

Discussion Paper

No. 2015-41 | June 04, 2015 | <http://www.economics-ejournal.org/economics/discussionpapers/2015-41>

Radical Uncertainty: Sources, Manifestations and Implications

Christian Müller

Abstract

This paper argues that radical uncertainty is the outcome of standard market activity. The theoretical findings are corroborated with empirical analyses. The model example is applied to asset pricing and radical uncertainty is found a solution to various asset pricing “puzzles”. In conclusion, radical uncertainty should form the basis of economic analysis.

(Published in Special Issue [Radical Uncertainty and Its Implications for Economics](#))

JEL F31 F47 C53

Keywords Rational expectations; uncertainty; subjectivity

Authors

Christian Müller, ✉ German University in Cairo, Faculty of Management Technology,
Department of Economics, 11835 New Cairo City, Egypt, mail@cmueller.ch

Citation Christian Müller (2015). Radical Uncertainty: Sources, Manifestations and Implications. Economics Discussion Papers, No 2015-41, Kiel Institute for the World Economy. <http://www.economics-ejournal.org/economics/discussionpapers/2015-41>

1 Introduction

The 2007–2009 financial crisis has generated an unprecedented surge in the use of the term ‘uncertainty’ among academics and the general public.¹ From an economist’s point of view there are two relevant possibilities to interpret this increased interest in uncertainty. First, there might be a new awareness of the role uncertainty plays in economics. Secondly, researchers express their own discomfort with established models and analyses which makes them adopt a more cautious language by using ‘uncertainty’ more often.

This paper argues that the second interpretation is, unfortunately, most likely correct. The simple reason for this conjecture is the systematic confusion of the meaning of uncertainty that has taken root despite the fact that the issue has been raised very early by Knight (1921) and Keynes (1936).

Today, the economics profession uses a wide range of terms for indicating the various degrees of uncertainty such as radical, fundamental uncertainty, ambiguity risk, and so on while the process of clarifying the terminology through deliberate reflection (Dequech, 2011) is not yet finished. This paper does not engage in this discussion but makes use of the more conventional duality of quantifiable and non-quantifiable uncertainty, the latter of which will be labelled (radical) uncertainty and the latter risk for the sake of convenience.

The main objective of the following analysis is to provide an analytical example to justify the fundamental lack of quantifiability of probabilities in economics. The backbone of this justification is the notion of subjectivity. In contrast to the mainstream approach which starts off with subjective (utility) expectations (?) and then transforms this subjectivity by rational expectations (Muth, 1961) to objective matters, this paper maintains

¹ A repec.org search among all featured journal articles for the term “uncertainty” in the title produced 1’526 hits for the period 2007–June 2010 (37 per month), but only 1’348 for the period 2003–2006 (28 per month).

subjectivity throughout.

To illustrate this point let us consider an example taken from the literature:

The efficient price process satisfies $dp^*(t) = \sigma(t)dw(t)$, where $w(t)$ is a standard Brownian Motion, σ is a random function that is independent of w , and $\sigma^2(t)$ is Lipschitz (almost surely). (Hansen and Lunde, 2006, p.4)

In this case the price process is driven by stochastic variables with well-defined probability distributions. Therefore, the authors conjecture, asset prices are uncertain. However, this is not quite true because all components of the process are measurable and hence, we should use the term “risk” instead. This said, Hansen and Lunde’s framework mirrors the standard in the literature. Quoting their work in this particular way is totally arbitrary and does by no means imply any critique of their achievements at large.

One more feature stands out. Hansen and Lunde’s (2006) process is objective, independent of human interference while it describes the outcome of genuinely human actions. Variants of the same account for numerous conditioning variables which in turn may also be the result of human action. A striking implication of this approach is that prices can be given even if no transaction takes place and no price is quoted. Consequently, human action, i.e. the labour of traders and intermediaries is not really necessary in this model world.

Amazingly, the assumption of objectivity of price processes and, in fact, of many other economic processes widely used in economics, is hardly contested in the literature (Pesaran, 1987, p.11 is a prominent exception).

The remainder of this paper will make a case for (true) uncertainty in economics. To that aim, a theoretical framework will be sketched and then the implications of uncertainty be empirically tested. On the basis of the empirical evidence it will be argued that

economists should make a serious effort to really account for uncertainty.

2 Objective asset pricing models and subjective interference

Suppose that asset prices, p^* would indeed follow a process as in (1)

$$(1) \quad dp^*(t) = \sigma(t)dw(t)$$

with $w(t)$ being a standard Brownian Motion, and σ is a random function that is independent of w , and $\sigma^2(t)$ is Lipschitz (almost surely). (Hansen and Lunde, 2006, p.4)

To obtain an idea of the prices we could switch on the computer, generate values for p^* and plot the result. As mentioned, prices are thus available without actual trades of the asset in question.

Let us now make the following thought experiment. We fix a point in time, say t and ask ourselves what price will result at t conditional on the past and all other information available to us. Depending on the decision theoretical framework, we might employ the von-Morgenstern-Neumann calculus of expected values, or the minimax, or maximin principles (Savage, 1951; Milnor, 1954), or something else. Sidelining the actual choice we may stick to economists' most popular custom and let a representative agent decide.²

The outcome of this exercise will be denoted p_t . In purely theoretical exercises this is as far as one has to go. As soon, as we want to link our theoretical considerations to reality we have to go one step further, however. We have to acknowledge that there exists a whole population of agents who all want to price the asset. We may thus write down

² Instead of a single agent we might also allow for finite heterogeneity of agents as long as there exists a fixed weighted average of decisions making principles.

the individual pricing problem:

$$(2) \quad p_i(t) = p(t) + \varepsilon_i(t)$$
$$\varepsilon_i(t) \sim (0, \zeta(t)^2)$$
$$\zeta(t)^2 < \infty$$

With $i = 1, \dots, N(t)$ denoting the agents in the population. Obviously, once we average about the agents, and letting N be large we obtain the backbone of basically all empirical approaches. On average, all agents agree on the correct price which is $p(t)$. It might be worth mentioning that the same mechanism needs to be assumed if we base our estimation on time series methods. In time series econometrics the key property required for estimation is ergodicity. Ergodicity makes only sense, however, if we consider our observations as being drawn from the true process. Hence, deviations of the true price from the actual price have to be considered small and negligible on average. This is equivalent to the requirement of the average of agents' pricing to converge to the true price.

It is now easy to operationalize uncertainty. Instead of letting $\varepsilon_i(t)$ follow a certain distribution with finite variance we simply drop this assumption. The empirical implications are straightforward. While in the standard situation an increase in $N(t)$ reduces the ambiguity of the true value of $p(t)$, we do not learn more about its true value under the alternative. In other words, if prices get more volatile the more subjects are involved in trading then prices are ruled by uncertainty yet not by risk.

3 The empirical approach

3.1 On the econometrics of quantifiable probabilities

Before actually turning to estimation we have to get a handle on the problem that $p(t)$ in (1) is defined in continuous time. Of course, with t denoting an infinitesimal point, there is no chance to actually observe any price or trader exactly at t . From now on we therefore regard t as a discrete, tiny time period of constant duration, say five or ten minutes, and $N(t)$ will be approximated by the number of actual price quotes within the same period.

We hence write

$$(3) \quad \bar{p}(t) = \frac{1}{N(t)} \sum_{i=1}^{N(t)} p(t)_i$$

$$(4) \quad \bar{\zeta}(t) = \frac{1}{N(t)} \sum_{i=1}^{N(t)} (p(t)_i - \bar{p}(t))^2$$

and the estimates $\bar{\zeta}(t)$ should be unrelated to $N(t)$ in the standard case while we should observe larger $\bar{\zeta}(t)$ the larger $N(t)$ under its uncertain alternative.³ The reason is easy to see. If uncertainty rules a true price does not exist and therefore, more opinions will not reveal more information about the true price. Quite to the contrary, the more opinions are voiced the more volatile prices will get.

One might wonder if the arguments apply to completely irrational agents as well. The answer is yes. As we still suppose the existence of a well-defined process $p(t)$ a positive association of $\bar{\zeta}(t)$ and $N(t)$ would simply indicate that agents are systematically unable to discover the true price process. However, if agents were irrational in that sense, a rational

³ Note that the use of the more popular notion of realised volatility due to Barndorff-Nielsen and Shepard (2002) and Andersen, Bollerslev, Diebold and Labys (2003), among others, instead of the simple variance formula is contingent on the existence of an objective pricing process and hence not appropriate here.

agent would not benefit from his or her superior skills. Knowing the true price when all others are trading on the wrong prices has no practical implications unless the true price will be realised with a given probability. With a positive relationship between $N(t)$ and $\bar{\zeta}(t)$ such a probability cannot be determined, however.

The dependence of $\bar{\zeta}(t)$ on $N(t)$ can be characterised as a subjective interference with objective reality. Looking at the empirical implications discussed above one might likewise say that subjective interference with objective reality may simply be regarded as subjective reality right away. Subjective reality, however, is nothing but uncertainty because humans are all individuals. The individual action is determined by creativity, fantasy, and conscious decision making subject to, of course, constraints and considering incentives. Consequently, nobody can reliably predict all human action and its outcome. This is one important force that makes the states of the environment uncertain, not risky. In short, if we can establish a link between $\bar{\zeta}(t)$ on $N(t)$ the distinction between objective reality with subjective interference and subjective reality becomes superfluous as in both cases uncertainty dominates.

Following the reasoning in the previous section we may now turn to the empirical test. We first write down the hypotheses and then discuss the empirical implementation. The following pair of hypotheses is going to be tested:

$$H_0 : \frac{\partial(\bar{\zeta}(t))}{\partial N(t)} \leq 0 \quad \text{v.} \quad H_1 : \frac{\partial(\bar{\zeta}(t))}{\partial N(t)} > 0$$

To shed light on the validity of H_0 we will proceed as follows. We employ high frequency data of stocks and exchange rates and (i) regress $\bar{\zeta}(t)$ on $N(t)$ in a linear model, (ii) use non-parametric methods to estimate the first derivative of $\bar{\zeta}(t)$ with respect to $N(t)$.

