demand factors.

The model advanced here gives the heterogeneous agent literature a further impulse by explicitly linking fundamental prices to real supply and demand factors of commodities. In doing so, it models the practice of speculators trading on fundamentals in detail and stays close to market reality. This is a major change to previous research, in which the fundamental price is modeled with a simple moving average. Moreover, the model is estimated on daily frequency data. This expectedly provides more detailed insights into price developments. Compared to previous research in which monthly data is used, this is another novelty. To estimate this more progressive model, state space modeling is used to estimate all coefficients directly on historic data.

The cacao market price is chosen to estimate the model on. The cacao market is of particular interest for several reasons. Cacao is a commodity whose dynamics provide potential insight into the dynamics of food commodities in general. In the light of recent high food prices and food riots, conclusions drawn may carry additional relevance for the debate on the effects of speculation on food availability to large parts of the world. Next, cacao is a food commodity that has evaded scrutiny so far, hardly any quantitative modeling has been done. Finally, the real demand and supply of cacao depend on only a limited number of factors. Cacao is not used to feed livestock, nor is it used to produce biofuel. On the supply side, there is no OPEC-like cartel while the production of cacao is still very concentrated; almost 70% of world cacao production comes from the south of West Africa (International Cocoa Organisation, 2012).

The origins of the data used to estimate the parameters in the model, and specific characteristics of that data are elaborated upon in the following section. Subsequently, the details of this cacao return model are discussed. This is followed by an implementation section that covers the operation of the filter applied to the specific cacao return state space model. In the next section the results of this empirical investigation are presented and evaluated. They reveal significant fundamentalist speculator behavior. However, the price patterns do not seem to coincide with expectations raised by the literature. The model further has multiple non-significant parameters and average out-of-sample performance. Based on these results, the conclusion is drawn that the cacao market price model evaluated here reveals significant dynamic speculator behavior that is difficult to interpret in line with what is suggested by literature. Further, it is concluded that HAM's can be modeled to lay closer to economic reality by using an explicitly modeled fundamental price. Further improvements along the lines directed by the estimation results, might lead to increased statistical performance as well as modeled price dynamics that lie closer to expectations raised by the literature.

2. Data description

The dataset used to estimate the model consists of daily time series ranging from 07 May 1999 to 31 January 2013. The variable of key interest is the daily cacao price. The daily cacao prices are constructed by taking the average of the prices in dollars of the nearest three expiries (ignoring the nearby if the 15th of the month before the maturity month of the nearby has passed) of daily cacao futures acquired from the Intercontinental Exchange (ICE). Futures prices instead of spot prices are used as proxy, as the academic literature shows that agricultural futures prices often lead agricultural spot prices; see for instance Hernandez & Torero (2010) and Reichsfeld & Roache (2011).

The exogenous variables in the model are a variable governing cacao stock levels and a variable governing rainfall in the cacao producing area. Cacao stock levels in bags of 145 lbs are proxied by cacao stock levels in warehouses in the United States monitored by the ICE. Missing values are interpolated using the average of the preceding and consecutive values. Values of stock levels that correspond to non-trading days are removed. The data on stock levels in ICE warehouses is retrieved from Bloomberg (the COCAIVTL Index). Daily data on average precipitation in millimeters in the main cacao producing areas (5N to 7N and -7E to 0E, see OECD & SWAC (2011)) is acquired from the Climate Explorer tool of the Royal Netherlands Meteorological Institute (KNMI).

Both the cacao market price series and the cacao stock level series exhibit high autocorrelation. Unit root tests with trend and intercept reveal that both contain a unit root. The results of the tests can be found in appendix A. First differences of both are taken to obtain stationary series.

Stylized facts

The cacao returns display typical characteristics of asset returns. First, they are not normally distributed; as can be seen in figure 1a, the histogram of cacao returns does not coincide with the red line indicating the shape of a normal distribution. While the returns seem to be nicely centered around zero, the peakedness of the histogram suggests that excess kurtosis is present. Moreover, the shape of the histogram is mildly asymmetric which indicates that skewness is not zero.

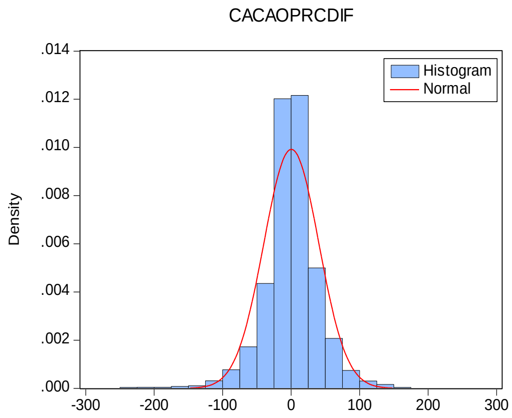

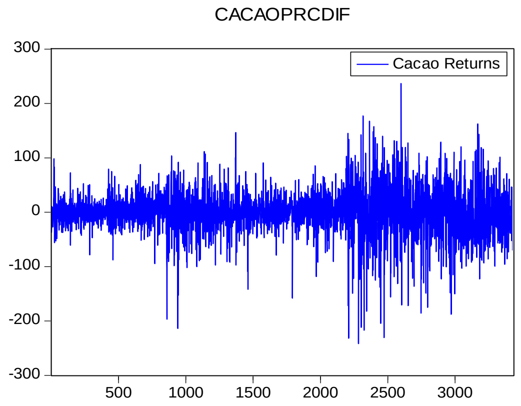

Figures 1a: Histogram of cacao returns; *figure 1b: cacao returns plotted over time*

The statistics in the table 1, show that the returns are indeed slightly skewed and exhibit excess Kurtosis. The J-B statistic is very high, which affirms that it is highly improbable that these returns are normally distributed.

Cacao Returns

Mean	0.324
Median	0.333
Std. Dev.	40.185
Skewness	-0.441
Kurtosis	7.274
Jarque-Bera	2725.145
Probability	0
Observations	3434

Table 1: Descriptive statistics cacao price return

Further, the right graph of figure 1 shows that there are periods with larger returns and periods with smaller returns following up on each other, which suggests that volatility clustering is present and that there is autocorrelation in the squared and absolute returns. The black and green lines in figure 2 below represent the empirical autocorrelation functions of the absolute cacao return and the squared cacao return. They confirm that there is significant correlation in the squared cacao returns and absolute cacao returns, as they lie outside the 95% confidence interval demarcated by the red lines. Figure 2 thus shows that volatility clustering is present.

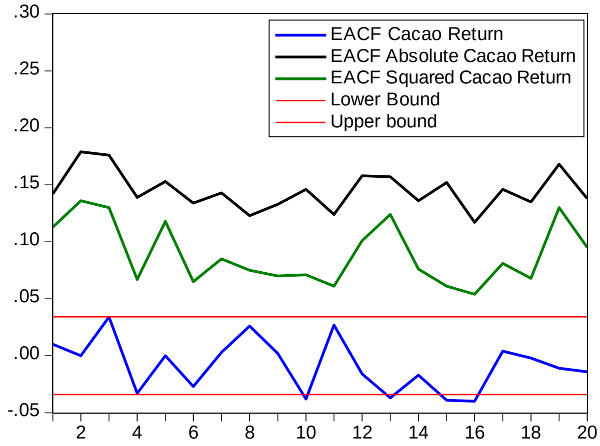

Figure 2: Plot of empirical autocorrelation functions of the normal, absolute and squared cacao returns

Finally, the right graph of figure 1 does not suggest that the cacao returns are correlated over time, as positive and negative returns do not seem to follow up on each other in a certain pattern. Figure 2 above confirms this. The blue line displays the empirical autocorrelation in cacao returns and remains well within the 95% confidence interval.

