ESTIMATION OF UTILITY FUNCTIONS: MARKET VS. REPRESENTATIVE AGENT THEORY

Kai Detlefsen [2]

Wolfgang K. Härdle [2]

Rouslan A. Moro [1,2]

[1] Deutsches Institut für Wirtschaftsforschung (DIW)
[2] Center for Applied Statistics and Economics (CASE), Humboldt-Universität zu Berlin

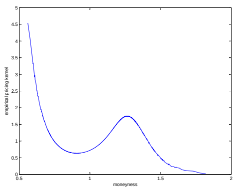

An investor observes the evolution of a stock price in the past and forms his subjective opinion about the future evolution of the price.

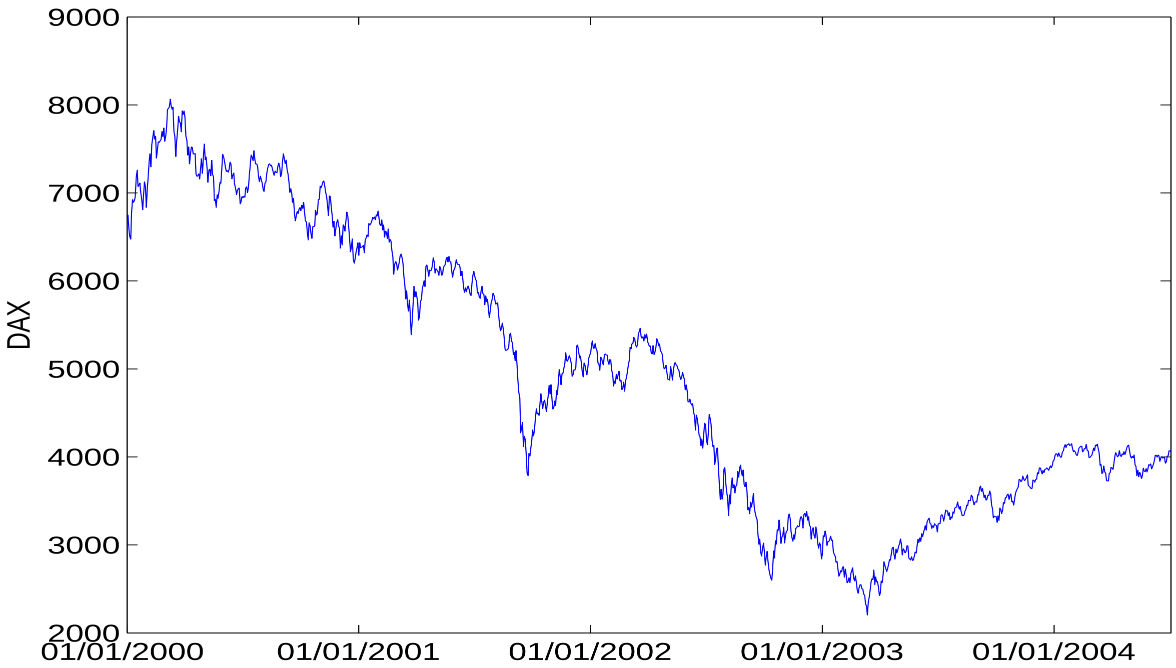

Figure 1: DAX, January 2000 - June 2004. Daily observations.

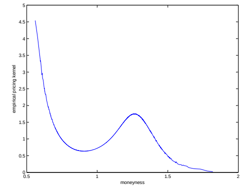

An opinion on the future value S_t of the stock at time t can be described by a density function p which is called **subjective density**, a.k.a. historical density or physical density.

This function can be estimated in many ways (parametric, nonparametric, ...).

Examples:

- Black-Scholes model (Nobel prize 1997): log normal distribution
- GARCH model (Nobel prize 2003, Engle): stochastic volatility
- non-parametric diffusion model (Ait-Sahalia 2000)

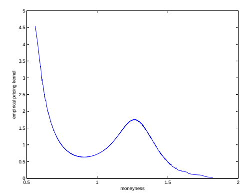

We model the logarithmic returns $\{r_t\}$ of the DAX by a GARCH(1,1) model:

$$r_t = \sigma_t Z_t$$

$$\sigma_t^2 = \omega + \alpha r_{t-1}^2 + \beta \sigma_{t-1}^2$$

From the logarithmic returns $r_i = \log(S_i) - \log(S_{i-1}),\ i = 1, \ldots, t$ and the starting stock price S_0 we can construct the final stock price by

$$S_t = S_0 \exp(\sum_{i=1}^{t} r_i).$$

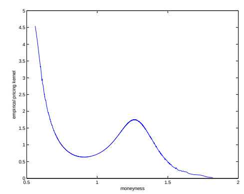

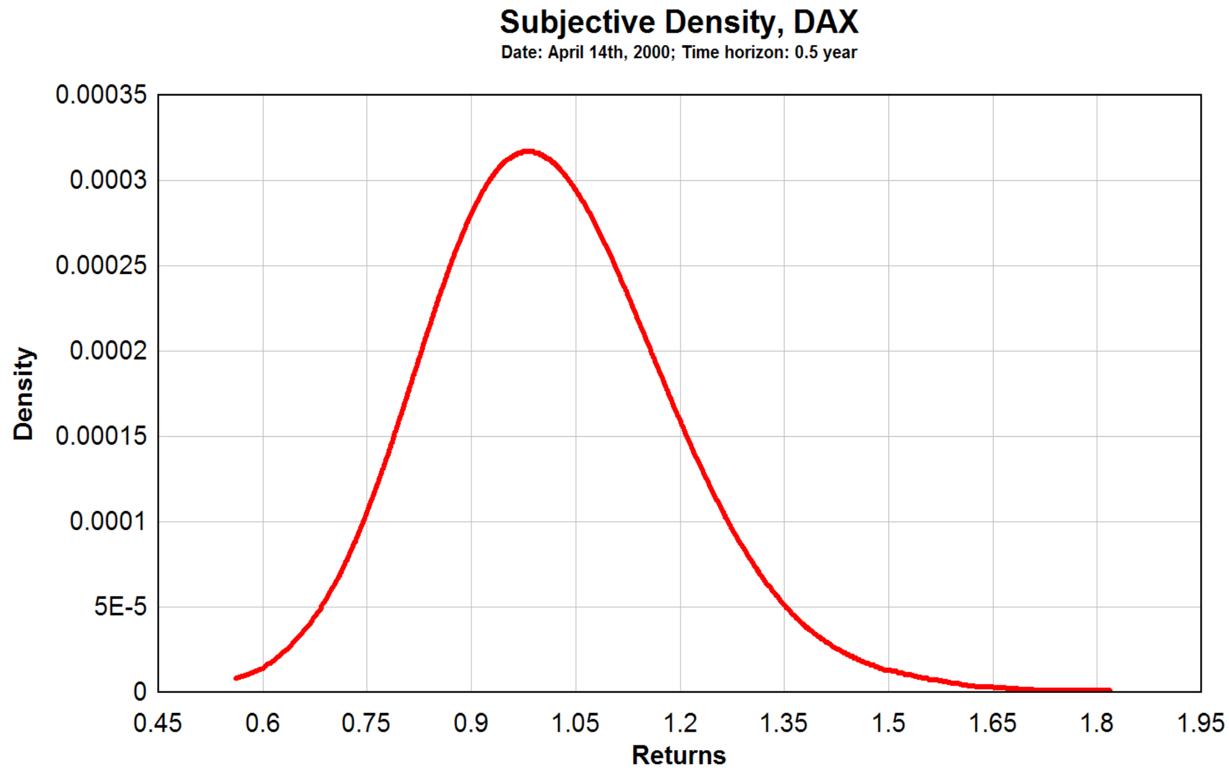

Figure 2: Subjective density on April 14th, 2000 for $\tau = 0.5$ ahead. In order to present the density independent of the starting stock price S_0 we do not plot $S_t \rightarrow \hat{p}(S_t)$ but $R_t \rightarrow \hat{p}(R_t S_0)$ (moneyness scale).

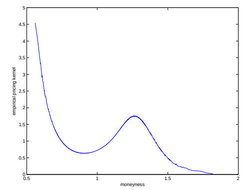

Besides the subjective density there is also a **state-price density** (SPD) q for the stock price implied by the market prices of options, a.k.a. **risk-neutral density**.

The state-price density differs from the subjective density because it corresponds to replication strategies and hence is a *martingale risk neutral measure*.

A person alone does not use in general a replication strategy but thinks in terms of his subjective density.