For the sake of brevity all details are referred to the appendix. At this point a summary of the results should suffice. Notice first that the workhorse of all asset pricing models,

the random walk hypothesis, easily passes the test.

To demonstrate the validity of the argument we generate data by a simple random walk

$$\begin{aligned} (5) \quad p_t &= p_{t-1} + \varepsilon_t \\ \varepsilon_t &\sim N(0, 1) \\ t &= 1, \dots, T. \end{aligned}$$

In order to mimic the actual data we let $T = 1000$ and draw between 2 and more than 950 data points randomly disregarding the remaining. Thus, the choices of the lengths of the time series within each 5-minute-interval is adopted from the observed data (in this case Credit Suisse data).

It is well-known, that the variance of a price such as in (5) depends on the *lengths* of the time series, ie the absolute time span within which the process is generated. It does not, however, depend on whether or not we do actually observe a realisation. Therefore, if we observe just a handful of price quotes during the five minutes, it should nevertheless be informative about the variance of the price process in this time interval. If we would observe a multiple of the small number of quotes this should lead to the same estimates about the true variance on average, albeit with a much higher precision.

Figure 1 demonstrates this feature by means of a simulated example in which the simulated process is designed as in (5) and thus represent the standard assumptions. For just a few observations within five minutes the variance-number-of-traders-relation is not estimated very precisely. The more trades are observed the tighter the confidence bands around the point estimates (left hand panel). Eventually, when the number of trades increases but on only a few occasions, the bands get wider again. Throughout, however,

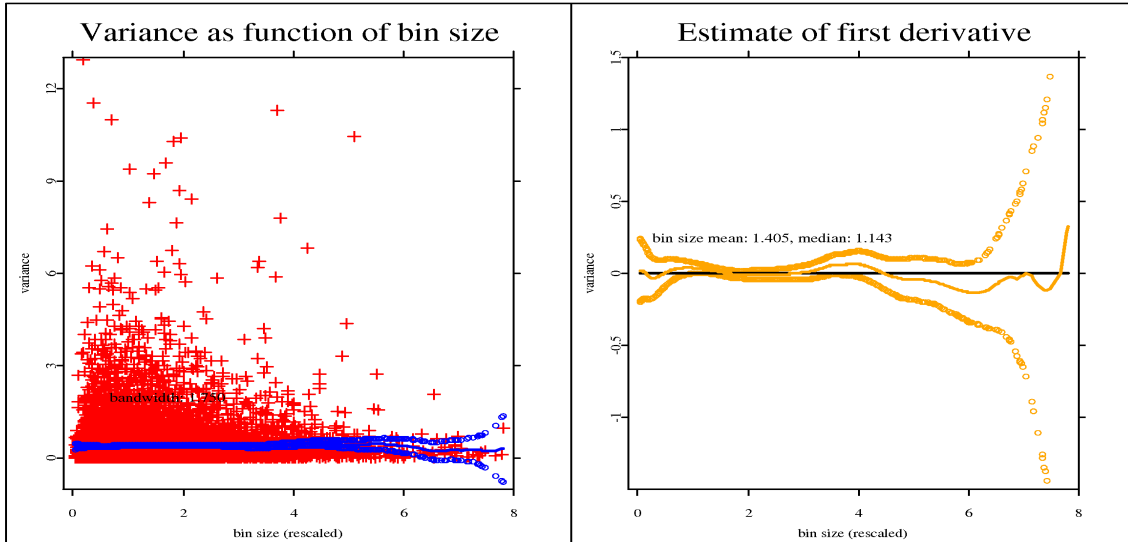


Figure 1: Random walk: Nonparametric estimation of bin size – variance relation and its first derivative

there is no indication of any significant impact of the number of trade(r)s on the price variance as the confidence band always covers the zero line.

Using the Knightian definitions the random walk approach belongs to the category of risk models. Figure 1 clearly shows that the first derivative (right panel) of $\bar{\zeta}(t)$ with respect to $N(t)$ not significantly different from zero. Hence, H_0 is accepted and uncertainty is ruled out.

In contrast, in all data examples H_0 can be safely rejected (see figure 2 below and 8-12 on pp. 35–40). It might be worth noticing that by the logic of science a single rejection of H_0 is sufficient to make the case for uncertainty. In this paper in more than six independent examples H_1 must be accepted.

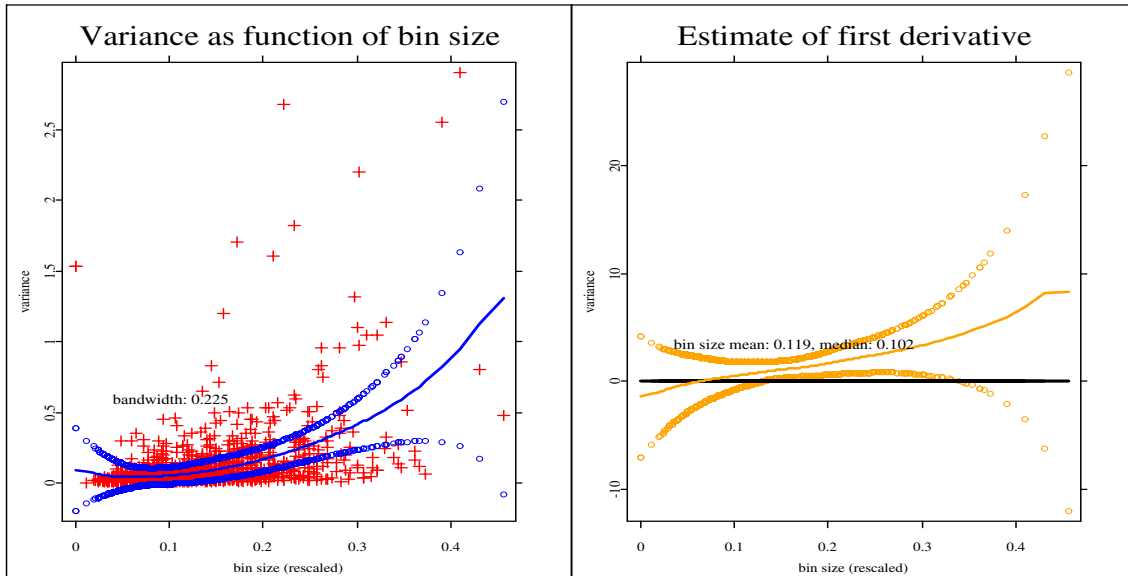


Figure 2: Nestlé tic data 2007: Nonparametric estimation of bin size – variance relation and its first derivative

3.2 A simple subjective asset pricing model

Having rejected H_0 , the objective asset pricing process, the question arises what drives asset prices instead. Therefore, we propose next a simple, very stylised subjective asset pricing model that is able to capture the empirical pattern encountered.

Let us assume an asset market that is characterised by infinite liquidity from an individual's point of view. Infinite liquidity might be justified by noting that a single investor is always small compared to the total or by supposing that credit markets work perfectly in the sense that a convincing investment idea will always meet sufficient means of financing.

Second, let us further assume a very large asset market such that there are always

enough assets available for selling or buying. For example, foreign exchange and stocks of large multinational companies would fall into this category. As we are only considering professional trading, each investor may switch her role between buyer and seller of the asset at any point in time.

3.2.1 The objective function

The investor at the asset market is assumed to act as an inter-temporal arbitrageuse. She buys (sells) if she thinks the asset price in the future to be higher (lower) than today's: $p_t < p_{t+1}$ ($p_t > p_{t+1}$), and the sum invested, x_t is either positive or negative.

Thus, the investor's sole objective is to make profit which probably characterises today's financial markets pretty well. Putting this approach in context one might remember the seven reasons for trading foreign exchange given by Friedman (1953). Only one of them (the seventh) was speculation. All other motives Friedman considered are related to some "real" economic activity such as raising the means for cross-border goods trade. Friedman goes on explaining the price mechanisms for all reasons except the last one. Let us therefore look at number seven.

Unfortunately, at time t , only p_t is known while p_{t+1} is not. The investor can, however, consider a certain range of values for p_{t+1} , and base her decision on the perceived properties, denoted $M_t(p_{t+1} | \mathcal{I}_t)$, of this range which is determined by experience, preferences, and many other factors she thinks worthwhile. These factors are summarised by \mathcal{I}_t . In a standard setting the investor would, for example, attach to each possible future price a probability giving rise to a probability distribution function $M_t(p_{t+1} | \mathcal{I}_t)$ with \mathcal{I}_t representing all the information available at t .

The investment decision is made about the amount of x_t to invest. Obviously, there

are three opportunities for the investor. She can either buy, sell, or do nothing. In order to progress the last option is not considered, it could, however, analytically be included by noting that inaction may incur opportunity costs. There are thus two possibilities left. If she sells the asset and tomorrow's price is higher than today's then she will have made a profit, otherwise she loses money, and vice versa.

It is possible to further extend the model by a feedback from the amount invested to $M_t(p_{t+1} | \mathcal{I}_t)$. Such an extension would not alter the main findings while explaining individual investment plans. For the sake of brevity we continue on a more conventional road side-lining the feedback issue for the time being. We are now ready to establish market equilibrium by a standard order book mechanism.

3.2.2 The investment rule

The individual rule is to invest as much and as long in the market as there is a difference between a subjectively predetermined value $m_{\mathcal{I}_t}(p_{t+1})$, and the spot price. We will call this price the individual reservation price. For example, if the investors formulate a probability distribution over p , then this value might correspond to a certain quantile of $M_t(p_{t+1})$. Thus, the general investment rule can be given as

1. Asset supply: $x_t^+ = x_t$ if $p_t > m_{\mathcal{I}_t}(p_{t+1})$.
2. Asset demand: $x_t^- = -x_t$ if $p_t \leq m_{\mathcal{I}_t}(p_{t+1})$.

3.2.3 The median investor

Assume a finite number $J \geq 1$ of *distinct* investors $j = 1, 2, \dots, J$, who individually form beliefs about the future asset price. They offer and demand the asset according to the

aforementioned rule. Demand and supply coincide under the following conditions.