3. Model

A Heterogeneous Agent Model for the cacao returns is constructed that utilizes agent behavior to explain cacao returns, building upon the work of Ellen & Zwinkels (2010). It consists of a part that directly models the effect of fundamental factors on the cacao returns and a part that models the behavior of two types of speculators active on the cacao commodity markets. The first type of speculator is the fundamentalist speculator, who trades against market price deviations from the fundamental cacao price. The second type of speculator is the chartist speculator, who pursues price trends. In this section, first the fundamental factors affecting cacao prices are explored. Instead of using a simple moving average as is common in the literature, the fundamental cacao return is determined by the interplay of the fundamental factors. Subsequently the behavior of fundamentalist speculators is modeled, followed by that of chartist speculators.

Fundamental factors are linked to cacao prices via cacao supply and cacao demand. Cacao is almost exclusively used to produce chocolate; the demand for cacao therefore mainly stems from chocolate producers. Demand for chocolate can be expected to be negatively dependent on the price of cacao (Ruf & Siswoputranto, 1995). Therefore real cacao demand (D_t^R) is formalized as:

$$D_t^R = b_1^R P_t \qquad 1)$$

Here b_1^R is expected to be negative, as a high cacao price implies a lower demand for cacao.

Cacao is supplied by cacao farmers, whose cacao production depends on environmental influences as well as returns on producing and selling cacao. Environmental influences have a direct effect on production. Rainfall is the most important influence; periods of draught of 3 months or longer endanger cacao production, while abundant and regular rain increase cacao production. High returns tend only to have an effect on the long-term, as newly planted cacao trees become productive after about 4 years (International Cocoa Organisation, 2013). Similarly, low returns have a more short-term effect as farmers cut down cacao trees to make room for other crops. Here, returns are approximated by the cacao price. Cacao market prices are used as producer prices seem to fluctuate on average with market prices (International Cocoa Organisation, 2012) and costs are assumed constant. Further, cacao supply is also dependent on global cacao inventory levels, as high inventory levels imply a large supply of cacao. (International Cocoa Organisation, 2013). Real supply (S_t^R) is therefore formalized below as:

$$S_t^R = b_1^S P_t + b_2^S Weather_t + b_3^S I_t \qquad 2)$$

As a high cacao price makes it more attractive for producers to supply more cacao, b_1^S is expected to be positive. $Weather_t$ is a continuous variable that is constructed as the average precipitation over the past 90 days. As it reflects the effects of rainfall upon cacao production, b_2^S is expected to have a positive sign. I_t replicates the effect of cacao inventory levels on supply; as such b_3^S is expected to have a positive sign as well.

As mentioned at the beginning of this section, the interplay of the fundamental factors described above does not only influence market prices, it also produces the fundamental cacao return (ΔF_{t+1}). It is modeled as:

$$\begin{aligned} \Delta F_{t+1} &= c + \vartheta_2^F (D_t^R - S_t^R) + \eta_{t+1} \qquad 3) \\ &= c + b_1\ P_t + b_2\ Weather_t + b_3\ \Delta I_t\ + \eta_{t+1} \end{aligned}$$

The fundamental price (F_t) is seen by the fundamentalists as the long-run intrinsic worth of cacao, and it steers the behavior of fundamentalist speculators. To be more precise, the demand for cacao of fundamentalist speculators (D_t^F) is driven by the difference between the fundamental price of cacao (F_t) and the current price of cacao (P_t). They expect the fundamental price (F_t) to be the price at which the commodity will be traded in the long-run. Thus, when the market price (P_t) is much lower (higher) than the fundamental price (F_t), fundamentalists expect market prices to rise (fall) and want to buy (sell) cacao; subsequently the demand for cacao from fundamentalist speculators will be large (small). They therefore have a stabilizing effect on the cacao prices. This is formalized below:

$$D_t^F = a^F(E_t^F(P_{t+1}) - P_t) \quad 4)$$

$$where$$

$$E_t^F(P_{t+1}) = P_t + b_1^F(P_t - F_t)^+ + b_2^F(P_t - F_t)^- \quad (4.1)$$

$$F_t \text{ is the fundamental price at time } t$$

$$b_1^F \in [-1,0], \qquad b_2^F \in [-1,0]$$

$(P_t - F_t)^+$ takes on the value of the difference $(P_t - F_t)$ if that difference is larger than 0. Similarly, $(P_t - F_t)^-$ takes on the value of the difference $(P_t - F_t)$ is smaller than 0. If there currently is a high market price relative to the fundamental price (i.e. $(P_t - F_t)^+ > 0$), a fundamentalist speculator will expect a lower market price next period as he expects market prices to mean-revert to the fundamental price. Thus b_1^F is expected to be negative. Moreover, as both prices cannot become negative, the difference between the market price and the fundamental price can never be larger than the market price. Hence b_1^F is expected to be larger than *-1*. On a similar note, b_2^F is expected to lie in the interval $[-1,0]$ as well. Two parameters are used, as Kahneman & Tversky (1979) have shown that individuals value avoiding losses more than obtaining gains. As actual prices cannot drop below 0, fundamentalists possibly respond more strongly to a market price that is below the fundamental price (as potential losses are limited) than to an actual price that is above the fundamental price. Similarly, short-sale constraints might restrict the possibilities of fundamentalist speculators to drive back overpricing. Hence, b_1^F is expected to be smaller than b_2^F.

On the other hand, chartist speculators look for trends in price changes and try to exploit these. Further, they ignore exogenous price information. Here a specific simple and general type of chartist speculator is used: the AR(1)-chartists. Chartists of this type form their expectation on the next periods price change by looking at the price change in the previous period. In other words, the demand for cacao from chartist speculators is driven by the difference between the current price (P_t) and the expected price ($E_t^c(P_{t+1})$) that is obtained through extrapolating the price increase of the previous period. These chartists ignore all information on fundamentals and push prices away from their fundamental price: small price changes are exacerbated as chartists perceive them as trends and trade in the same direction. The chartist demand is formalized below:

$$D_t^c = a^c(E_t^c(P_{t+1}) - P_t), where: \quad 5)$$

$$E_t^c(P_{t+1}) = P_t + b_1^c(P_t - P_{t-1})^+ + b_2^c(P_t - P_{t-1})^- \quad (5.1)$$

If the price change of last period was positive (i.e. $(P_t - P_{t-1})^+ > 0$), chartists expect the actual price to be larger next period as they see price changes as price trends. Thus b_1^c is expected to be positive. On a similar note, b_2^c is expected to be positive as well. Again b_1^c and b_2^c are allowed to take different values to account for loss aversion and short-sale restrictions; b_1^c is expected to be larger than b_2^c.