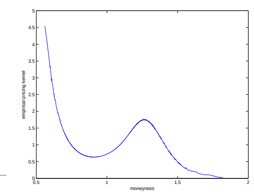

We use the Heston continuous stochastic volatility model, which can be regarded as an industry standard for option pricing models.

The Heston model is given by

$$\frac{dS_t}{S_t} = rdt + \sqrt{V_t} dW_t^1$$

where the volatility process is modelled by a square-root process:

$$dV_t = \xi(\eta - V_t)dt + \theta\sqrt{V_t} dW_t^2,$$

and W^1 and W^2 are Wiener processes with correlation ρ.

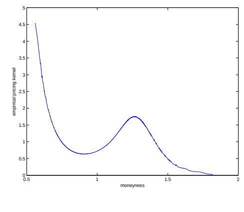

Using option prices with time-to-maturity between 0.25 and 1 and moneyness between 0.5 and 1.5 we get the following estimate for the orisk-neutral state-price density for $\tau = 0.5$ years ahead.

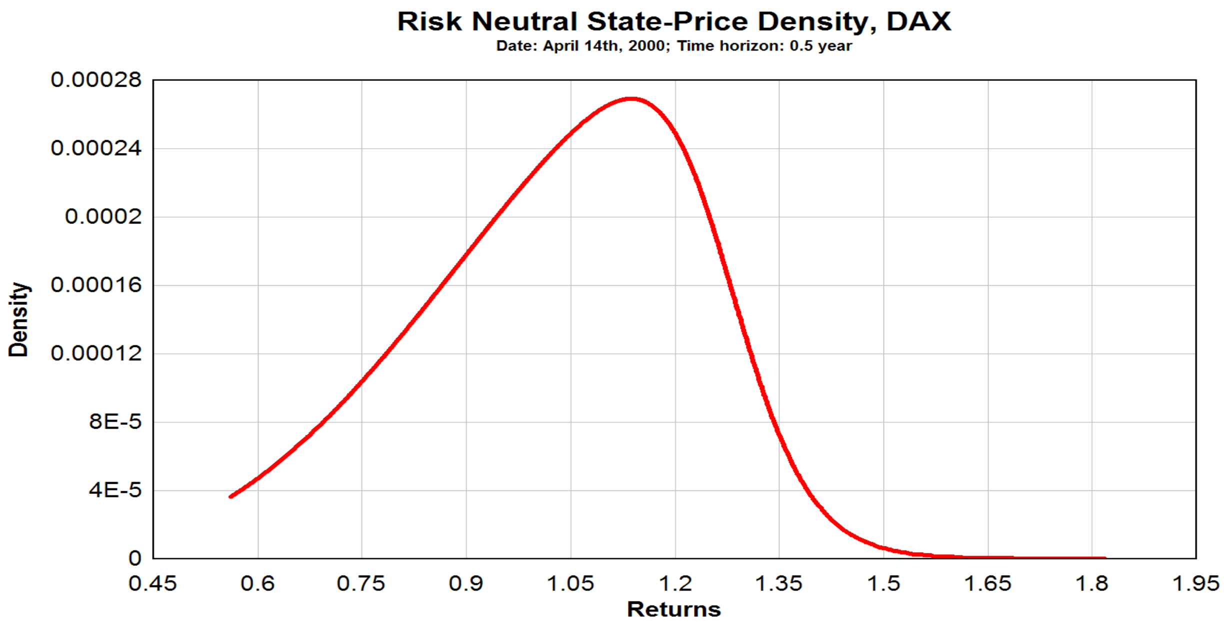

Figure 3: State-price density q_t, $r_{0.5} = 4.06\%$, April 14th, 2000.

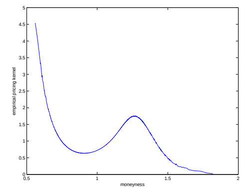

The **pricing kernel** $m(S_t)$ is defined as:

$$m(S_t) = \exp(-rt)\frac{q(S_t)}{p(S_t)}$$

where r is the interest rate with maturity t.

An estimate of the pricing kernel is called **empirical pricing kernel**. We use the estimate:

$$\hat{m}(S_t) = \exp(-rt)\frac{\hat{q}(S_t)}{\hat{p}(S_t)}$$

where $\hat{q}$ and $\hat{p}$ are the estimated risk-neutral and subjective densities.

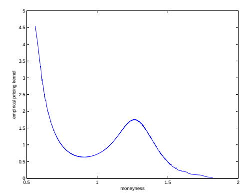

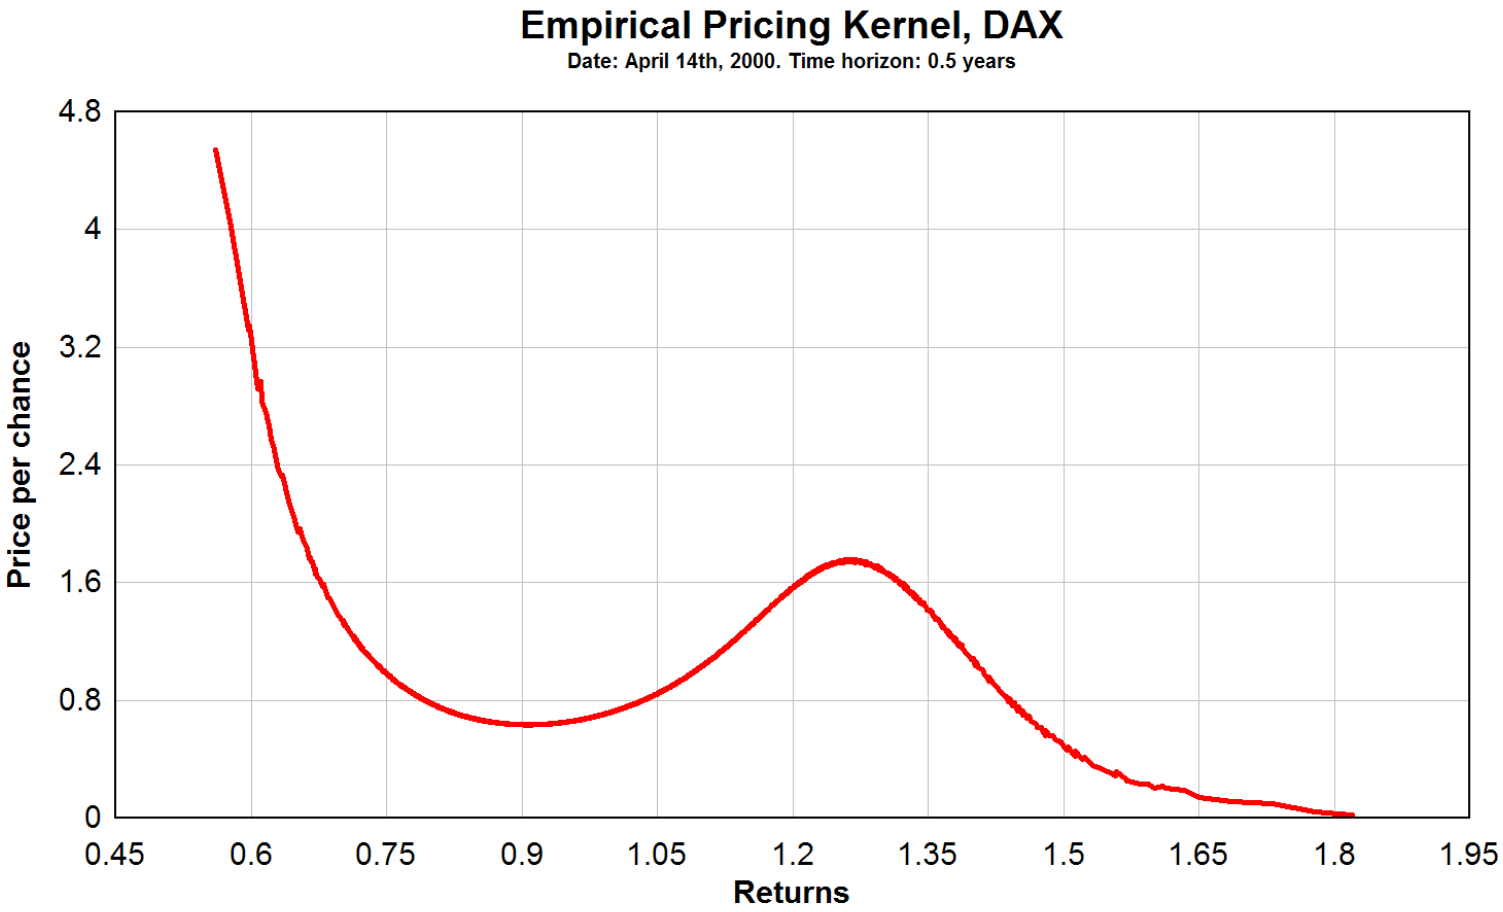

Figure 4: Empirical pricing kernel for $\tau = 0.5$, $r_{0.5} = 4.06\%$, April 14th, 2000.