DEFINITION 1 (Market clearing and equilibrium price). *The market clears if*

$$\sum_{i=1}^J x_{t,i} = 0.$$

The market is in equilibrium if

$$p_t \leq p_t^+ = \min\{m_{\mathcal{I}_t}^{(1)}(p_{t+1}), m_{\mathcal{I}_t}^{(2)}(p_{t+1}), \dots, m_{\mathcal{I}_t}^{(i)}(p_{t+1}), \dots\} \forall i = 1, 2, \dots, J^+.$$

and

$$p_t \geq p_t^- = \max\{m_{\mathcal{I}_t}^{(1)}(p_{t+1}), m_{\mathcal{I}_t}^{(2)}(p_{t+1}), \dots, m_{\mathcal{I}_t}^{(j)}(p_{t+1}), \dots\} \forall j = 1, 2, \dots, J^-.$$

where J^+ and J^- count the suppliers and the sellers of the asset respectively. Any price $p_t, p_t^- \leq p_t \leq p_t^+$ is an equilibrium price.

Ordering all sets $\{m_{\mathcal{I}_t}^{(i)}(p_{t+1}), x_{t,i}\}$, $i = 1, 2, 3 \dots, J$ from smallest to largest $m_{\mathcal{I}_t}^{(\cdot)}(p_{t+1})$ obtains the market price as $p_t = m_{\mathcal{I}_t}^{(*)}(p_{t+1})$ where $m_{\mathcal{I}_t}^{(*)}(p_{t+1})$ corresponds to the median $x_{t,*}$ of the ordered sequence.

An interesting case for the market solution is $J = 1$. It follows that $x_{t,1} = 0$ and hence no transaction takes place. Nevertheless, the spot price is defined. However, as no transaction takes place, the ‘true’ spot price cannot be observed, instead the last period’s will feature in the statistics.

This situation would arise if *all* agents expect the same future spot price, have the same probability distribution in mind including identical attitudes toward risk and accordingly want to either go short or long. The latter makes sure that the asset is neither supplied nor demanded.

As an example consider international investment banks which trade collateralised debt obligations. Since these papers are in general not traded at exchanges quantitative methods are employed to price them. The more investors rely on similar or even identical pricing models the lower J will be. Therefore, the spread of (similar) quantitative pricing models may have contributed to the 2007 / 2008 financial crises, when the securities' market effectively collapsed.⁴

3.3 Properties of the subjective asset pricing model

It remains to be shown that the prices generated by the subjective asset pricing model do indeed share main features with the actual data. Given the model, all information sets, all individual probability distributions and the market solution, p_t can be calculated. We define

$$\mu^{(J-1)} := \frac{1}{J-1} \sum_{i=1}^{J-1} m_{\mathcal{I}}^i.$$

The $\mu^{(J)}$ and $p_t^{(J)}$ are defined accordingly. Being the J th investor the pre-condition for objectivity would therefore be

$$(6) \quad \mu^{(J)} = \mu^{(J-1)}.$$

Equation (6) implies that the individual reservation prices converge to a fixed number, such that if the pool of investors was growing, the observed quantiles would converge to a stationary number.

⁴ These markets finally died when it became apparent that all models generated too high prices. As a result only one model ($J = 1$) survived. This model made all owners of collateralised debt obligations trying to sell.

The answer to the question whether there is a J for which condition (6) hold is, however, no. Notice first that the market solution requires $p_t^{(J)} - p_t^{(J-1)} = \gamma \neq 0$.⁵ Then write

$$(7) \quad \mu^{(J)} = \frac{1}{J-1} \sum_{i=1}^{J-1} p_t^{(J-1)} - \frac{1}{J(J-1)} \sum_{i=1}^{J-1} p_t^{(J-1)} + \frac{1}{J} p_t^{(J-1)} + \frac{1}{J} \sum_{i=1}^J \gamma$$

$$(8) \quad = \mu^{(J-1)} + \frac{1}{J} p_t^{(J-1)} - \frac{1}{J} \mu^{(J-1)} + \gamma$$

to see that only the two middle terms in (8) disappear for large J , whereas the last remains no matter how large J gets. Therefore, for non-degenerate values of γ a limiting value for $\mu^{(J)}$, $J \rightarrow \infty$ does not exist. Hence, p_t is nonstationary in J and an objective distribution probability does not exist. Instead, the distribution always depends on J implying that it is inherently subjective. Nonstationarity w.r.t. time of the first moment and hence every higher moment follows directly.

This is exactly what the actual data properties imply: the variance of the prices does not depend on time but on the frequency of trades, and hence by analogy on the number of traders. Therefore, the subjective asset pricing model is able to replicate the empirical observations.

3.4 Caveats and further research proposals

The empirical investigation showed that the data has properties which one would expect if investors behaved according to the median model. However, the median model poses the absence of an objective price process and this absence is impossible to prove since the non-existing can not be proven to not exist. Therefore, the empirical evidence can

⁵ Strictly speaking, $\gamma \neq 0$ only holds for sure if the new investor's investment exceeds $x_{t,*}$, the median investor's investment. Otherwise, several investors have to enter. The line of argument is neither affected by this special case nor by conditioning γ on J .

be interpreted as an *as if* behaviour. Investors behave *as if* there was no objective price process. This interpretation also holds the key for reconciling experimental evidence of investors' behaviour.

The first stylised fact is the so-called irrational behaviour in artificial asset markets (see inter alia Smith, Suchanek and Williams, 1988; Cipriani and Guarino, 2005). It has likewise be demonstrated that experienced traders can push the market price towards its fundamental value and hence eradicate irrational prices (see e.g. Dufwenberg, Lindqvist and Moore, 2005; Drehmann, Oechsler and Roeder, 2005; Hussam, Porter and Smith, 2008, to name but a few). Notably, all these experiments use a design in which an (implicit) objective price process is induced. For example, the traded asset may yield a return with a given probability each period. Therefore, irrationality in such a situation might be used as an argument against the median model. I prefer a different interpretation, however. The participants in these experiment behave exactly as they would have done in the real world: they trade as if there was no objective price process. By contrast, expert traders are able to discover the induced pricing rule and hence tend to behave rationally. Therefore, these experiments do not lend support to the standard approach. The decisive question is how do experts trade in the absence of an objective price process? Thus, the need for an accordingly set up experiment remains and economists might have a closer look at optimal decision making under the subjective probability approach in general.

Therefore, in the light of more realistic properties of the theoretical model I consider the conclusion of rational "irrationality" of asset markets the more plausible one.

Alternatively, one may regard each tic as a piece of information itself. In the logic of my argument every investor would represent an indispensable piece of information. The standard REH approach would thus have to include all investors in the information

set.⁶ Then, the standard approach and my model would generate data which would be observationally equivalent.

Interestingly, recent research into the impact of public announcements on asset markets (Carlson and Lo, 2006; Omrane and Heinen, 2009, among many) shows that the breaking of publicly available news clearly leaves its traces in the data by increased trading activity, for instance. At the same time, however, similar traces are also found when no new information arrives. Researchers very often link this phenomenon to the presence of private information and other market frictions. The remarkable distinction to the subjective modelling approach hence is that the subjective approach does not rely on two unproven assumptions. The standard approach starts with positing the existence of an objective pricing process and when implausible consequences arise, yet another assumption is made to fix the problem. The subjective model requires fewer such assumptions while yielding more explanatory power and should therefore be regarded superior to the objective modelling approach.

4 Summary and conclusions

Not the least forced by the 2007 – financial crisis economists and the public at large are aware of the fact that radical uncertainty needs to be accounted for in economic analysis. This paper presents a simple subjective asset pricing model that aims at justifying the dominance of uncertainty as opposed to quantifiable risk in market transactions. Next to the analytical arguments, an example taken from the empirical asset pricing literature demonstrates the results. Asset pricing issues are a perfect role model because it is haunted

⁶ Arguably, with such modification the idea of a representative agent disappears.

by a number of puzzling results such as the lunch break puzzle, the gone-fishing-effect, and seemingly “irrational” behaviour of investors at large. This paper suggests to put these puzzles in the context of real uncertainty for solving them. Using the simple subject asset pricing model it can be shown how these data features can be generated without recurring to auxiliary assumptions regarding the existence of unobservable private information, for example. The main implications of the subjective asset pricing model have found overwhelming empirical support.

A Empirical hypothesis testing

A.1 The data

In the first exercise we use stock prices of frequently and internationally traded stocks: Nestlé and Credit Suisse. Nestlé, is a Swiss company which is one of the largest enterprises in Europe. Likewise, Credit Suisse is one of the biggest banks on the continent.

The data at hand covers two distinct periods. The first stretches over January and February 2007 (Nestlé only) which can be considered a quiet and ‘normal’ market period. The second sample starts on January, 1st and ends on July 31, 2009.

These three data sets comprise more than 84’400 (Nestlé, 2007 sample) and more than 1.3 million (Nestlé and Credit Suisse each, 2009 sample) observations. These prices are next aggregated into 1650 (Nestlé, 2007 sample) and 14675 (Nestlé and Credit Suisse each, 2009 sample) ten and five minutes bins respectively.

Table 1 on page 19 summarizes the data characteristics, while figure 3 provides a plot of the 2007 Nestlé data. In the top panel we see a cross plot of the data while the bottom panel presents a non-parametric density estimate of the number of observations,

Table 1: Data characteristics

per bin	mean	min	max	variance
Nestlé share prices January and February 2007				
Total number of bins: 1650				
no. tics	49.38	.0	188.0	729.07
variance	.09	.0	2.91	0.04
Nestlé share prices January – July 2009				
Total number of bins: 14675				
no. tics	89.47	2.0	773.0	3327.3
variance	.0007	0.0	.059	1.8e-6
Credit Suisse share prices January – July 2009				
Total number of bins: 14675				
no. tics	88.684	2.0	955.0	4268.9
variance	.003	0.0	.280	348.7e-6

Sources: Swiss stock exchange, own calculations.

that is the sizes of the bin. These two plots already do suggest that the variance tends to increase with the number of observed trades. Turning to formal methods this impression is corroborated.