Whether the speculators in the market follow a fundamentalist strategy or a chartist strategy depends on how they assess the performance of both strategies. In this research, the speculators are assumed to base their assessment on past performance of the strategies, as is often done in the HAM literature (see for instance De Jong, Verschoor, & Zwinkels (2009)). Past performance of the strategies is measured using the squared forecasting error over the previous K periods, as in Ellen & Zwinkels (2010). If a certain strategy has performed well over the past period, more speculators tend to follow that strategy. The fraction of all the speculators that opt for the fundamentalist strategy (W_t) is determined by a multinomial switching rule as proposed in Brock & Hommes (1997), but with the adjustments made by Ellen & Zwinkels (2010) to allow for easier estimation as well as easier comparison of the intensity of choice parameter γ across different markets and over time, if need be there. This fraction (W_t) is calculated as follows:

$$W_t = \left(1 + \exp\left(-\gamma\left(\frac{A_t^C - A_t^F}{A_t^F + A_t^C}\right)\right)\right)^{-1} \qquad 6)$$

$$where$$

$$A_t^F = -\sum_{k=1}^{K}(E_{t-k-1}^F(P_{t-k}) - P_{t-k})^2 \qquad (6.1)$$

$$A_t^C = -\sum_{k=1}^{K}(E_{t-k-1}^C(P_{t-k}) - P_{t-k})^2 \qquad (6.2)$$

The fraction of all speculators that adhere to the chartists strategy is then equal to $1 - W_t$. The intensity of choice parameter γ determines how much past performance affects the choice between strategies for next period. If $\gamma = \infty$, a slightly better performance of one particular strategy will already result in all speculators adopting that strategy next period; if $\gamma = 0$ past performance will not matter at all and the fraction of speculators that adopt a fundamentalist strategy will be equal to the fraction of speculators that adopt a chartist strategy, namely 0.5. The number of days K on which the performance of the strategies is assessed is determined by estimating the full model described in this section for different values of K (K = 1,2,…,5,10). The K for which the Akaike Information Criterion

(AIC) value is smallest is subsequently chosen, where the AIC is calculated as follows (Heij, de Boer, Franses, Kloek, & van Dijk, 2004):

$$AIC(p) = \log(s_p^2) + \frac{2p}{n} \qquad 7)$$

$$where\ p = 8\ is\ the\ number\ of\ incl.\ regr.,$$

$$n = 3435\ \ is\ the\ number\ of\ observations, and\ s_p^2\ is\ MLE\ of\ the\ error\ variance$$

The table below lists the AIC value for the different values of K. As the table shows, the AIC value is smallest for K = 2. Hence the implementation and results sections elaborate upon the model estimated with K = 2.

	K=1	K=2	K=3	K=4	K=5	K= 10
AIC Value	7.3792	7.3745	7.3778	7.3782	7.3791	7.3820

Table 2: AIC Values for different values of K

The total demand for cacao (D_t^m) consists of the weighted demand of the fundamentalist speculators, the weighted demand of the chartist speculators and the demand stemming from the real economy. This is formalized below:

$$D_t^m = D_t^R + W_t D_t^F + (1 - W_t) D_t^C\ , where: \qquad 8)$$

$$W_t = \left(1 + \exp\left(-\gamma\left(\frac{A_t^C - A_t^F}{A_t^F + A_t^C}\right)\right)\right)^{-1}\ , and$$

$$A_t^F = -\sum_{k=1}^{K}(E_{t-k-1}^F(P_{t-k}) - P_{t-k})^2,\ A_t^C = -\sum_{k=1}^{K}(E_{t-k-1}^C(P_{t-k}) - P_{t-k})^2$$

As mentioned above, the fundamental cacao return (ΔF_{t+1}) is determined by the interplay of non-speculative supply and demand. It is a function of excess real demand:

$$\Delta F_{t+1} = c + \vartheta_2^F(D_t^R - S_t^R) + \eta_{t+1} \qquad 9)$$

$$= c + b_1\ P_t + b_2\ Weather_t + b_3\ \Delta I_t\ + \eta_{t+1}$$

Finally, the price change is a function of excess total demand and an error term. By plugging in the equations above and some rewriting, we obtain the final model for cacao price changes that is estimated:

$$\Delta P_{t+1} = c + \vartheta(D_t^m - S_t) + \varepsilon_{t+1} = \vartheta(D_t^R + W_t D_t^F + (1 - W_t)D_t^C - S_t) + \varepsilon_{t+1}$$

$$= c + \vartheta\,(b_1^R P_t + W_t a^F (E_t^F(P_{t+1}) - P_t) + (1 - W_t)a^c (E_t^c(P_{t+1}) - P_t) - b_1^S Weather_t - b_2^S \Delta I_t) + \varepsilon_{t+1}$$

$$= c + b_1 P_t + b_2 Weather_t + b_3 \Delta I_t + W_t(\alpha_1 (P_t - F_t)^+ + \alpha_2 (P_t - F_t)^-) + (1 - W_t)(\beta_1 (P_t - P_{t-1})^+ + \beta_2 (P_t - P_{t-1})^-) + \varepsilon_{t+1} \qquad 10)$$

$$= f(\boldsymbol{Data}_t, F_t, F_{t-2}, \boldsymbol{\theta}) + \varepsilon_{t+1}, where\ \theta\ represents\ all\ the\ parameters\ to\ be\ estimated$$

4. Estimation

As can be deduced from the above description, the model is a state-space model where the fundamental price of cacao, which is unfortunately not observed, is the process state. Only the actual trading price influenced by speculators is observed. To estimate state-space models, there is one widely used tool: the Kalman filter (Orderud, 2005).

The Kalman filter is a recursive algorithm that is an optimal estimator for linear systems. It demands the assumption that the errors in both the state and process equations are normally-distributed. It subsequently makes use of certain properties of joint normally distributed variables, to express the state and process equations recursively. Next, the parameters can be estimated using Maximum Likelihood estimation (Hamilton, 1994).

The model under consideration here is non-linear, due to the specification of the fraction of fundamentalist speculators W_t in terms of the state variable $F_t$and the multiplication of W_t with the speculator demand function. This implies that the Kalman filter is not directly applicable and needs to be modified. One proposed modification that makes the Kalman filter applicable to non-linear models is called the extended Kalman filter (EKF). It requires the assumption that a non-linear model can adequately be approximated by a linear model using a Taylor series expansion and treats the linear approximation as the true model to which the Kalman filter can be applied (Hamilton, 1994).

The EKF can be difficult to implement, as it requires the calculation of the Jacobian matrix. Moreover, if implemented, it can produce an unstable filter, given that the functions are not approximately linear on the time scale of update intervals, as well as large errors as it provides only a “first order” approximation (Julier & Uhlmann, 1997; Wan & Merwe, 2000). Julier & Uhlmann (1997) have developed an alternative method called the unscented Kalman filter (UKF) to overcome these issues. Instead of linearizing the function, the UKF approximates the probability distribution by propagating a set of sigma points through the ‘true’ non-linear function. These sigma points are chosen such that they possess the same first and second moments as the distribution from which they come. This method typically decreases computational complexity and increases estimation accuracy. Moreover, it allows for non-Gaussian or non-additive noises (Orderud, 2005; Julier, 1999; Wan & Merwe, 2000). Because of its lower computational complexity and increased estimation accuracy, the UKF is applied here.

5. Implementation

The UKF is implemented following the two stage method described in Durbin & Koopman (2012). The parameters in the model are estimated iteratively using maximum likelihood estimation. First the UKF is run with a set of parameters to obtain residuals. Subsequently the parameter estimates are improved by changing them such that the sum of squared errors becomes smaller, which is equivalent to maximizing the log-likelihood. This is repeated until the parameter estimates have converged. The parameters are assumed to be given each time the filter is run.