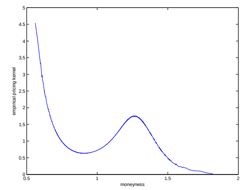

Problems

- How to explain the non-monotonicity of the pricing kernel?
- What type of utility functions can generate observed pricing kernels and prices?
- What happens if the hypothesis of the existence of the representative investor is abandoned?
- How can we experimentally estimate individual pricing kernels and utility functions?

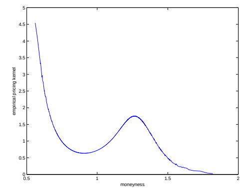

Outline of the Talk

1. Motivation ✓
2. Pricing equation and pricing kernel (stochastic discount factor)
3. Pricing kernel estimation with the Heston and GARCH(1,1) models
4. Decomposition of the market utility function
5. Individual utility functions
6. Market aggregation mechanism
7. Estimation of the distribution of investor types
8. Behavioural experiment design
9. Outlooks

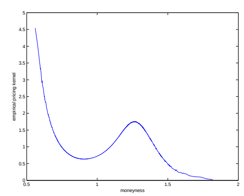

Utility Maximisation Problem

$$\max_{\{\xi\}} U(C_t) + \mathrm{E}_t^* \left[\beta U(C_{t+1}) \right] \tag{1}$$

s.t. $C_t = e_t - P_t \xi$

$$C_{t+1} = e_{t+1} + X_{t+1} \xi$$

where X_{t+1}	–	a pay-off profile of an asset at $t+1$
P_t	–	the price of the asset at t
ξ	–	portfolio position
β	–	discount factor
e_t, e_{t+1}	–	wages at t and $t+1$
E_t^*	–	risk neutral expectation at time t

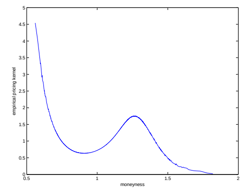

Pricing Equation

If the utility function depends only on state variables and the discount factor $\beta = const$, the price of **any** security paying X_{t+1} at time $t+1$ is:

$$P_t = \mathrm{E}_t\left[\beta \frac{U'(C_{t+1})}{U'(C_t)} X_{t+1}\right] = \mathrm{E}_t\left[m_t^* X_{t+1}\right] \tag{2}$$

where the pricing kernel (PK) a.k.a. stochastic discount factor or price per chance is:

$$m_t^*(C_t, C_{t+1}) = \beta \frac{U'(C_{t+1})}{U'(C_t)} = \mathrm{const}_t \cdot U'(C_{t+1})$$

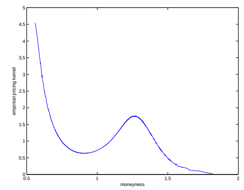

Pricing Kernel Projection

Pricing equation: $P_t = \mathrm{E}_t \left[m_t^*(C_t, C_{t+1}) X_{t+1} \right]$

Pricing equation using the projection of the PK onto asset pay-offs X_{t+1}:

$$P_t = \mathrm{E}_t \left[m_t(X_{t+1}) X_{t+1} \right], \tag{3}$$

where the PK projection is:

$$m_t(X_{t+1}) = \mathrm{E}_t \left[m_t^*(C_t, C_{t+1}) | X_{t+1} \right]$$

Since pricing with m_t^* and m_t is equivalent, we denote $m_t(X_{t+1})$ as the pricing kernel, $U_t(X_{t+1})$ and $U_t'(X_{t+1})$ as a utility and marginal utility function respectively

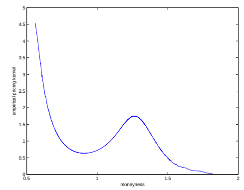

We can write the risk-neutral pricing equation as:

$$P_t = \beta \int_0^\infty X_{t+1} dQ(X_{t+1})$$

where $Q_t(X_{t+1})$ is the observed risk-neutral distribution of returns X_{t+1} at time $t+1$. It is equivalent to

$$P_t = \beta \int_0^\infty X_{t+1} \frac{q_t(X_{t+1})}{p_t(X_{t+1})} dP(X_{t+1})$$

where $P_t(X_{t+1})$ is a subjective distribution, or

$$P_t = \int_0^\infty m_t(X_{t+1}) X_{t+1} dP(X_{t+1}) = \mathrm{E}_t \left[m_t(X_{t+1}) X_{t+1} \right],$$

where the pricing kernel $m_t(X_{t+1}) = \beta \frac{q_t(X_{t+1})}{p_t(X_{t+1})}$

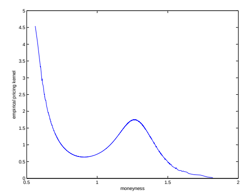

Estimation of the Pricing Kernel

The empirical pricing kernel is:

$$\hat{m}_t(X_{t+1}) = \beta \frac{\hat{q}(X_{t+1})}{\hat{p}(X_{t+1})},$$

where $\hat{q}$ and $\hat{p}$ are the estimated risk-neutral and historical subjective densities; $\beta = e^{-r}$ is a discount factor.

PK is estimated with parametric models:

- the risk neutral density q_t from option prices with the Heston model
- the historical subjective density p_t from stock prices with the GARCH(1,1) model

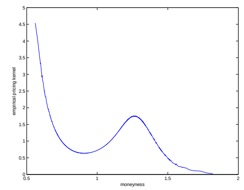

Estimation of the Risk Neutral Density q_t

Risk neutral density q_t is estimated from DAX option prices using the stochastic volatility Heston model:

$$\frac{dS_t}{S_t} = rdt + \sqrt{V_t}dW_t^1$$

where the volatility process is:

$$dV_t = \xi\left(\eta - V_t\right)dt + \theta\sqrt{V_t}dW_t^2$$

W_t^1, W_t^2 – Wiener processes with correlation ρ

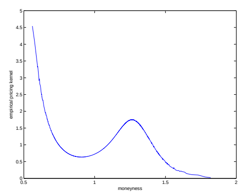

The parameters in the Heston model can be interpreted as:

- ξ – mean-reversion speed
- η – long-term variance
- V_0 – short-term variance
- ρ – correlation
- θ – volatility of volatility

η and V_0 control the term structure of the implied volatility surface (i.e. time to maturity direction).

ρ and θ control the smile/skew (i.e. moneyness direction).

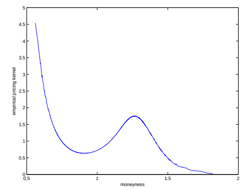

The state-price density is derived from European option prices that may be represented in an implied volatility surface:

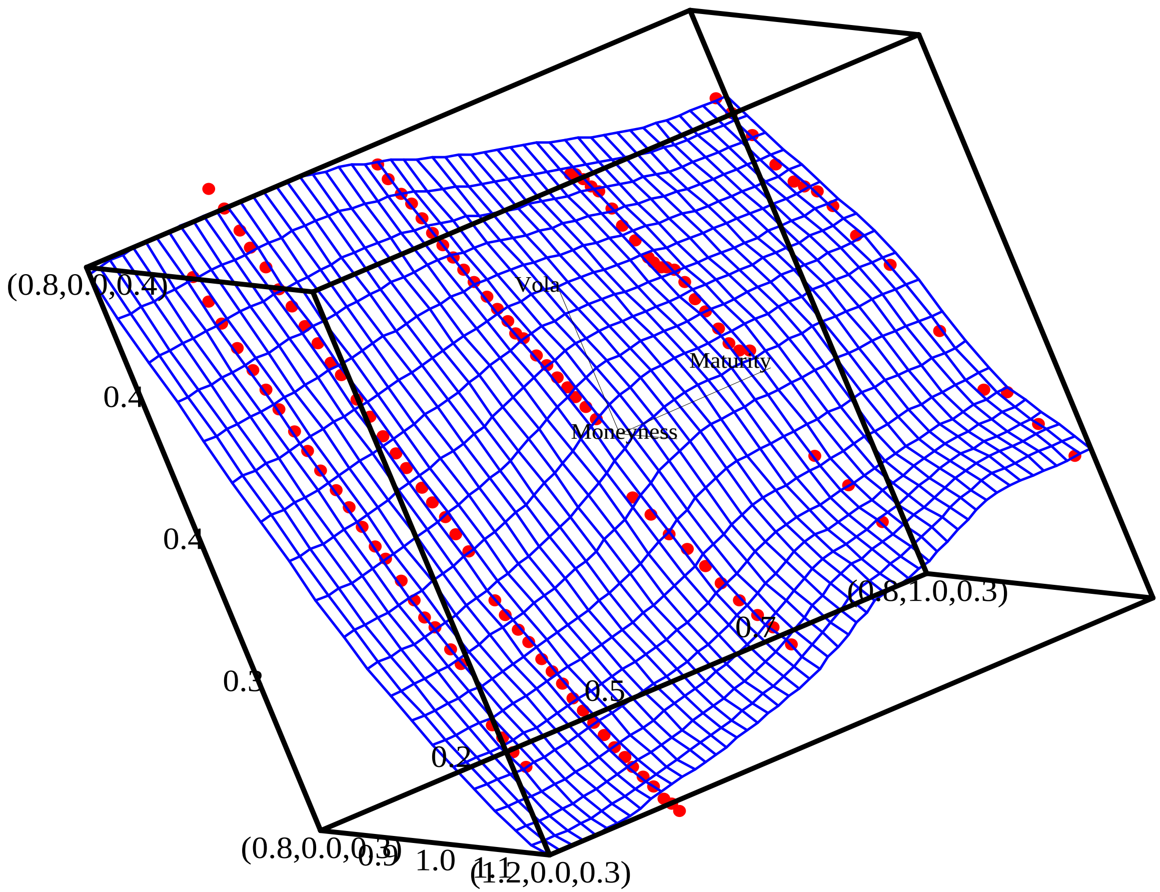

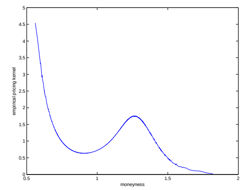

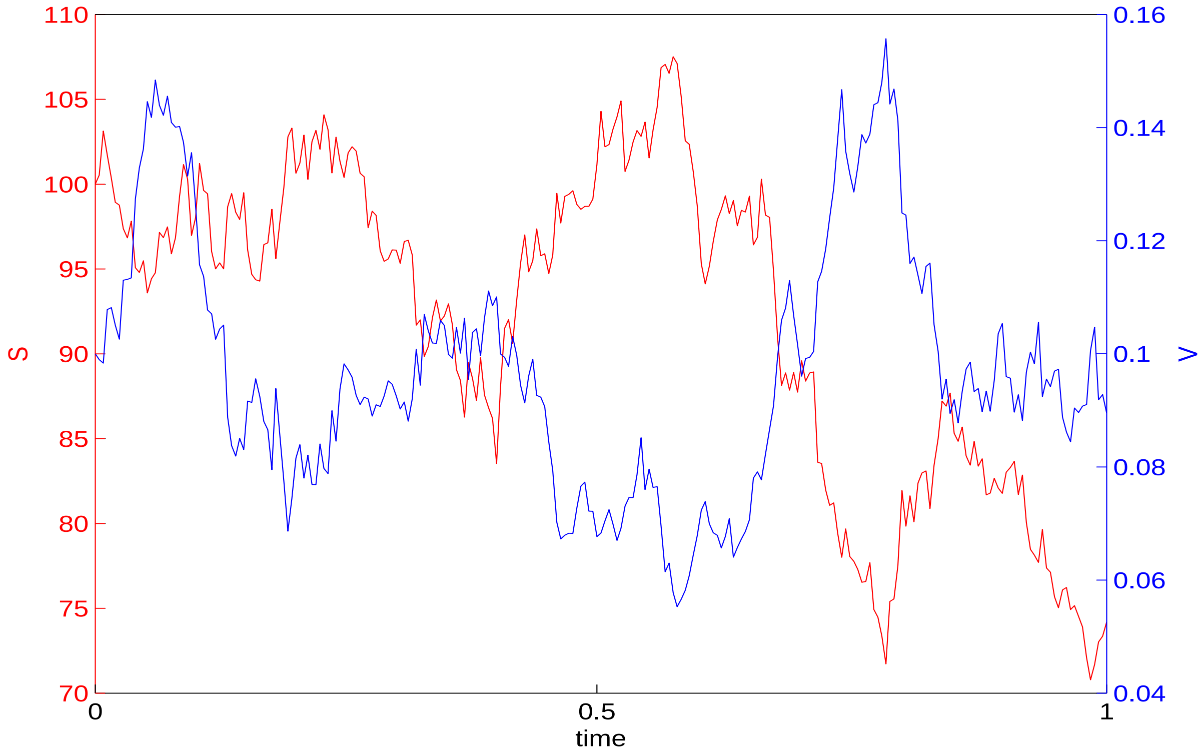

Figure 5: Simulated paths in the Heston model for the parameters $V_0 = 0.1$, $\eta = 0.08$, $\xi = 2$, $\theta = 0.3$, $\rho = -0.7$. S – stock process, V – variance process.

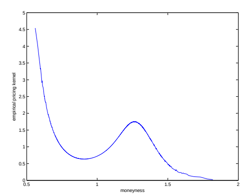

We estimate the parameters of the state-price (objective) density by minimising the MSE of the implied volatilities:

$$\frac{1}{n}\sum_{i=1}^{n}(IV_i^{model} - IV_i^{market})^2$$

where IV^{model} and IV^{market} refer to model and market implied volatilities; n is the number of observations on the surface.

Typically, we observe prices of options with the time to maturity $\tau \in [0.25; 1]$ years and moneyness $K/S_0 \in [0.5; 1.5]$.

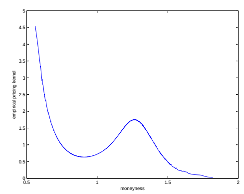

The plain vanilla prices are calculated by a method of Carr and Madan:

$$C(K,T) = \frac{\exp\{-\alpha \ln(K)\}}{\pi} \int_0^{+\infty} \exp\{-\mathbf{i}v \ln(K)\} \psi_T(v) dv$$

for a damping factor $\alpha > 0$. The function ψ_T is given by

$$\psi_T(v) = \frac{\exp(-rT)\phi_T\{v - (\alpha+1)\mathbf{i}\}}{\alpha^2 + \alpha - v^2 + \mathbf{i}(2\alpha+1)v}$$

where ϕ_T is the characteristic function of $\log(S_T)$.

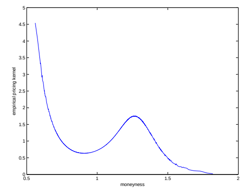

Estimation of the Subjective Density p_t

The logarithmic returns r_t of DAX are modelled with the GARCH(1,1) model:

$$r_t = \sigma_t Z_t$$

$$\sigma_t^2 = \omega + \alpha r_{t-1}^2 + \beta \sigma_{t-1}^2$$

From the logarithmic returns $R_i = \log(S_i) - \log(S_{i-1}),\ i = 1, \ldots, t$ and the starting stock price S_0 we can construct the final stock price as:

$$S_t = S_0 \exp(\sum_{i=1}^{t} r_i).$$

The model is fitted by maximising the likelihood function

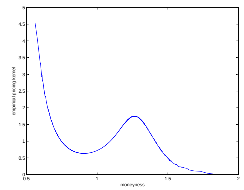

We estimate the subjective density p in a forward rolling time window of the length of two years:

- Fit the GARCH(1,1) model for DAX returns
- Simulate N time series of the returns (N=5000)
- Compute the final N DAX prices
- Evaluate $\hat{p}$ using kernel density estimation with the Gaussian kernel