A.2 Linear models

A.3 Empirical evidence across time, ...

The following regressions analysis sheds light on the relationship between variance and number of tics. Because the functional form of this relation is unknown I use a seventh order ($i_{max} = 7$) Taylor approximation of the true functional relationship between ζ_t and N_t to begin with:

$$(9) \quad \bar{\zeta}_t = \alpha_0 + \alpha_1 N_t + \alpha_2 N_t^2 + \dots + \alpha_{i_{max}} N_t^{i_{max}} + \epsilon_t$$
$$\epsilon_t \sim i.i.d.(0, \sigma^2)$$

Table 2: Estimation results: Coefficient estimates and residual standard deviation

i_{max}	α_0	α_1	α_2	α_3	α_4	α_5	$\hat{\sigma}_\epsilon$
Nestlé share prices January and February 2007							
1	-0.059 (-6.45)	0.003 (19.0)	n.a. n.a.	n.a. n.a.	n.a. n.a.	n.a. n.a.	0.178493 -
4	0.171 (4.29)	-0.01 (-3.60)	0.0002 (3.73)	-1.8e-6 (-3.29)	5.5e-9 (3.41)	n.a. n.a.	0.172287 -
Nestlé share prices January – July 2009							
3	0.054 (1.53)	0.006 (7.91)	6.6e-6 (1.60)	2.1e-8 (3.67)	n.a. n.a.	n.a. n.a.	1.22 -
Credit Suisse share prices January – July 2009							
5	0.0 -	0.0196 (7.75)	0.0003 (7.39)	-1.96e-6 (-9.93)	4.5e-9 (12.8)	-2.8e-12 (-14.2)	5.16271 -

Coefficient estimates and corresponding t -values in parentheses below.

Although the variables exhibit a time subscript the regressions are essentially cross section regressions. There may be occasions on which there are periods of generally higher or particular low $\bar{\zeta}_t$ around but under the null hypothesis (the standard approach) this should not be related to N_t .

Applying standard model reduction technologies such as general-to-specific F-testing and selection criteria (Akaike, Final Prediction Error, Schwarz) I derive a suitable representation of the data. In most cases the optimal order seems to be four. Next, the first derivative with respect to the number of observations within each bin is calculated and evaluated for the data range. The following table collects the optimally fitting models and in one instance (Nestlé share prices in 2007) also a model variant where a simple linear model is estimated ad hoc.

Yet another, nonparametric estimation of the relationship between bin size and variance is reported in the appendix. All methods deliver the same results qualitatively.

In the case where a simple linear model is estimated a standard t-test can be used for evaluating the validity of the subjective model. The null hypothesis maintains the standard case while the alternative corresponds to the subjective asset pricing model.

$$(10) \quad H_0 : \alpha_1 = 0 \quad vs. \quad H_1 : \alpha_1 > 0.$$

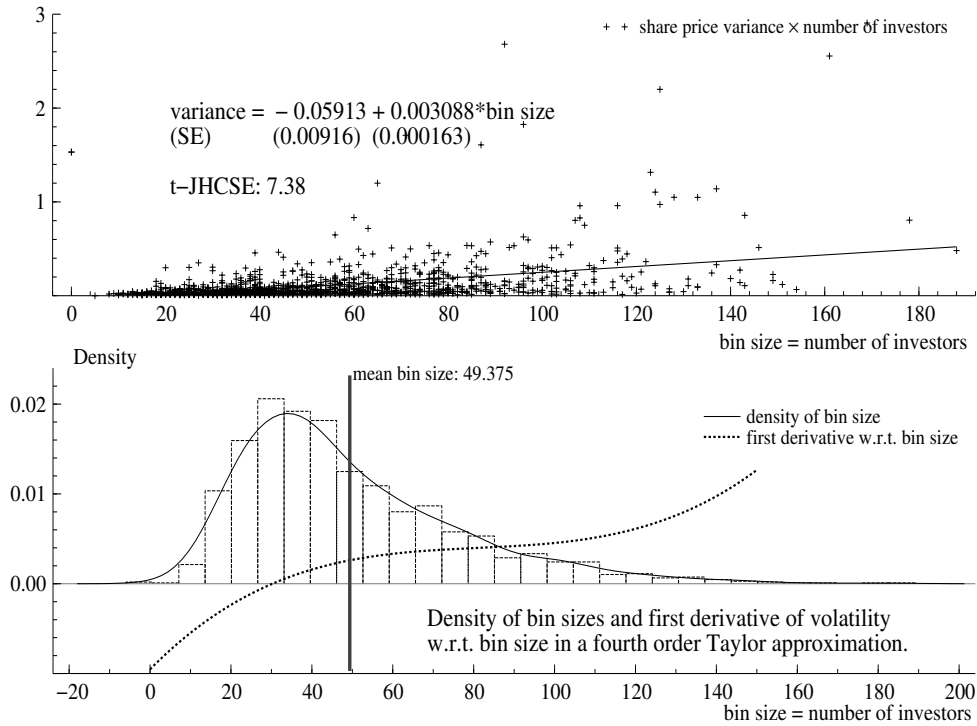


Figure 3: Nestlé 2007: Data plot and graphical estimation results

The estimation results are reported in table 2. They point strongly to a positive relationship between the number of trades and the variance of the price. This is in stark contrast to the usual conviction that more trades would reveal more information about the true price. Instead of increasing the precision with which we measure the price by using more observations it does in fact decrease.




Figure 4: Nestlé 2009: Data plot and graphical estimation results

Because it is not easy to gauge the first derivative with respect to the number of observations from the coefficient estimates, we provide plots of the derivatives. It turns out that the first derivative is positive around the mean. This can be inferred from the lower panel of figures 3 to 5 where the dotted (figure 3), or smooth solid line (figures 3, 4) marks the function of the *first derivative*.⁷ All in all there is little doubt that instead of increasing the precision of our price measure the precision decreases when more trades take place for any fixed information set.




Figure 5: Credit Suisse 2009: Data plot and graphical estimation results

⁷ The first derivative is normalised to match the density estimate scale. This adjustment does not affect its position relative to the zero line.

Interestingly, Lyons (2001), and Evans and Lyons (2002) observe similar effects when they report the tremendous increase in the measure of fit of their exchange rate model. The key variable they introduce is order flow data leading to an increase of up to 64% in the measure of regression fit. Moreover, the variables which are in line with economic theory are insignificant on all but one occasion. The result is similar to the present since (cumulated) order flows are under fairly plausible assumptions proportionate to the number of investors. Given the median model no wonder therefore, that Evans and Lyons are able to explain a larger share of the variance.

The evidence presented here, could be challenged on grounds of endogeneity bias. If the number of investors was dependent on the variance of the price process, then the regression coefficients of equation (9) would not be reliable. Therefore, recent papers such as Ané and Ureche-Rangau (2008) investigate the hypothesis that both number of trades (rather: trading volume) and volatility are jointly determined by a latent number of information arrivals. In our context this would imply that the five (ten) minutes time interval was not short enough for keeping the information set constant. In the particular case of Ané and Ureche-Rangau the data is daily price and volume of stocks which certainly justifies modelling information arrivals. However, the general question whether or not trading volume / number of traders is exogenous to the volatility remains.

In support of our regression approach we would like to point to the well-known lunchtime volatility decline.⁸ In fact, for every major asset market, be it stock markets, foreign exchange markets, or bond markets intra-day volatility assumes an U-shape (see e.g., Ito, Lyons and Melvin, 1998; Hartmann, Manna and Manzanares, 2001, and the references therein). Thus, following an exogenously determined decline in the number of investors

⁸ Similar observations have been made with respect to the holiday season. See e.g. Hong and Yu (2009) on the so-called gone-fishing-effect.

(traders) the volatility decreases justifying the assumption of weak exogeneity of numbers of investors. The same U-shape pattern can be found in our data. For the sake of brevity we do not report the details. They are available on request, however.

In sum, the empirical evidence is more in favour of the model presented in section 3.2 than in line with the traditional approach.

A.4 ... , space and markets

The previous sections provide evidence for abandoning the standard macro finance approach in favour of an alternative model that maintains individual rational behaviour while emphasising the role of subjective rationality on the macro level.

However, there are at least two possibilities to match the data evidence with the traditional view. One possibility is offered by infinite variance Lévy processes as price generating processes. These processes also feature a higher variance the more data we observe holding the information set constant. As regards the discrimination between the subjective model and Lévy processes there is little one can do except from experiments. Therefore, objective Lévy processes and the subjective asset pricing model probably generate data with very similar basic characteristics.

The second explanation could be that the five / ten minutes time interval is not short enough for actually keeping the information set constant. If so, the increase in the variance as more observations enter the intervall might simply be a reflection of a variation in the information set.

Is this argument sufficient word of comfort for returning to the standard approach? In my opinion it is not. The reason is very simple. While shares like Nestlé's are traded every other second many of those assets which can be considered alternative investments

and hence conditioning variables in portfolio models, for example, may be traded far less frequently. As an example consider the Swiss bond market. A safe alternative to the Swiss shares would be Swiss government bonds. It can happen that those bonds are not traded at all within hours. Therefore, the assumption made before finds support that within the five / ten minutes time interval the information set remains constant.

Turning the argument around we would need to carefully synchronise the data of interest and the information set, before we take up the standard approach again. Therefore, an inevitable test of macro finance model would have to look at the high frequency data and make sure that during those time spells where the conditioning variables *do not* change the corresponding number of investors *do not* have explanatory power for the variance of the dependent variable. So far, the standard procedure would be to synchronise observation data by using “suitable” time aggregates such as days, weeks, months, or quarters. I do hazard the guess that the synchronisation exercise, however laborious, would always produce the same result namely nonstationarity with respect to the number of trades.