In the first stage of the filter the expected fundamental return values ($E(\Delta\boldsymbol{F}_t|\Delta P_t)$**)** and covariance at time t ($Var(\Delta\boldsymbol{F}_t|\Delta P_t)$**)** are updated using the market return observation at time t+1 (ΔP_{t+1}):

$$\Delta\boldsymbol{F}_{t|t+1} = E(\boldsymbol{\Delta F}_t|\Delta P_{t+1}) = E(\Delta\boldsymbol{F}_t|\Delta P_t) + \boldsymbol{V}_{Fv,t+1}V_{vv,t+1}^{-1}v_{t+1} \qquad 11)$$

$$\boldsymbol{V}_{t|t+1} = Var(\Delta\boldsymbol{F}_t|\Delta P_{t+1}) = Var(\Delta\boldsymbol{F}_t|\Delta P_t) - \boldsymbol{V}_{Fv,t+1}V_{vv,t+1}^{-1}\boldsymbol{V}_{Fv,t+1}{}'$$

$$where$$

$$v_{t+1} = \Delta P_{t+1} - \overline{\Delta P_{t+1}} \qquad (11.1)$$

$$\overline{\Delta P_{t+1}} = \sum_{i=0}^{2m} w_i f(\boldsymbol{Data}_t, \boldsymbol{sp}_{t,i}, \boldsymbol{\theta}) \qquad (11.2)$$

$$\boldsymbol{V}_{Fv,t+1} = \sum_{i=0}^{2m} w_i(\boldsymbol{sp}_{t,i} - \Delta\boldsymbol{F}_{t|t})\left(f(\boldsymbol{Data}_t, \boldsymbol{sp}_{t,i}, \boldsymbol{\theta}) - \overline{\Delta P_{t+1}}\right) \qquad (11.3)$$

$$V_{vv,t+1} = \sum_{i=0}^{2m} w_i\left(f(\boldsymbol{Data}_t, \boldsymbol{sp}_{t,i}, \boldsymbol{\theta}) - \overline{\Delta P_{t+1}}\right)^2 + H, \qquad (11.4)$$

$$where\ H = Var(\varepsilon) = \sigma_\varepsilon^2\ and\ m = 3$$

Recall that in the model under consideration ΔP_{t+1} depends on the fundamental return at time t (ΔF_t) via $(P_t - F_t)^+$ and $(P_t - F_t)^+$, as well as on the fundamental return at time t-2 (ΔF_{t-2}) and t-3 (ΔF_{t-3}) via W_t. Put differently, the ΔP_{t+1} contains information on ΔF_t, ΔF_{t-2}, and ΔF_{t-3}. Hence ΔP_{t+1} is used to update ΔF_t, ΔF_{t-2} as well as ΔF_{t-3}. The bold symbols for the fundamental return at time t thus refer to the vectors:

$$E(\Delta\boldsymbol{F}_t|\Delta P_t) = \Delta\boldsymbol{F}_{t|t} = \begin{bmatrix} \Delta F_t \\ \Delta F_{t-2|t-1} \\ \Delta F_{t-3|t} \end{bmatrix} \qquad 12)$$

$$E(\Delta\boldsymbol{F}_t|\Delta P_{t+1}) = \Delta\boldsymbol{F}_{t|t+1} = \begin{bmatrix} \Delta F_{t|t+1} \\ \Delta F_{t-2|t+1} \\ \Delta F_{t-3|t+1} \end{bmatrix}$$

In the vector $\boldsymbol{V}_{Fv,t+1}$, and the scalars $V_{vv,t+1}$ and ΔP_{t+1}, $\boldsymbol{sp}_{t,i}$ is placed instead of $\Delta\boldsymbol{F}_{t|t}$. It is a vector of sigmapoints, which are points centered around $\Delta F_{t|t}$, created using the following equations:

$$\boldsymbol{sp}_{t,0} = \Delta\boldsymbol{F}_{t|t} \qquad 13)$$

$$\boldsymbol{sp}_{t,i} = \Delta\boldsymbol{F}_{t|t} + \sqrt{m+k}\sqrt{\boldsymbol{V}_{t|t}(i,i)}^{*}$$

$$\boldsymbol{sp}_{t,i+m} = \Delta\boldsymbol{F}_{t|t} - \sqrt{m+k}\sqrt{\boldsymbol{V}_{t|t}(i,i)}^{*}$$

$$where\ i = 1{:}m\ and\ m = 3$$

$$and\ \sqrt{\boldsymbol{V}_{t|t}(i,i)}^{*}\ is\ a\ mx1\ matrix\ with\ \sqrt{\boldsymbol{V}_{t|t}(i,i)}\ at\ place\ i\ and\ 0\ at\ the\ other\ places$$

These sigmapoints carry weights w_i which reflect their respective probabilities. They are calculated such that they approximately reflect probabilities from a normal distribution. They are calculated as follows:

$$w_0 = \frac{k}{m+k} \qquad 14)$$

$$w_{1:2m} = \frac{1}{2(m+k)}$$

Together, the sigmapoints and weights approximate a probability distribution, in this case the normal distribution. The sigmapoints of the fundamental returns are propagated through the non-linear function $f\big(\boldsymbol{Data}_t, \boldsymbol{sp}_{t,i}, \boldsymbol{\theta}\big)$ to produce different realizations of the required statistics, which are combined by using the weights that approximate their respective probabilities.

In the second stage of the filter, the next values of the vector of fundamental market returns $\Delta\boldsymbol{F}_{t+1|t+1}$ and of the covariance matrix of fundamental market returns $\boldsymbol{V}_{t+1|t+1}$ are predicted. The fundamental market returns are assumed not to be correlated over time, as they should be driven by fundamental factors alone. Further, at time t+1, ΔF_{t-1} has already been updated with the information available at time t. Hence the predicted vector of fundamental market returns can be written as:

$$\Delta \boldsymbol{F}_{t+1|t+1} = \begin{bmatrix} c + b_1^F P_t + b_2^F Weather_t + b_3^F \Delta I_t \\ \Delta F_{t-1|t} \\ \Delta F_{t-2|t+1} \end{bmatrix} \qquad 15)$$

The error terms η_t of the fundamental market equation are assumed to be homoscedastic and to be uncorrelated over time. Under these assumptions the predicted covariance matrix $\boldsymbol{V}_{t+1|t+1}$ is thus constant over time and the off-diagonal elements are equal to zero:

$$\boldsymbol{V}_{t+1|t+1} = \begin{bmatrix} \sigma_{\eta_{t+1}} & 0 & 0 \\ 0 & \sigma_{\eta_{t-1|t}} & 0 \\ 0 & 0 & \sigma_{\eta_{t-2|t+1}} \end{bmatrix} \qquad 16)$$

$$where\ \eta_{t+1}\ refers\ to\ the\ error\ term\ at\ time\ t\ before\ it\ is\ updated,$$
$$\eta_{t-1|t}\ refers\ to\ the\ error\ term\ at\ time\ t-1\ after\ it\ is\ updated\ once$$
$$and\ \eta_{t-2|t+1}\ refers\ to\ the\ error\ term\ at\ time\ t-2\ after\ it\ is\ updated\ twice$$

In the light of this model and implementation, fundamental price smoothing is unnecessary. In smoothing, the updated fundamental returns are adjusted with information on the full sample to obtain the best fundamental return estimates. The updating makes sense if the full sample contains additional information on the fundamental returns; in particular this would be the case if the fundamental return model were containing an autoregressive component and the fundamental returns were thus being linked together. However, in the fundamental return model the fundamental return is a function of exogenous variables and is only linked to three market returns, namely the market returns at time t+1, t+3 and t+4. The full sample does not contain additional information over the two market returns; hence smoothing the updated fundamental returns does not improve them and is unnecessary.

The model is evaluated in-sample by calculating t-statistics of the parameters in the market equation to get an impression of their significance, as well as by computing the coefficient of determination. The model is further evaluated out-of-sample by assessing the bias, accuracy and efficiency of one-day-ahead forecasts and comparing them to those statistics obtained from a random walk model, an AR(1) model and a HAM model with a moving average of the cacao return with a window of 500 observations as fundamental return. Finally, some implicated return/fundamental return dynamics, as well as speculator proportions are evaluated.