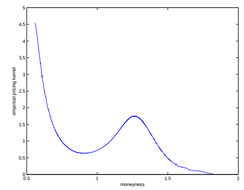

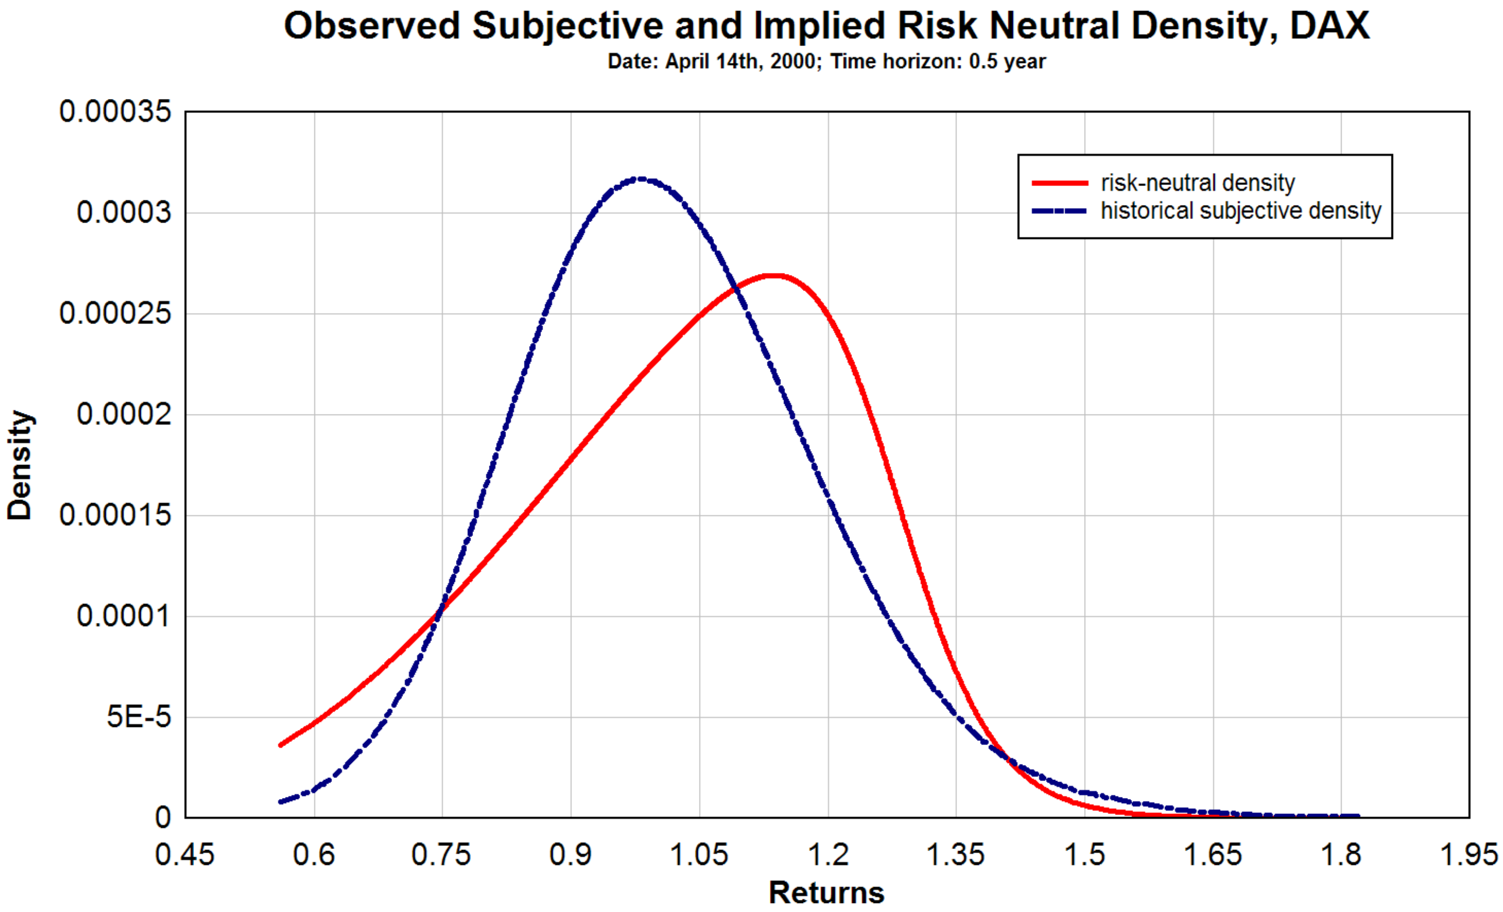

Figure 6: Estimated price densities for $\tau = 0.5$ year, April 14th, 2000.

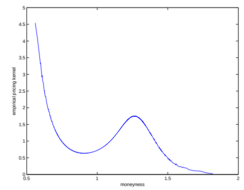

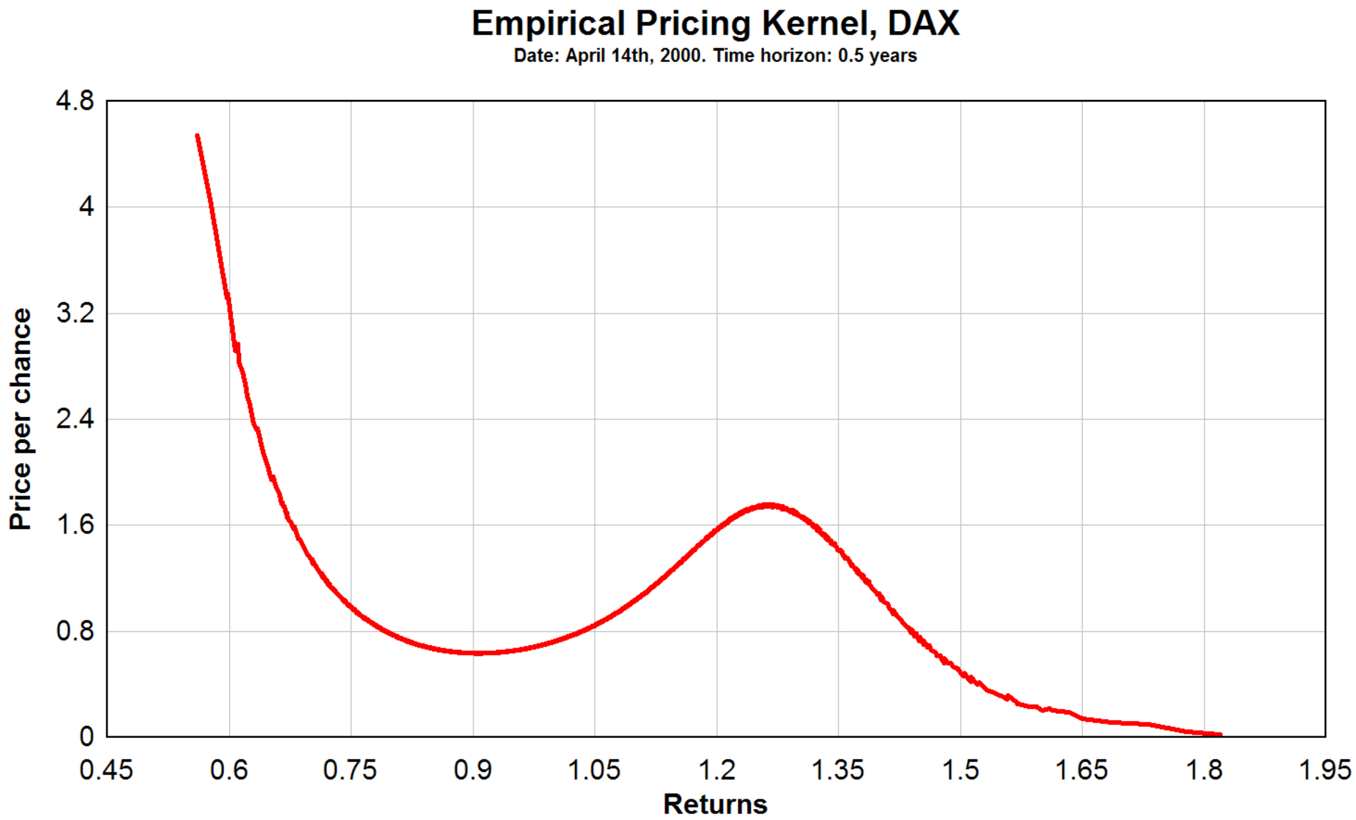

Figure 7: Estimated pricing kernel for $\tau = 0.5$ year, $r_{0.5} = 4.06\%$, April 14th, 2000.