Luckily, high quality data which permits such synchronisation exercise is becoming more readily available. Very recently Akram, Rime and Sarno (2008) have investigated arbitrage on foreign exchange markets, for example. Their high frequency data set consists of matched spot, forward (forward swap) and deposit interest rate data for the currency pairs British Pound / US Dollar, Euro / US Dollar, Japanese Yen / US Dollar. This data will be used in the following to corroborate the previous findings.

Akram et al.’s (2008) main data source is Reuters which is an advantage for the British Pound but less so for the Euro and the Yen as Reuters is not the main trading platform in these latter two cases. Moreover, the Japanese Yen is most heavily traded when Reuters does not collect the data. Therefore, we will only look at the British Pound and the Euro

pairs.

Even though Akram et al. (2008) collect observations at the highest possible frequency available to them there are occasions on which quotes for the swap, the spot, and the interest rates do not occur simultaneously. Therefore, the variables with the lowest trading activity set the limits. The most important effect on the data sample is a difference in the number of observations despite an exact match of the sample period.

Of course, in order to test the model we need to track the market activity as closely as possible. Whenever there are quotes for, say, the spot rate while there are no changes in the interest rate we lose information. That's why we again restrict our analysis to the largest information sets.

The variable of interest is the arbitrage opportunity defined by the covered interest parity condition given below

$$(11) \quad fx_t = fx_t^e \frac{i_t}{i_t^*} + e_t.$$

Equation (11) has it that the spot exchange rate (denoted fx_t) must equal the forward rate (fx_t^e) up to deposit interest rate (i_t) on domestic assets discounted by the foreign interest rate (i_t^*) of the same maturities as the forward contract. As regards the actual data bid and ask prices are available. Using ask and bid quotes provides a much more reliable picture of true arbitrage opportunities. Consequently, for each currency pair we obtain two deviation measures.

A nonzero e_t indicates arbitrage opportunities. Akram et al.'s (2008) analysis focusses on the properties of e_t . They show for example that sizeable arbitrage opportunities exist but these are all very short lived. For the sake of brevity I do not describe the data in detail. All those details are reported in Akram et al. (2008), the data has been downloaded from

Table 3: CIP deviation data characteristics Feb – Sep 2004

per bin	mean	min	max	variance
POUND / USD ask 12 months				
Total number of bins: 19727				
no. tics	139.21	1.0	1524.0	7323.25
variance	9.25	0.0	1772.89	519.85
POUND / USD bid 12 months				
Total number of bins: 19727				
no. tics	139.21	1.0	1524.0	7323.25
variance	10.38	0.0	1649.45	531.97
POUND / USD ask 6 months				
Total number of bins: 19711				
no. tics	131.61	1.0	1521.0	7239.00
variance	2.05	0.0	2282.09	358.70
POUND / USD bid 6 months				
Total number of bins: 19711				
no. tics	131.61	1.0	1521.0	7239.00
variance	2.09	0.0	2182.55	337.41
EURO / USD ask 12 months				
Total number of bins: 19735				
no. tics	129.74	1.0	558.0	5955.00
variance	6.20	0.0	9695.46	4858.83
EURO / USD bid 12 months				
Total number of bins: 19735				
no. tics	129.74	1.0	558.0	5955.00
variance	5.14	0.0	9594.91	4728.03
EURO / USD ask 6 months				
Total number of bins: 19713				
no. tics	117.14	1.0	559.0	5781.86
variance	0.91	0.0	33.08	2.67
EURO / USD bid 6 months				
Total number of bins: 19713				
no. tics	117.14	1.0	559.0	5781.86
variance	0.91	0.0	29.80	2.06

Sources: Akram et al. (2008), own calculations.

Table 4: Estimating CIP deviation variance: Coefficient estimates and residual standard deviation

i_{max}	α_0	α_1	α_2	α_3	α_4	$\hat{\sigma}_\epsilon$
POUND / USD ask 12 months						
4	7485.59 (522.3)	-15.31 (8.94)	0.2237 0.0456	-0.0004 7.78e-7	2.36e-7 3.558e-8	18705.7 -
POUND / USD bid 12 months						
1	7095.85 (380.6)	20.677 (2.51)	n.a. n.a.	n.a. n.a.	n.a. n.a.	16786.22 -
POUND / USD ask 6 months						
1	1478.29 (74.46)	3.562 (0.52)	n.a. n.a.	n.a. n.a.	n.a. n.a.	3466.82 -
POUND / USD bid 6 months						
3	1008.62 (196.1)	12.875 (3.451)	-0.0455 (-0.017)	7.0e-5 (2.41e-5)	n.a. n.a.	3718.51 -
EURO / USD ask 12 months						
2	1682.37 (86.96)	6.86 (2.320)	6.05 (0.012)	n.a. n.a.	n.a. n.a.	6799.21 -
EURO / USD bid 12 months						
3	1667.63 (188.8)	38.67 (3.95)	-0.1445 0.0233	0.0003 0.839e-5	n.a. n.a.	7269.83 -
EURO / USD ask 6 months						
4	240.727 (45.56)	14.25 (1.476)	-0.0955 0.0144	0.0003 5.11e-5	-2.19e-7 5.765e-8	1590.91 -
EURO / USD bid 6 months						
4	229.241 (40.21)	13.285 (1.303)	-0.086 0.0127	0.0002 4.51e-5	-2.03e-7 5.088e-8	1404.12 -

Regressions of variance on number of traders (tics). Coefficient estimates and corresponding t -values in parentheses below.

Dagfinn Rime's website. Rime also kindly provided advise in handling and interpreting the data.

In what follows we will look at derived values for e_t for the two currency pairs Pound / US Dollar, and Euro / US Dollar. For each of these two pairs e_t is calculated for bid and ask spot rates respectively. I investigate forward contracts for twelve and six months because these are the most liquid markets and we therefore most likely obtain a fair picture of the whole market. Taken together, eight data sets are available for analysis.

The observation period is February 13 to September 30, 2004, weekdays between 07:00 and 18:00 GMT which provides up to 2.7 million observations per currency pair and quote (bid or ask). This data is again bundled into five minutes bins.

After going through the same steps of analysis as before it turns out that the standard approach can again be rejected in basically all cases. The first derivative of the function describing the relationship between bin size and variance is positive around mean / median, and relying on nonparametric analysis, there is convincing evidence for this derivative to be significantly positive.




Figure 6: UIP one year ask (top panel) and six months bid (bottom panel) British Pound / US Dollar 2004: Data plot and graphical estimation results

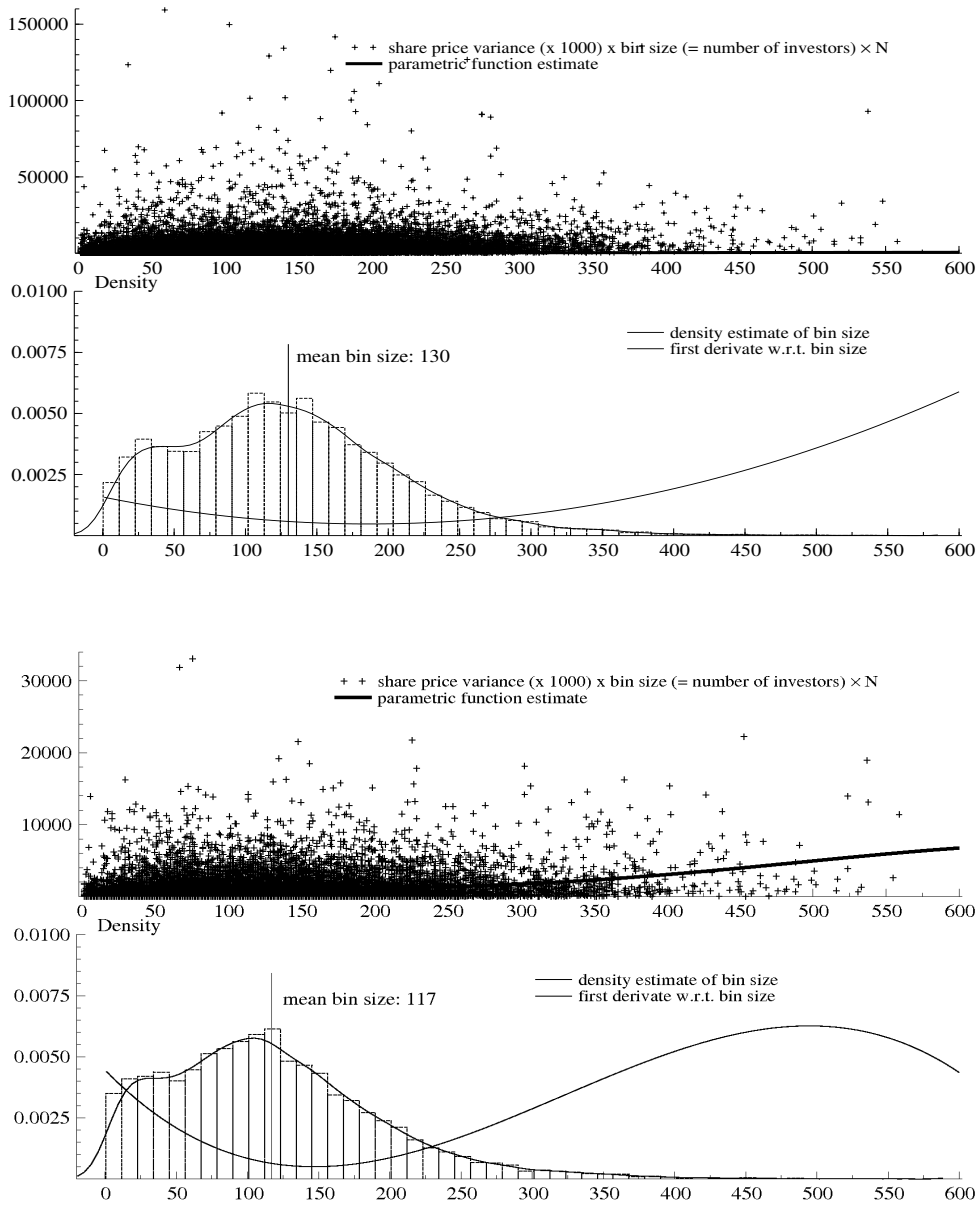


Figure 7: UIP one year bid (top panel) and six months ask (bottom panel) Euro / US Dollar 2004: Data plot and graphical estimation results

B Nonparametric estimation of the bin size – variance relationship

Equation (9) defines a parametric function of the relation between bin size (the approximation of number of traders) and the variance of the asset price within those five / ten minute time bins. The according results lend support to the hypothesis of a positive association between the number of trades and the variance of the asset price. Nevertheless, one may wonder to what extent these results depend on the specific parametric functional forms used. Therefore, I report the outcome of a nonparametric, local quadratic estimation of the relation between bin size and variance.