The forecasts are made by dividing the dataset in an in-sample section consisting of the first 2595 observations, or equivalently the sample from 06-05-1999 until 30-09-2009, and an out-of-sample section consisting of the remaining 840 observations, or equivalently the sample from 30-09-2009 until 31-01-2013. First, the parameters of all four models are estimated in-sample. Subsequently, for the state space model the UKF is applied and the value of $\overline{\Delta P_{T+1|T}}$ is taken as the one-day-ahead forecast, where T refers to the last in-sample day. For the next forecast, the in-sample section is augmented with day $T+1$ after which the UKF is rerun with the same parameters to produce $\overline{\Delta P_{T+2|T+1}}$. For the other models, one-day-ahead forecasts are similarly produced, with iteratively making a forecast after which the in-sample section is augmented one day at the time. After 30 forecasts, all parameters are re-estimated on the in-sample section that is now 30 observations larger. This is continued until 840 1-step-ahead forecasts are made.

The bias of the forecasts is evaluated by calculating whether $E\left[e_{t+1|t} = 0\right]$. The accuracy of the forecasts is scrutinized by calculating the mean squared prediction errors (MSPE's) of the forecasts, that is: $E\left[e_{t+1|t}^2\right] = E\left[\left(\Delta P_{T+1} - \overline{\Delta P_{T+1|T}}\right)^2\right]$. The MSPE of the model under scrutiny is compared to the MSPE's of the other model by means of Diebold-Mariano-tests. This test examines whether the MSPE's differ significantly from one another. First, the loss differential $d_{t+1} = e_{i,t+1|t}^2 - e_{j,t+1|t}^2$ is calculated. Under the null of equal forecast accuracy, d_{t+1} is expected to be zero. The DM statistics is then calculated as: $DM = \frac{\text{mean}(d)}{\sqrt{\text{var}(d)/P}} \sim N(0,1)$, where P is the number of forecasts and equals 840. The efficiency of the forecasts is evaluated by means of Mincer-Zarnowitz regressions. Efficiency implies that the forecast errors cannot be explained with any information available at time t. The Mincer-Zarnowitz regression takes the form: $\Delta P_{T+1} = \alpha + \beta\overline{\Delta P_{T+1|T}} + \varepsilon_{t+1}$. Efficiency is tested by means of an F-test with restrictions $\alpha = 0$ and $\beta = 0$.

6. Results

The parameter estimation on the full sample yields the following coefficients (coefficients followed by * are significant at the 5% significance level, coefficients followed by ** are significant at the 1% significance level):

Parameter Values

Parameters	c	b_1	b_2	b_3	α_1	α_2	β_1	β_2
Values	-2.66144	-0.00035	0.01184*	-0.00001	-0.26096**	-0.00120	0.04419	0.02012

Table 3: Parameter estimates

Parameter Values

Parameters	σ_{η_t}	$\sigma_{\eta_{t\|t+1}}$	$\sigma_{\eta_{t\|t+3}}$	σ_{ε_t}	γ
Values	305.0	558.2	39.2	1,144.1	4.4735

Table 3 (continued): Parameter estimates

As cacao price can have a negative effect on the cacao return via supply and a positive effect on the cacao return via demand in our model, b_1 could have taken on both a positive and a negative value. The negative sign indicates that the negative effect on cacao returns via supply seems to slightly outweigh the positive effect via demand. The positive sign of b_2 is unexpected as the theorized positive effect of rainfall on supply implies a negative effect on cacao returns. Possibly, too much rainfall negatively affects cacao production. The negative sign of b_3 fits with the expectation that more stock implies more supply and thus has a negative effect on cacao returns. Further, both of the α's have a negative sign, which reflects the theorized effect of fundamental speculators pushing down too high market prices and pushing up too low ones, back to fundamental values. Next, the β's have positive signs, which corresponds with the theorized effects of chartist speculators pursuing price trends.

The relative size of α_1 with respect to α_2 is unexpected. Due to loss-aversion and short-sale restrictions α_1 was expected to be larger than α_2, while here it is found that reverse holds. On the other hand, the sizes of β_1 and β_2 are in line with expectations, as β_1 is larger than β_2. The value of b_1 governing the effect of cacao market prices on cacao market returns, shows that per dollar price increase the return is decreased by one-hundredth cent. Next, the parameter b_2 displays that the cacao return increases with about one-sixth cent per millimeter extra average precipitation in the cacao producing areas. Finally, the parameter b_3 indicates that the cacao return decreases with about one-thousandth cent per bag of cacao inventory.

The intensity of choice parameter γ has a relatively high value if compared to the value it attains in the work of Ellen & Zwinkels (2010), where it attains values of 1.1940 and 1.3566. This implies that a slightly better performance of a certain strategy causes a larger shift in the proportion of fundamentalist and chartist speculators. The relation between strategy performance and the fraction of fundamentalist speculators for $\gamma = 4.4735$ is plotted below in figure 3:

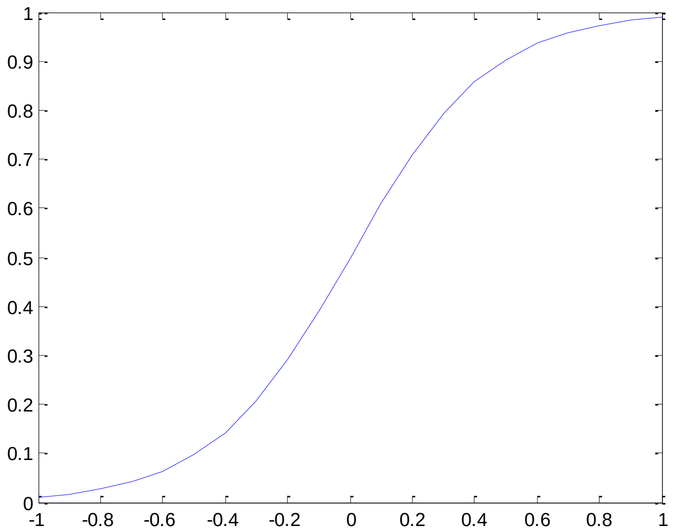

Figure 3: Plot of the fraction of fundamentalist speculators Wt (Y-axis) versus relative performance $\left(\frac{A_t^C - A_t^F}{A_t^F + A_t^C}\right)$ *(X-axis)*

A difference with Ellen & Zwinkels (2010) that might explain this, is the utilization of daily frequency data instead of monthly frequency data. Speculators that evaluate their strategies monthly might be less prone to switch. Another explanation might also be that speculator behavior differs across markets.

The values of the variance parameters of the error of the fundamental equation show an increased variance after the first update and a lower variance afterwards. This suggests that variance is added after the first update. As a consequence of the small disturbances after the final update, the parameters in the fundamental return equation are all highly significant. Note that the fundamental return ΔF_t in the state space model does not depend on a lagged value but is only determined by exogenous variables. As an alternative to estimating by filtering, the model can also be estimated in a different way. See the appendix B for details.

To be able to perform correct t-tests on the parameters in the market equation, the assumptions of homoscedasticity and no autocorrelation of the residuals of the market return equation are scrutinized, for details see appendix C. The residuals show to be heteroscedastic but to possess no autocorrelation.