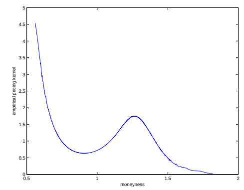

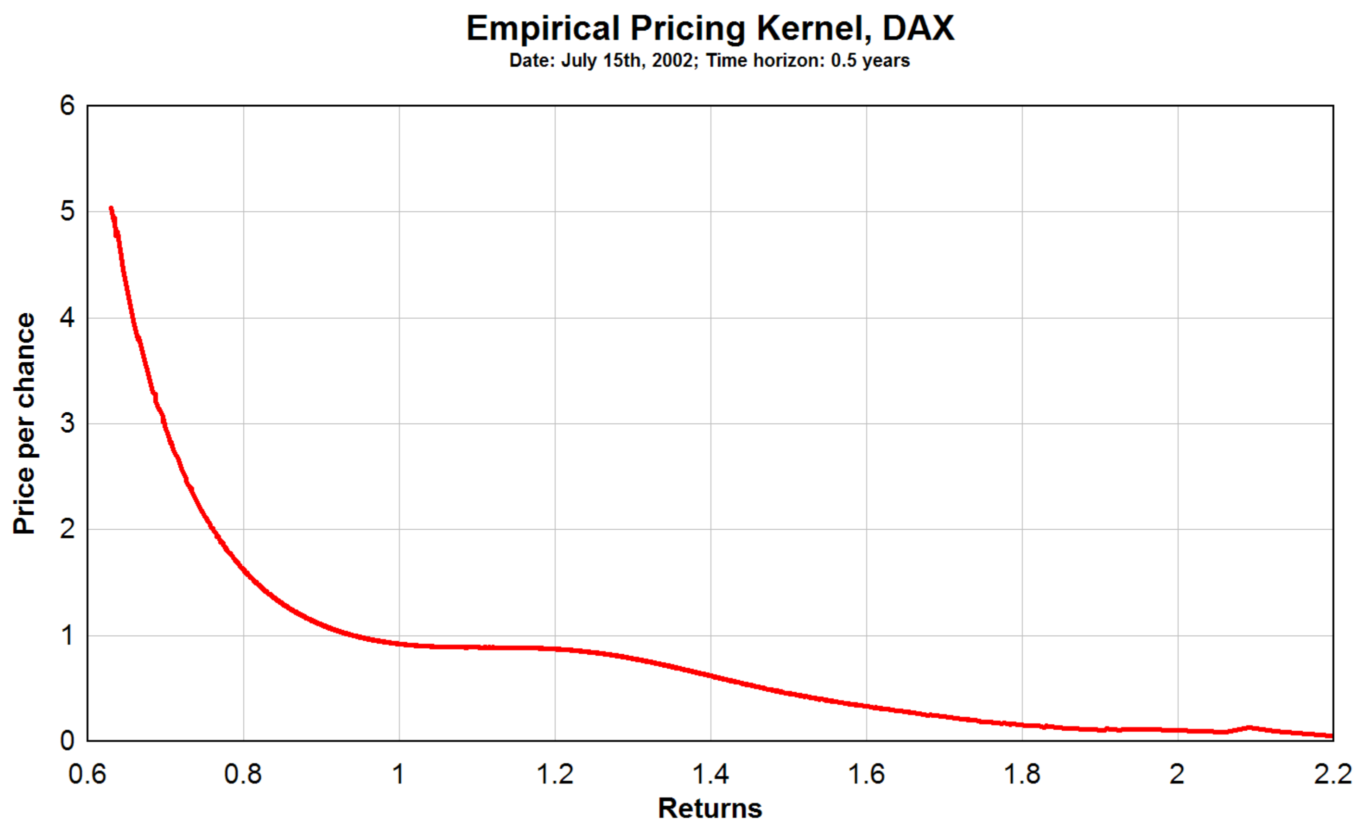

Figure 8: Estimated pricing kernel for $\tau = 0.5$ year, $r_{0.5} = 3.50\%$, July 15th, 2002.

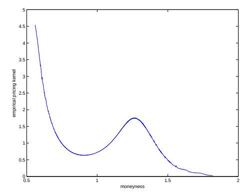

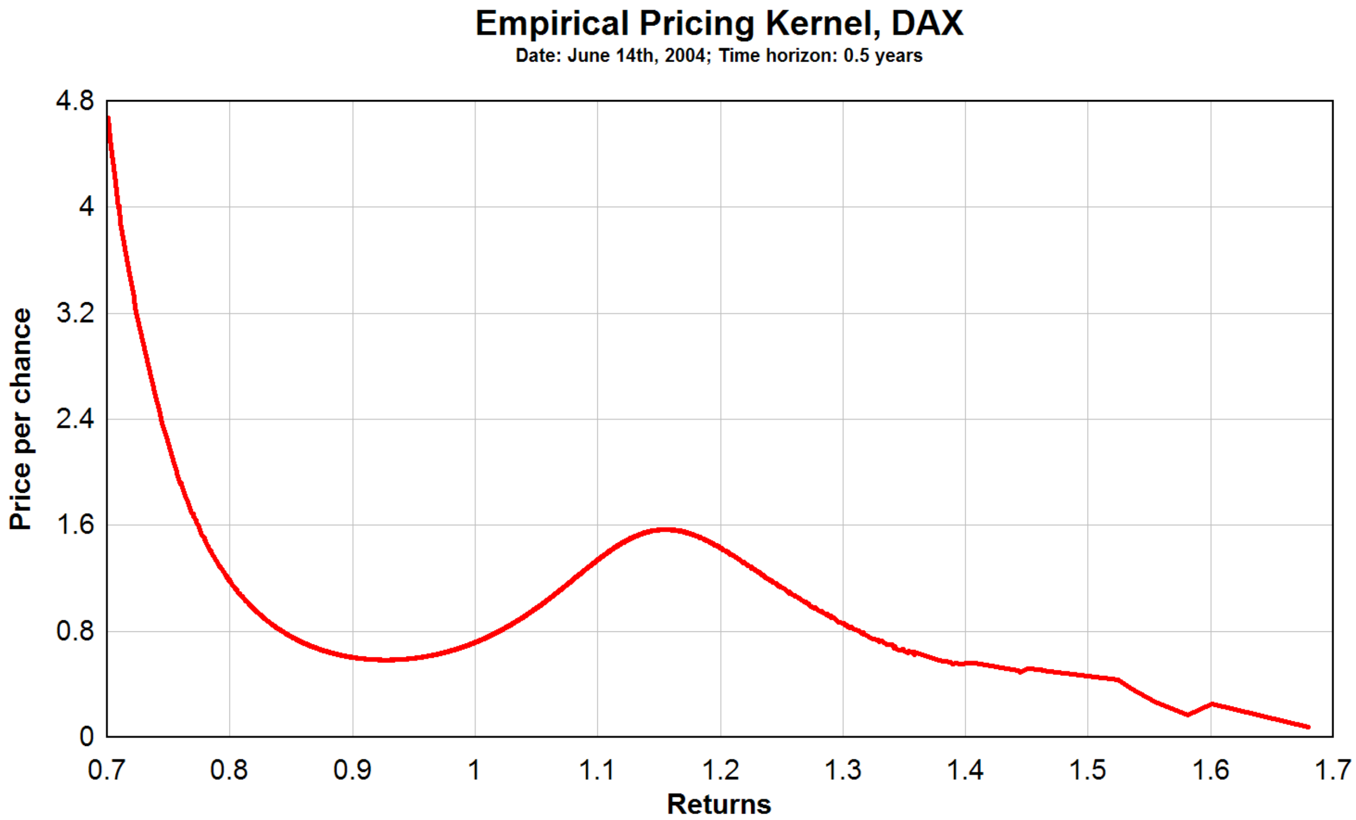

Figure 9: Estimated pricing kernel for $\tau = 0.5$ year, $r_{0.5} = 2.23\%$, June 14th, 2004.