The estimation is based on the software XploRe which is specifically designed for analysing financial market data by means of non- and semi-parametric functions.⁹ In particular, I make use of the procedure “lplocband” of the “smoother” library applying the Epanechnikov kernel. The kernel bandwidth is chosen manually because the automatic procedures always selected the lowest possible bandwidth within the pre-defined range. These lower bands were close to the minimum distance between any two explanatory variable data points. The results do not change qualitatively, however, within a large range of bandwidths.

Before turning to the empirical evidence let me reconcile the results which could be expected under the null hypothesis, the standard approach. Figure 1 plots observations that are generated by simulating 14675 random walks of length 955. In the next step, between 2 and 955 data points of these random walks are selected randomly from each

⁹ The software is available free of charge from <http://lehre.wiwi.hu-berlin.de/Professuren/quantitativ/statistik/xploRe>, the code and the data are available on request from the author.

of the 14675 data sets. These observations mimic the five minutes bins. Accordingly, the variance of these bins is estimated and set in relation to the number of artificial observations entering the bin. This simulation procedure thus draws on the actual Credit Suisse data and clearly demonstrates that even under the random walk hypothesis for price data the relationship between bin size and variance should be completely stochastic; the first derivative estimate frequently crosses the zero line, and the 95 percent confidence bands safely enclose zero.

By contrast, the empirical relationships do look pretty different. For example figure 2 (p. 10) shows that the estimated first derivative is significantly larger than zero around the mean bin size in the case of the 2007 Nestlé data. Very similar pictures emerge for the other data sets.

In some instances (see figure 11 on p. 39), there are also hints for another phenomenon. In these instances the relationship between variance and bin size seems to be negative. This situation occurs when trading volume is low (small bin sizes) and gives rise to the possibility of dependent observations. For example, when trading activity is low, several consecutive trades may be exercised by the same trader(s).

In order to render the estimation feasible, i.e. avoiding numerical problems, the independent variable was divided by twice the maximum value of the bin size. If that was not sufficient to overcome numerical problems both variables were normalised by their respective empirical standard deviations. This linear transformation cannot affect the relation between the independent and the dependent variable.

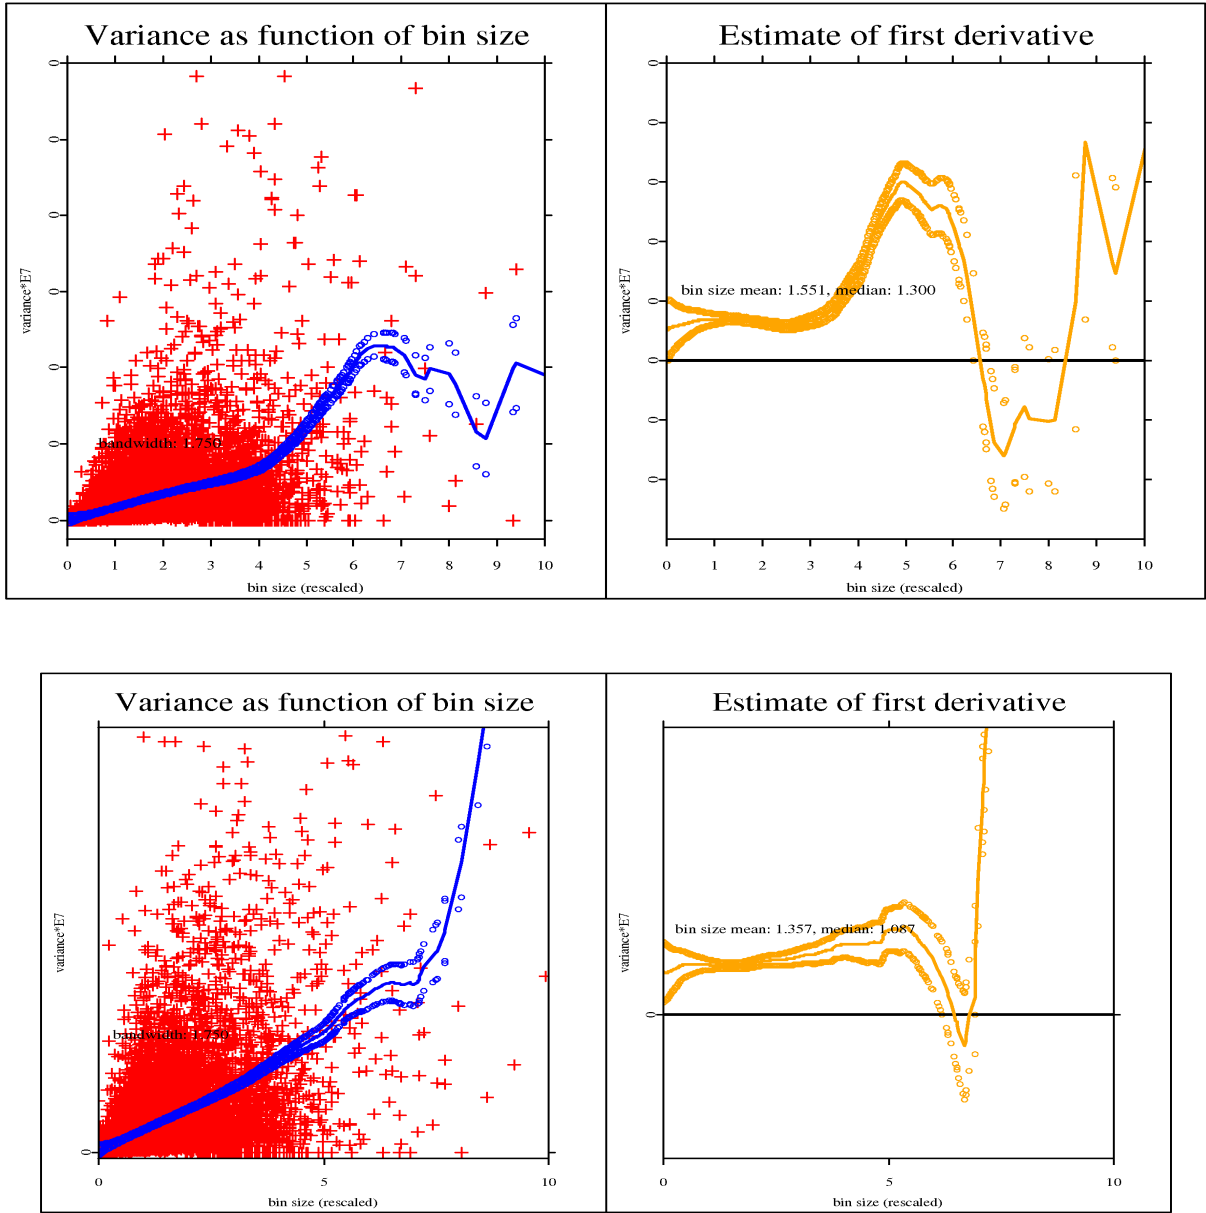


Figure 8: Nestlé tic data 2009 (top panel) and Credit Suisse tic data 2009 (bottom panel):

Nonparametric estimation of bin size – variance relation and its first derivative

Pretty much in line with the parametric estimation the variance increases with the bin size. The panel on the left shows an upward trend in the variance for growing bin sizes and the panel to the right confirms that the first derivative of the relationship is significantly larger than zero around the mean bin size and for sizes larger than the mean. Therefore, the hypothesis derived from the median model receives support once more.

C Evidence from foreign exchange markets

The following graphs depict the results for the data compiled by Akram et al. (2008). Here, the data is always standardized such that the empirical variance of dependent and independent variable is one. As before, this linear transformation cannot affect their relationship.