As there is heteroscedasticity present in the residuals of the market return equation, White standard errors are used to compute t-statistics for the betas. The White estimate of the covariance matrix of the betas is calculated as in Heij et al. (2004):

$$\widehat{var}(b) = (X'X)^{-1}\left(\sum_{i=1}^{n} e_i^2 x_i' x_i\right)(X'X)^{-1} \quad 17)$$

$$where\ X\ is\ the\ matrix\ of\ regressors, e_i^2\ is\ the\ squared\ residual\ and$$
$$x_i\ is\ a\ row\ vector\ with\ observation\ i\ of\ X.$$

For the covariance of the betas of the market return equation, the matrix of regressors X is constructed with the twice updated fundamental values. The square roots of the elements on the diagonal of the two covariance matrices are used to calculate T-statistics for the betas:

$$t = b/SE \quad 18)$$

The results are displayed in table 4 below (t-statistics with * are significant at the 5% significance level, t-statistics with ** are significant at the 1% significance level):

T-statistics Betas Cacao Return Equation

Parameter	c	b_1	b_2	b_3	α_1	α_2	β_1	β_2
T-statistic	-0.85001	-0.30569	2.28754*	-0.93437	-3.34155**	-0.30902	0.57870	0.24580

Table 4: t-statistics of parameter estimates

Inspection of the market equation results reveals that of the parameters governing fundamentalist speculator behavior α_1 is significantly different from 0 at the 1% significance level. This implies that there is significant speculation that drives cacao prices from overpricing back to fundamental values; the latter of which are described by the exogenous variables cacao price, rainfall and inventory levels and the parameters c, b_1, b_2, b_3. Looking back at the market equation, it is interesting to see that of these variables nonetheless only rainfall seems to significantly (at the 5% significance level) influence the market price directly. The betas governing chartist speculation β_1 and β_2 are both not significant, implying that price deviations from the fundamental values are not due to trend extrapolating AR(1) chartists.

The coefficient of determination is calculated for the market return equation in the following manner:

$$R^2 = 1 - \frac{SSR}{SST} \quad 19)$$

$$where\ SSR\ is\ the\ sum\ of\ squared\ residuals\ and\ SST\ is\ the\ sum\ of\ squared\ returns$$

This yields the following results:

	R^2
Market Equation	0.0146

Table 5: Coefficient of determination of the market equation

These results show that about 1.5% of all the dynamics in the daily market returns is explained by this model. This is may seem a low percentage; nonetheless if financialization of the commodity markets has made the cacao market (partly) adopt stock market characteristics, then this level already exceeds the level of explained dynamics that may reasonably be expected (Rapach & Zhou, 2012). This combined with the non-significance of certain parameters mentioned above, suggests that the model could perform better in a more trimmed down version.

The forecasts (P=840) made to compare the different models result in the following statistics:

	Unbiasedness		**Accuracy**	**Efficiency**		
	Mean Forecast Error		**MSPE**	**Mincer-Zarnowitz regression**		
	$E[e_{t+1\|t} = 0]$	*p*-value	$E[e^2_{t+h\|t}]$	α	β	*p*-value Joint F-test
HAM-model	0.0241	0.9807	2449.6	-2.0780	-0.4400	0.0000
RW-model	-0,0296	0.9763	2427.0	-1.0567	-	-
AR1-model	-0,0352	0.9719	2429.5	8.4670	-15.7674	0.0000
HAM-MA-model	-0.0103	0.9918	2438.9	-3.7205	-4.2860	0.0000

Table 6: Out-of-sample results

The forecast statistics show that all models are unbiased. Moreover, all models seem to possess about the same amounts of accuracy. The results of the formal test of equal accuracy are displayed in table 7 below. They show that there is no significant difference in accuracy between the state space model and other models:

DM-statistics

	RW	AR(1)	HAM-MA
HAM-model	0.4402	0.3873	0.2182

Table 7: Diebold-Mariano statistics

Further, the parameter estimates in the Mincer-Zarnowitz regressions all seem to hint at inefficiency. Note that there is only one parameter estimate in the RW equation as including the RW forecast in the regression would imply including a second constant, causing problems for the matrix inversion required to calculate the OLS parameters. As a result, efficiency cannot be determined for the RW-model using a Mincer-Zarnowitz regression. For the other models very low P-values are found. Hence efficiency is rejected for the forecasts of those models.

In figure 4a below the cacao fundamental price is plotted, together with the cacao price constructed from the values of $\Delta P_{t+1|t}$ with starting point the cacao market price at day 1, and the fundamental price constructed from thrice updated values of fundamental returns ($\Delta F_{t|t+4}$) with starting point the cacao market price at day 1. In figure 4b below, the difference between the cacao return and the fitted cacao return over time is plotted.

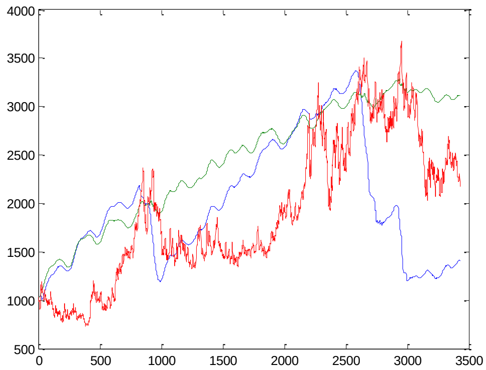

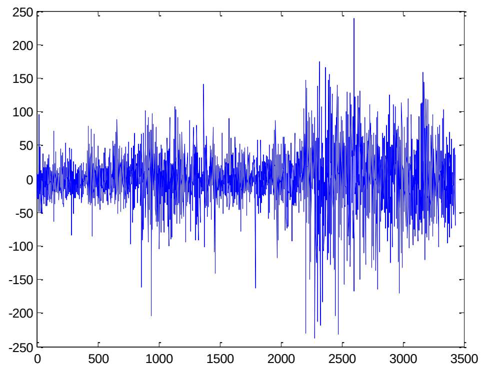

Figure 4a: Plot of the cacao price (red), the fitted price (blue) and the fundamental price (green) (Y-Axis) versus time (X-Axis)

Figure 4b: Plot of the difference between the cacao return and the fitted cacao return (Y-Axis) versus time (X-Axis)

The fitted cacao price constructed from the values of $\Delta P_{t+1|t}$ with starting point the cacao market price at day 1 seems to deviate quite substantially from the observed cacao price. This can be attributed to the fact that only about 1.5% of the observed cacao returns is explained by the model, and that the fitted cacao price is initialized once at day 1. In figure 4b, it is shown that the difference between the cacao return and the fitted cacao return seems to be centered around 0. This implies that the fitted return does not seem to structurally deviate from the market return.

Looking at figure 4a, it is noteworthy that the fundamental price of cacao seems to be very high. It suggests that there have only been four very short periods of overpricing and that in general cacao market (and fitted) prices lie much below fundamental levels. Specifically, according to the graph the periods of overpricing take place in July 2002, May 2008, March 2010, and January 2011. They are immediately followed by a strong dip in the fitted cacao prices, due to the very high and significant coefficient α_1 governing fundamentalist behavior. This is contradictory to the literature referred to in the introduction that suggests multiple extended price bubbles taking place on the cacao market.

In figure 5 below, the fraction of fundamentalist speculators active on the cacao market over time is depicted. The fraction plotted here is a moving average of 90 days. As can be seen from the graph, in general there are approximately the same amounts of chartist and fundamentalist speculators active. During periods in which prices rise above fundamental levels, the fraction of speculators that adheres to a fundamentalist strategy seems to drop meaningfully. This finding seems to reinforce the belief that periods of overpricing can be attributed to speculators engaging in herd behavior and following price trends instead of fundamentals. However, a closer inspection reveals that the fraction of fundamentalist speculators drops together with the fitted price after overpricing. This suggests that price bubbles are not caused by a dropping fraction of fundamentalist speculators. Rather this proposes that periods of overpricing are followed by a larger influence of chartist speculators on the cacao market.