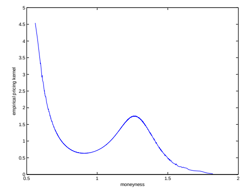

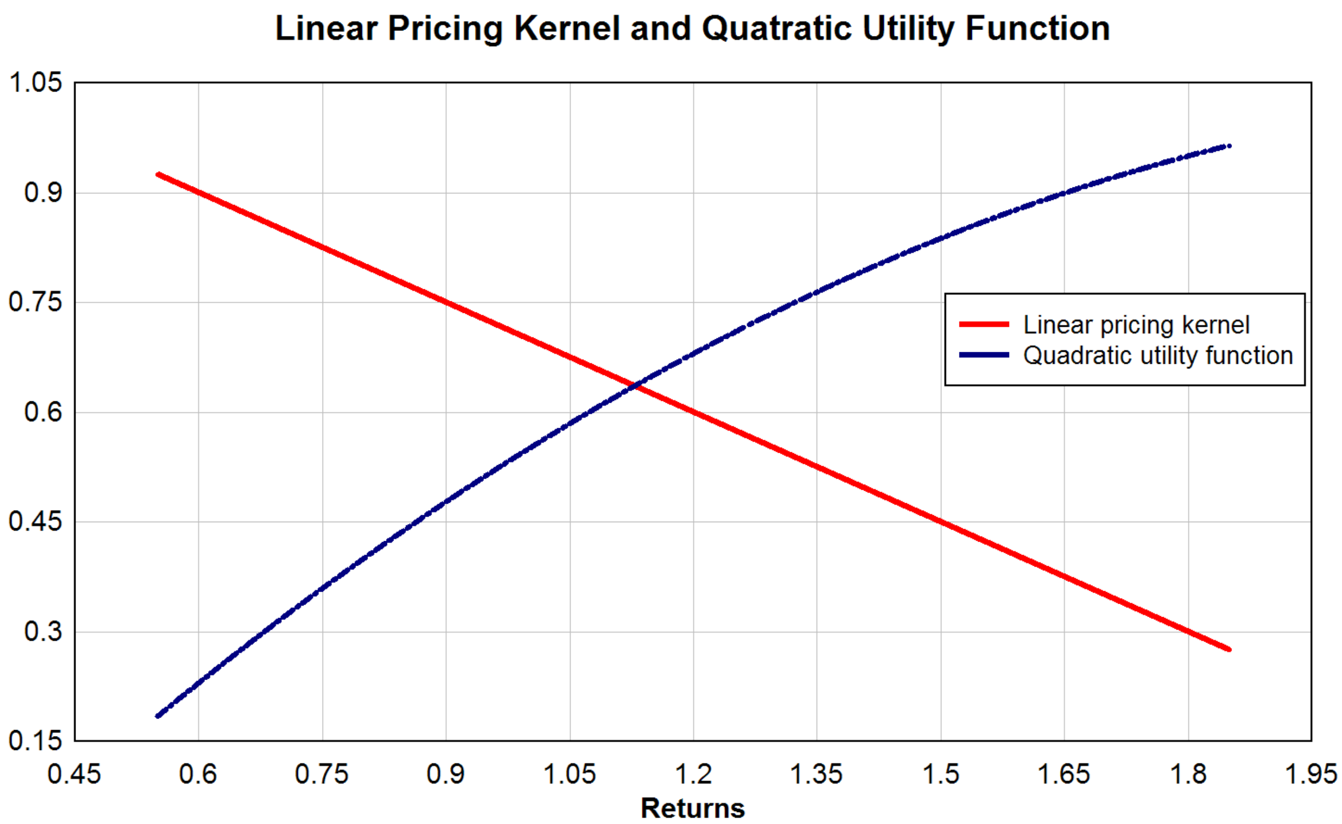

Figure 10: Linear pricing kernel and quadratic utility function (CAPM model). $U(X_{t+1}) = -aX_{t+1}^2 + bX_{t+1} + c$.

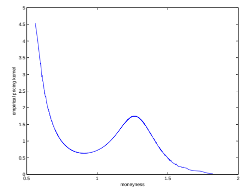

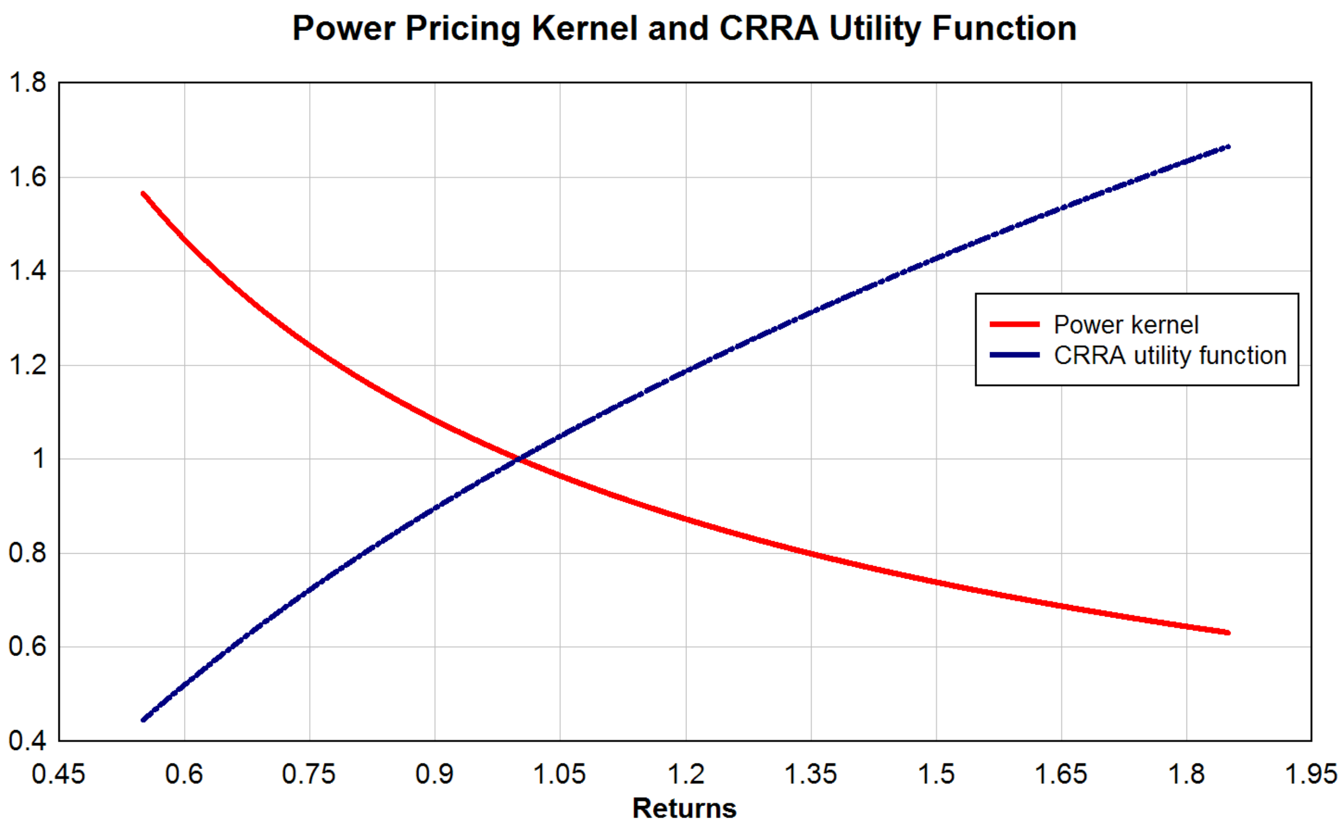

Figure 11: Power pricing kernel and CRRA utility function. $U(X_{t+1}) = a\frac{X_{t+1}^{1-\gamma}}{1-\gamma}$.