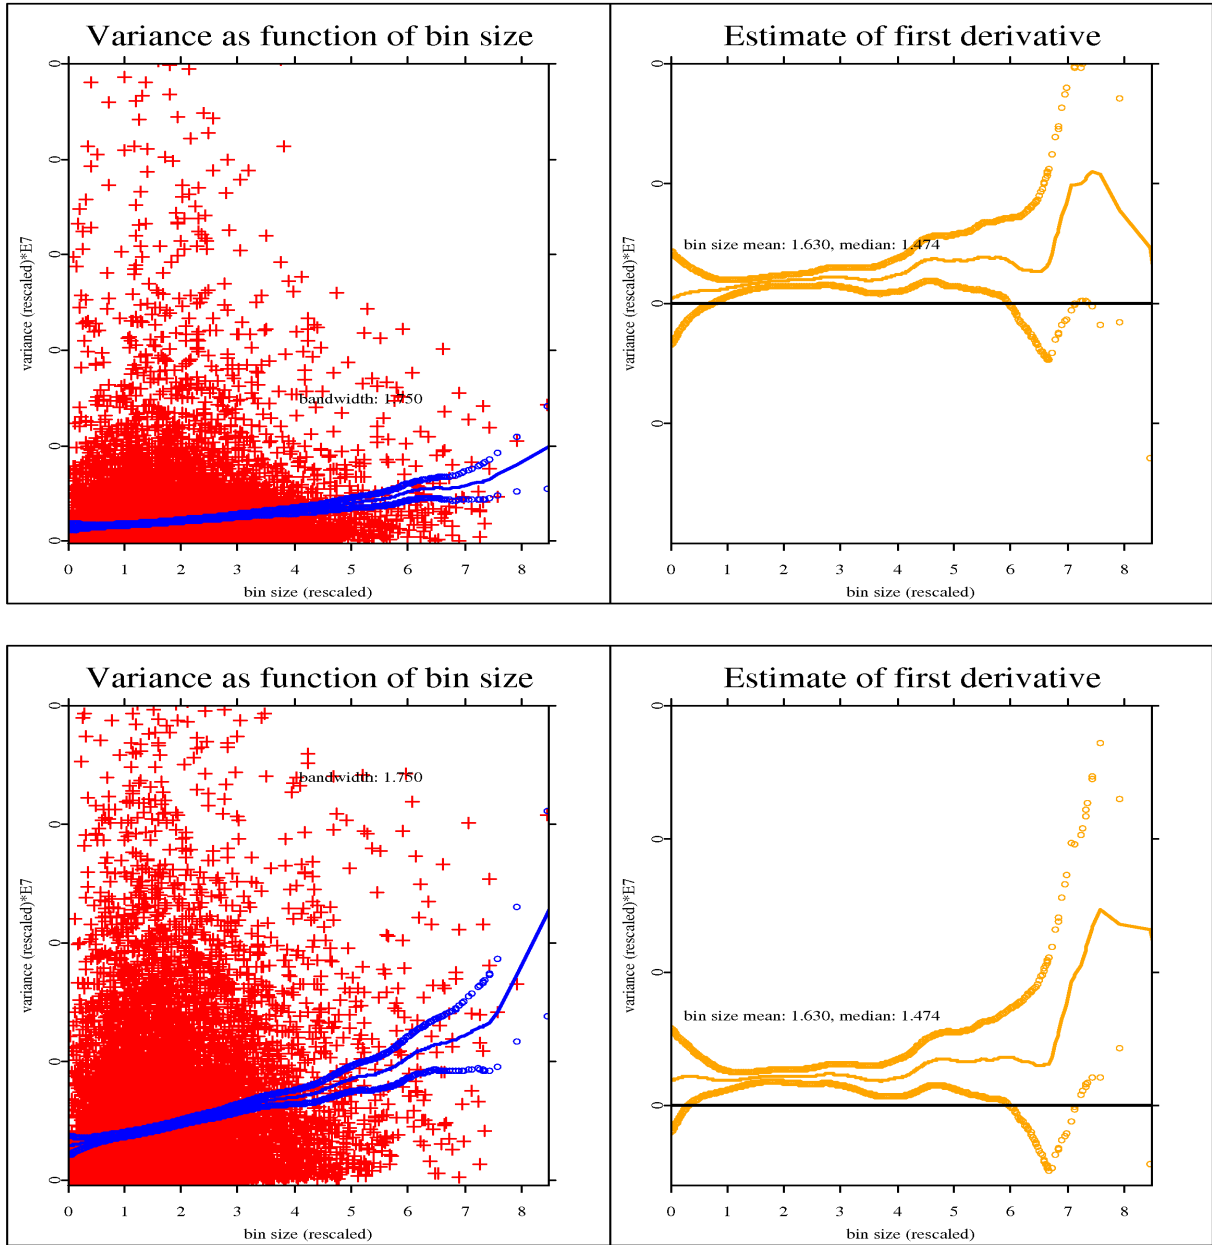


Figure 9: Pound one year ask (top panel) and bid (bottom panel): Nonparametric estimation of bin size – variance relation and its first derivative

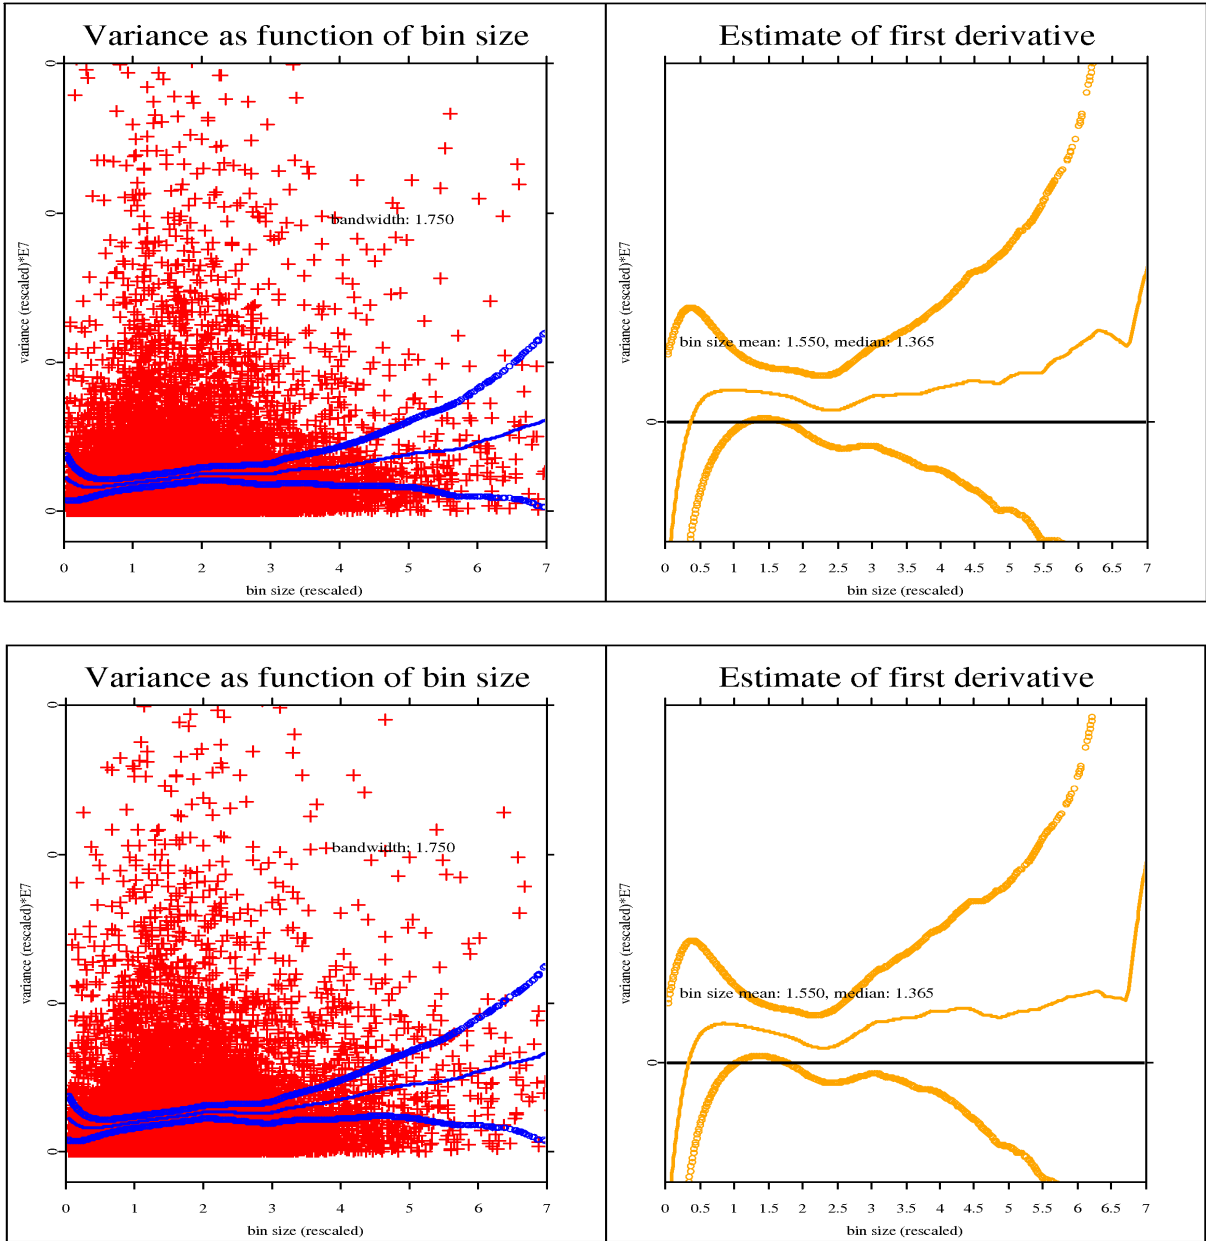


Figure 10: Pound 6 months ask (top panel) and bid (bottom panel): Nonparametric estimation of bin size – variance relation and its first derivative