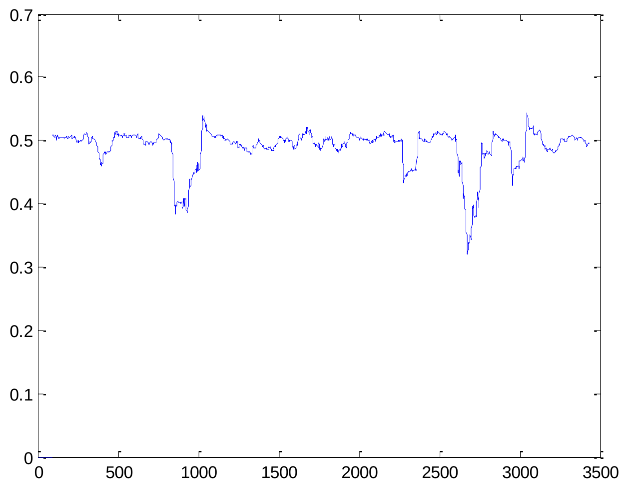

Figure 5: Plot of the fraction of fundamentalist speculators Wt (Y-axis) over time (X-axis)

7. Conclusion

In this paper a Heterogeneous Agent Model for the cacao market is developed modeling the behavior of two types of speculators to explain cacao returns over fundamental returns. The model explicitly links fundamental returns to fundamental factors to better model speculator behavior and the cacao market returns. Moreover, it is estimated on daily data to obtain detailed insights into price dynamics.

The model estimates stay close to theoretical expectations. The results show that the signs of the estimated parameters are all in line with expectations except for one fundamental factor, precipitation. The relative sizes of the speculator parameters are as expected for the parameters governing the behavior of the chartist speculators, but are reversed for the parameters governing the behavior of fundamentalist speculators. The intensity of choice parameter governing switching behavior shows that dynamical switching takes place based on past performance.

However, the price dynamics are hard to interpret. The fraction of fundamentalist speculators is on average equal to the fraction of chartist speculators. Periods of overpricing are followed by, rather than coincide with, a period in which the fraction of fundamentalist speculators is relatively low. This suggests that periods of overpricing are not caused by too few fundamentalist speculators keeping prices around fundamental levels and to many chartist speculators pushing prices away from fundamental levels; the results rather suggests the causation works the other way around. Further, the periods of overpricing that occur according to the model are much too short compared to the periods of overpricing mentioned by the literature.

Statistically, the model does not outperform. The parameter that models the effect of fundamentalist speculation that goes against overpricing is shown to be significant, as well as the parameter governing the effect of rainfall on the cacao market price. The model as a whole has a coefficient of determination of 1.46%. Generally, out-of-sample performance is not significantly different from a random walk model, an AR(1) model, or a HAM model with moving average.

These results thus show that it is difficult to apply the heterogeneous agent framework to higher frequency data than monthly frequency data and produce bubble dynamics that are consistent with what is suggested by literature. Nonetheless, the results do show that an explicitly modeled fundamental price can be incorporated into HAM´s which in turn can be estimated using a non-linear filtering method such as the Unscented Kalman Filter. This shows that HAM's can be specified to lay closer to economic reality than in previous research and produce results that thus offer a more detailed insight into cacao price developments.

That the model in its current form can be further improved to better capture the return dynamics of cacao, is suggested by in-sample non-significance of parameters, relatively high R-squared and average out-of-sample performance. The results indicate that fundamental factors do not necessarily need to be included directly into the cacao return model. Further, the non-significance of speculator behavior parameters suggest that the precise form of the equations governing fundamentalist and chartist behavior need further improvement; additional research might find functional forms that more adequately model speculator behavior. This might also change the unexpectedly late drop in the fraction of fundamentalist speculators following overpricing, as well as the high fundamental price.

Summing up, the cacao market price model evaluated here reveals significant dynamic speculator behavior that is difficult to interpret in line with what is suggested by literature. Further, it shows that HAM's can be modeled to lay closer to economic reality by using an explicitly modeled fundamental price. Further improvements along the lines directed by the estimation results, might lead to increased statistical performance as well as modeled price dynamics that lie closer to expectations raised by the literature.

References

Banerjee, A. V. (1992). A Simple Model of Herd Behavior. *The Quarterly Journal of Economics*, *107*(3), 797–817.

Barberis, N., & Thaler, R. (2002). A Survey of Behavioral Finance.

Brock, W. A., & Hommes, C. H. (1997). A Rational Route to Randomness. *The Econometric Society*, *65*(5), 1059–1095. Retrieved from http://www.jstor.org/stable/2171879 .

De Jong, E., Verschoor, W. F. C., & Zwinkels, R. C. J. (2009). Behavioural heterogeneity and shift-contagion: Evidence from the Asian crisis. *Journal of Economic Dynamics and Control*, *33*(11), 1929–1944. doi:10.1016/j.jedc.2009.06.002

Durbin, J., & Koopman, S. J. (2012). *Time Series Analysis by State Space Methods* (2nd ed.). Oxford: Oxford University Press.

Ellen, S. Ter, & Zwinkels, R. C. J. (2010). Oil price dynamics: A behavioral finance approach with heterogeneous agents. *Energy Economics*, *32*(6), 1427–1434. doi:10.1016/j.eneco.2010.03.003

Flassbeck, H., Bicchetti, D., Mayer, J., & Rietzler, K. (2011). Price formation in financialized commodity markets: The role of information. *United Nations Publication*, *UNCTAD/GDS*.

Gorton, G., & Rouwenhorst, K. G. (2005). Facts and Fantasies about Commodity Futures. Retrieved from http://fic.wharton.upenn.edu/fic/papers/06/0607.pdf

Hamilton, D. (1994). State-Space Models. In R. F. Engle & D. L. McFadden (Eds.), *Handbook of Econometrics, Volume IV* (Vol. 1, pp. 3041–3080). Elsevier Science B.V.

He, X.-Z., & Westerhoff, F. H. (2005). Commodity markets, price limiters and speculative price dynamics. *Journal of Economic Dynamics and Control*, *29*(9), 1577–1596. doi:10.1016/j.jedc.2004.09.003

Heij, C., de Boer, P., Franses, P. H., Kloek, T., & van Dijk, H. K. (2004). *Econometric Methods with Application in Business and Economics*. Oxford: Oxford University Press.

Hernandez, M., & Torero, M. (2010). *Examining the Dynamic Relationship between Spot and Future Prices of Agricultural Commodities*.

Hommes, C. (2006). Heterogeneous Agent Models in Economics and Finance. In L. Tesfatsion & K. L. Judd (Eds.), *Handbook of Computational Economics, Volume 2: Agent-based Computational Economics* (Vol. 2).

International Cocoa Organisation. (2012). *The World Cocoa Economy: Past and Present*. Retrieved from http://www.icco.org/about-us/international-cocoa-agreements/cat_view/30-related-documents/45-statistics-other-statistics.html

International Cocoa Organisation. (2013). ICCO Website. Retrieved May 12, 2013, from http://www.icco.org/

Julier, S. J. (1999). The Scaled Unscented Transformation. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.46.8404

Julier, S. J., & Uhlmann, J. K. (1997). A New Extension of the Kalman Filter to Nonlinear Systems. In *The Proceedings of AeroSense: The 11th International Symposium on Aerospace/Defense Sensing, Simulation and Controls*. Orlando, Florida.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk, *47*(2), 263–292.