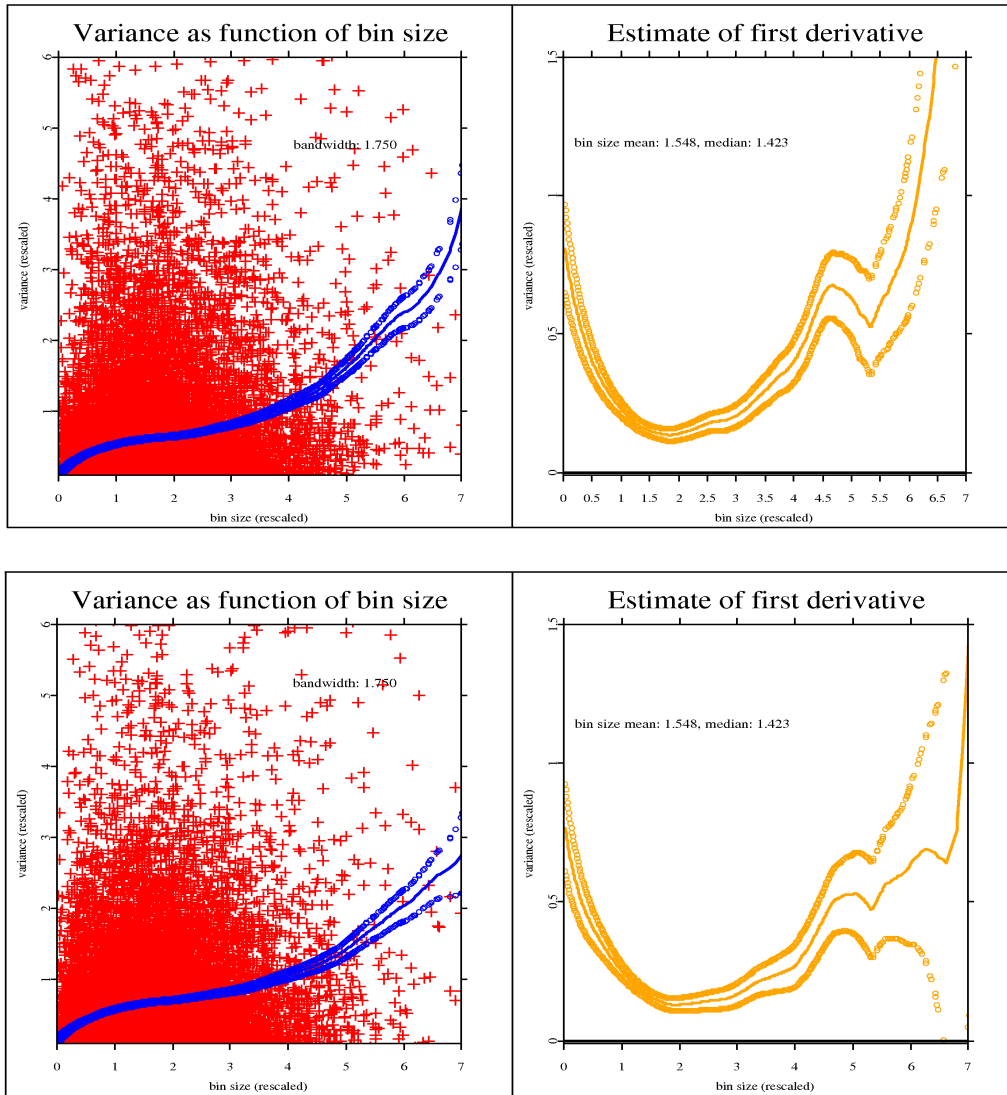


Figure 11: Euro 6 months ask (top panel) and bid (bottom panel): Nonparametric estimation of bin size – variance relation and its first derivative

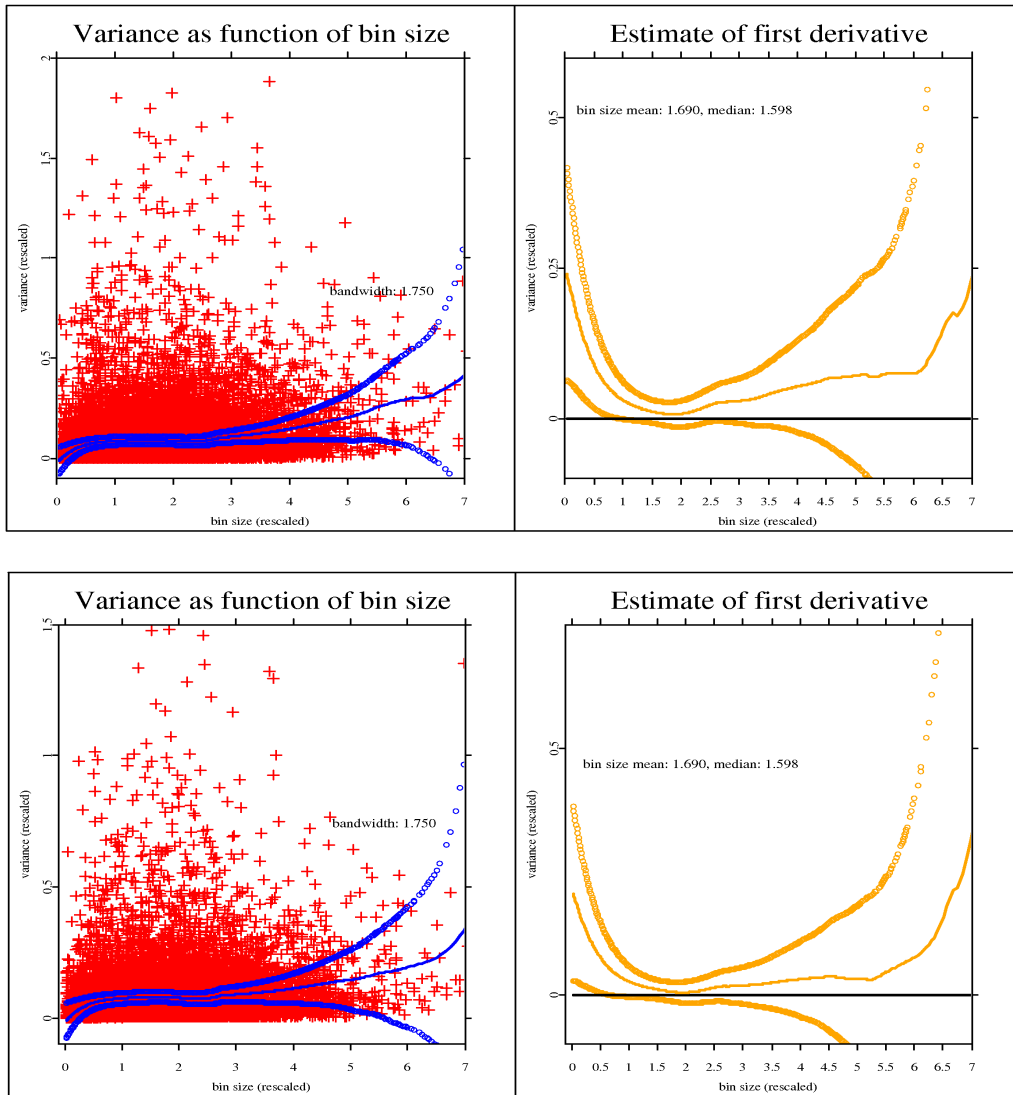


Figure 12: Euro 12 months ask (top panel) and bid (bottom panel): Nonparametric estimation of bin size – variance relation and its first derivative

References

- Akram, Q. F., Rime, D. and Sarno, L. (2008). Arbitrage in the foreign exchange market: Turning on the microscope, *Journal of International Economics* **76**(2): 237–253.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. and Labys, P. (2003). Modeling and Forecasting Realized Volatility, *Econometrica* **71**(2): 579–625.
- Ané, T. and Ureche-Rangau, L. (2008). Does trading volume really explain stock returns volatility?, *Journal of International Financial Markets, Institutions and Money* **18**(3): 216–235.
- Barndorff-Nielsen, O. E. and Shepard, N. (2002). Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Modells, *Journal of the Royal Statistical Society, Series B* **64**: 253 – 280.
- Carlson, J. A. and Lo, M. (2006). One minute in the life of the DM/US\$: Public news in an electronic market, *Journal of International Money and Finance* **25**: 1090 – 1102.
- Cipriani, M. and Guarino, A. (2005). Herd Behavior in a Laboratory Financial Market, *American Economic Review* **95**(5): 1427–1443.
- Dequech, D. (2011). Uncertainty: A Typology and Refinements of Existing Concepts, *Journal of Economic Issues* **45**(3): 621–640.
- Drehmann, M., Oechsler, J. and Roeder, A. (2005). Herding and Contrarian Behavior in Financial Markets: An Internet Experiment, *American Economic Review* **95**(5): 1403–1426.

- Dufwenberg, M., Lindqvist, T. and Moore, E. (2005). Bubbles and experience: An experiment, *American Economic Review* **95**(5): 1731–1737.
- Evans, M. D. D. and Lyons, R. K. (2002). Order flow and exchange rate dynamics, *Journal of Political Economy* **110**(1): 170 – 180.
- Friedman, M. (1953). The Case of Flexible Exchange Rates, in M. Friedman (ed.), *Essays in Positive Economics*, University of Chicago Press, Chicago, pp. 157–203.
- Hansen, P. R. and Lunde, A. (2006). Realized variance and market microstructure noise, *Journal of Business & Economic Statistics* **24**: 127–161.
- Hartmann, P., Manna, M. and Manzanares, A. (2001). The microstructure of the euro money market, *Journal of International Money and Finance* **20**(6): 895–948.
- Hong, H. and Yu, J. (2009). Gone fishin’: Seasonality in trading activity and asset prices, *Journal of Financial Markets* **12**(4): 672–702.
- Hussam, R. N., Porter, D. and Smith, V. L. (2008). Thar She Blows: Can Bubbles Be Rekindled with Experienced Subjects?, *American Economic Review* **98**(3): 924–937.
- Ito, T., Lyons, R. K. and Melvin, M. T. (1998). Is There Private Information in the FX Market? The Tokyo Experiment, *Journal of Finance* **53**(3): 1111–1130.
- Keynes, J. M. (1936). *The General Theory of Employment, Interest, and Money*, Macmillan, London and New York.
- Knight, F. H. (1921). *Risk, Uncertainty, and Profit*, Schaffner and Marx; Houghton Mifflin Company, Boston, MA.

- Lyons, R. K. (2001). Foreign exchange: macro puzzles, micro tools, *Pacific Basin Working Paper Series 01-10*, Federal Reserve Bank of San Francisco.
- Milnor, J. W. (1954). Games Against Nature, in C. H. Coombs, R. L. Davis and R. M. Thrall (eds), *Decision Processes*, Wiley, New York, pp. 49 – 60.
- Muth, J. F. (1961). Rational expectations and the theory of price movements, *Econometrica* **29**: 315–335.
- Omrane, W. B. and Heinen, A. (2009). Is there any common knowledge news in the euro/dollar market?, *International Review of Economics & Finance* **18**(4): 656 – 670.
- Pesaran, H. M. (1987). *The Limits to Rational Expectations*, Basil Blackwell, Oxford.
- Savage, L. (1951). The Theory of statistical decision, *Journal of the American Statistical Association* **46**: 55 – 67.
- Smith, V. L., Suchanek, G. L. and Williams, A. W. (1988). Bubbles, crashes, and endogenous expectations in experimental spot asset markets, *Econometrica* **56**(5): 1119–51.

Please note:

You are most sincerely encouraged to participate in the open assessment of this discussion paper. You can do so by either recommending the paper or by posting your comments.

Please go to:

<http://www.economics-ejournal.org/economics/discussionpapers/2015-41>

The Editor