Lagi, M., Bar-yam, Y., Bertrand, K. Z., & Bar-yam, Y. (2011). *The Food Crises : A quantitative model of food prices including speculators and ethanol conversion*. Retrieved from http://ssrn.com/abstract=1932247

OECD, & SWAC. (2011). *Emerging good practice in combating the worst forms of child labour in West African cocoa growing communities*. Retrieved from http://www.oecd.org/swac/publications/49069653.pdf

Orderud, F. (2005). Comparison of Kalman Filter Estimation Approaches for State Space Models with Nonlinear Measurements. Trondheim.

Radetzki, M. (2006). The anatomy of three commodity booms. *Resources Policy*, *31*(1), 56–64. doi:10.1016/j.resourpol.2006.06.003

Rapach, D. E., & Zhou, G. (2012). *Forecasting Stock Returns Forthcoming in the Handbook of Economic Forecasting , Volume 2* (Vol. 2).

Reichsfeld, D. A., & Roache, S. K. (2011). Do Commodity Futures Help Forecast Spot Prices ?

Reitz, S., & Westerhoff, F. (2007). Commodity Price Cycles and Heterogeneous Speculators. *Empirical Economics*, *33*, 231–244.

Rosegrant, M. (2008). *Biofuels and Grain Prices : Impacts and Policy Responses*. Retrieved from http://www.ifpri.org/sites/default/files/publications/rosegrant20080507.pdf

Ruf, F., & Siswoputranto, P. S. (1995). *Cocoa Cylces: The Economics of Cocoa Supply* (pp. 1–54). Woodhead Publishing.

Wan, E. A., & Merwe, R. Van Der. (2000). The Unscented Kalman Filter for Nonlinear Estimation. In *Adaptive Systems for Signal Processing, Communications, and Control Symposium 2000. AS-SPCC. The IEEE 2000* (pp. 153–158).

Westerhoff, F., & Reitz, S. (2005). Commodity price dynamics and the nonlinear market impact of technical traders: empirical evidence for the US corn market. *Physica A: Statistical Mechanics and its Applications*, *349*(3-4), 641–648. doi:10.1016/j.physa.2004.11.015

Appendices

Appendix A. Unit Root tests

The Augmented Dickey Fuller test is applied to the time series Cacao Price, Stock Level and Precipitation. The null hypothesis is that a Unit Root is present. As the table below shows, the null hypothesis is only rejected for the time series Precipitation. This implies that a unit root is present in both the Cacao Price time series and the Stock Level time series.

Unit Root tests

	ADF test statistic	P_{null}
Cacao Price	-2.7960	0.1988
Stock Level	-1.6428	0.7759
Precipitation	-8.1138	0.0000

Appendix B

A different way of estimating the model is possible. Below are all the equations that relate to the fundamental return (ΔF_{t+1}) and fundamental price (F_t), where the error term in the fundamental return equation is removed:

$$\Delta P_{t+1} = \Delta F_{t+1} + W_t(\alpha_1(P_t - F_t)^+ + \alpha_1(P_t - F_t)^-) + (1 - W_t)(\beta_1(P_t - P_{t-1})^+ + \beta_2(P_t - P_{t-1})^-) + \varepsilon_{t+1}$$

$$where:$$

$$\Delta F_{t+1} = c + b_1 P_t + b_2 Weather_t + b_3 \Delta I_t$$

$$F_t = F_{t-1} + \Delta F_t$$

$$W_t = f(A_t^F, A_t^C, \gamma)$$

$$A_t^F = -\sum_{k=1}^{K}(E_{t-k-1}^F(P_{t-k}) - P_{t-k})^2$$

$$E_t^F(P_{t+1}) = P_t + b_1^F(P_t - F_t)^+ + b_2^F(P_t - F_t)^-$$

As can be seen from these equations, at time t there is no uncertainty with regard to the fundamental return at time t+1 (ΔF_{t+1}) and the fundamental price at time t (F_t). They can be plugged into the equation governing the cacao market return ΔP_{t+1}. This equation can subsequently be estimated using MLE.

Appendix C. Analysis of residuals of the cacao market return equation

The residuals of the market return are calculated as follows:

$$\varepsilon_{t+1|t+3} = \Delta P_{t+1|t+3} - f(\boldsymbol{Data}_t, F_t, F_{t-2}, \boldsymbol{\theta})$$

The graph below depicts the residuals over time. As can be seen from the graph, the residuals of the earlier observations seem to be smaller than those of the later observations; this would suggest the presence of heteroscedasticity. Further, the graph does not hint at autocorrelation in the residuals.

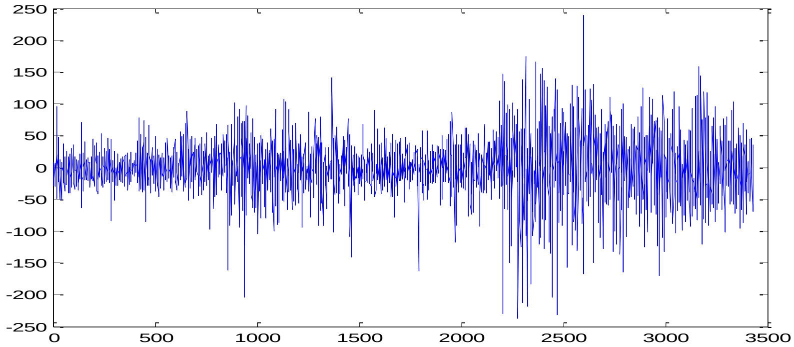

The assumption of homoscedasticity in the market returns is examined in a similar way as above using the GoldFeld-Quandt test as described in Heij et al. (2004). The first third of observations is compared to the third third of observations; this time $k = 8$, as this is the number of regressors in the market return equation. As can be seen from the table below, homoscedasticity is rejected also here:

Goldfeld-Quandt test

s_1^2	869.275
s_3^2	2838.36
F	3.2652
P_{null}	0.000

The specific form of the heteroscedasticity present is further investigated by means of plots of the empirical autocorrelation functions of the absolute residuals of the market equation and of the squared residuals of the market equation below. The upper and lower bound represent the boundaries of a 95% confidence interval. The left figure indicates that there indeed is substantial correlation that only declines very slowly over time; this implies the presence of volatility clustering. The right figure shows that the first 5 lags are most important in explaining the autocorrelation, with two other prominent spikes in both the absolute and squared returns indicating that the 13th and 20th lags are also responsible for the high autocorrelation.

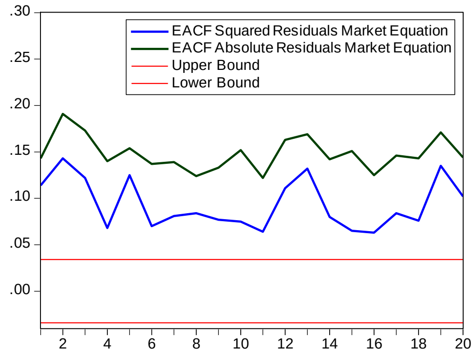

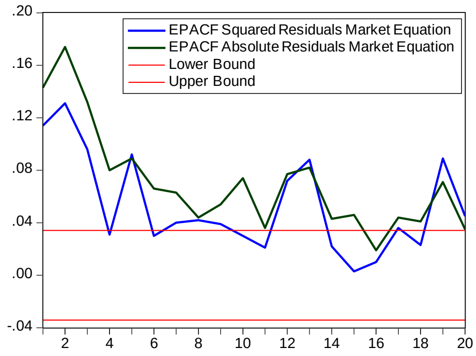

Further, the assumption of no autocorrelation in the residuals of the market return equation is tested by plotting the empirical autocorrelation function of the residuals of the market equation. The first 20 empirical auto correlations are displayed in the graph below, together with a 95% confidence interval. As the figure shows, there does not seem to be a particular autoregressive pattern present in the residuals.