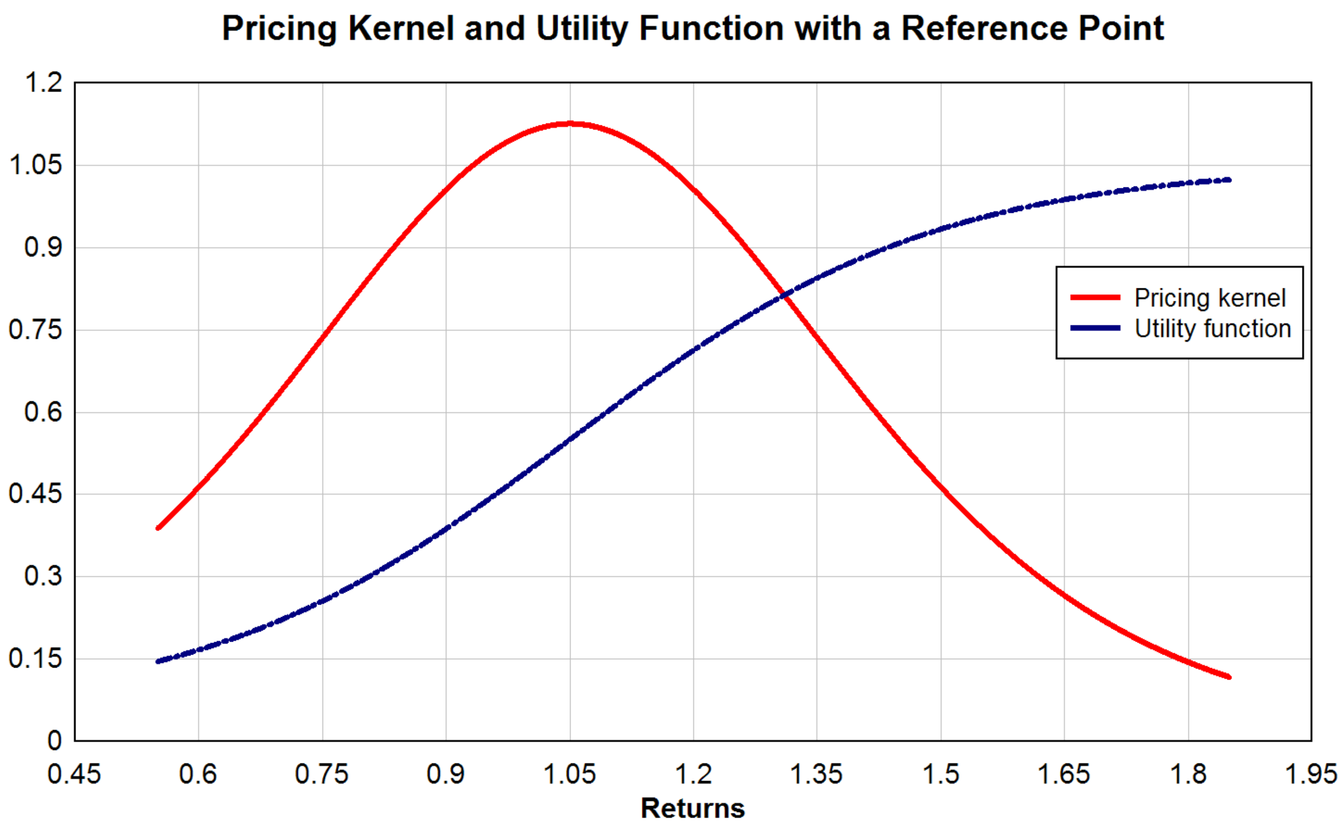

Figure 12: Pricing kernel and utility function suggested by Kahneman and Tversky based on behavioural experiments.

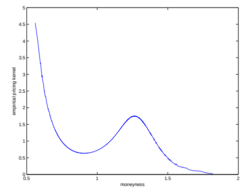

Estimation of the Market Utility Function

Utility function derived from the market data is the **market utility function**. It requires the assumption about the existence of a **representative investor**

$$m_t(X_{t+1}) = \text{const}_t \cdot U'(X_{t+1}) \tag{4}$$

Since a cardinal utility function can be defined up to a linear transformation, the constant can be neglected

$$U_t(X_{t+1}) = \int_{\inf(X_t)}^{X_{t+1}} m_t(s) ds$$

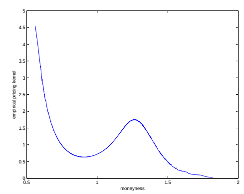

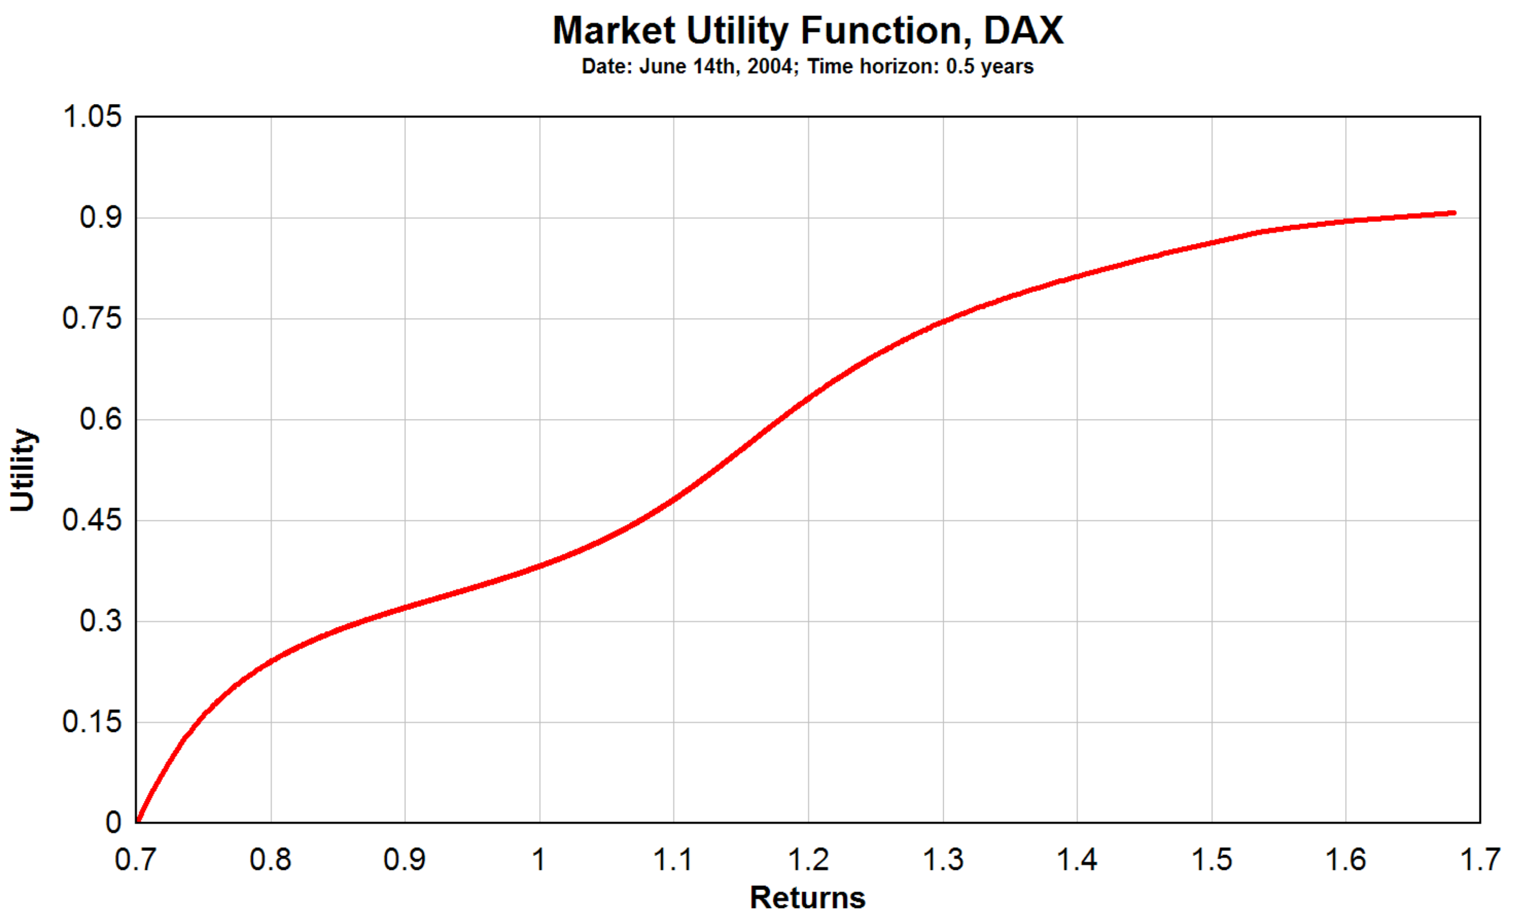

Figure 13: Market utility function, DAX, $\tau = 0.5$ years, June 14th, 2004.

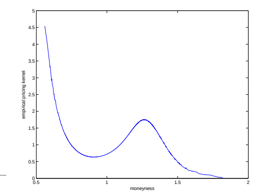

Decomposition of the Utility Function

Observation: the portions of the utility function below $X_{t+1} = 1$ and above $X_{t+1} = 1.15$ are very well approximated with shifted CRRA functions, $k = 1, 2$:

$$U_t^{(k)}(X_{t+1}) = a_k \frac{(X_{t+1} - c_k)^{\gamma_k - 1}}{\gamma_k - 1} + b_k,$$

where the shift parameter is c_k. The CRRA function becomes infinitely negative for $X_{t+1} = c_k$ and is extended as $U_t^{(k)}(X_{t+1}) = -\infty$ for $X_{t+1} < c_k$, i.e. investors by all means will avoid the situation when $X_{t+1} < c_k$. For a standard CRRA utility function $c_k = 0$.

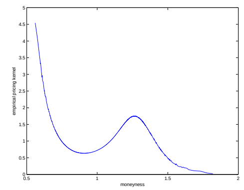

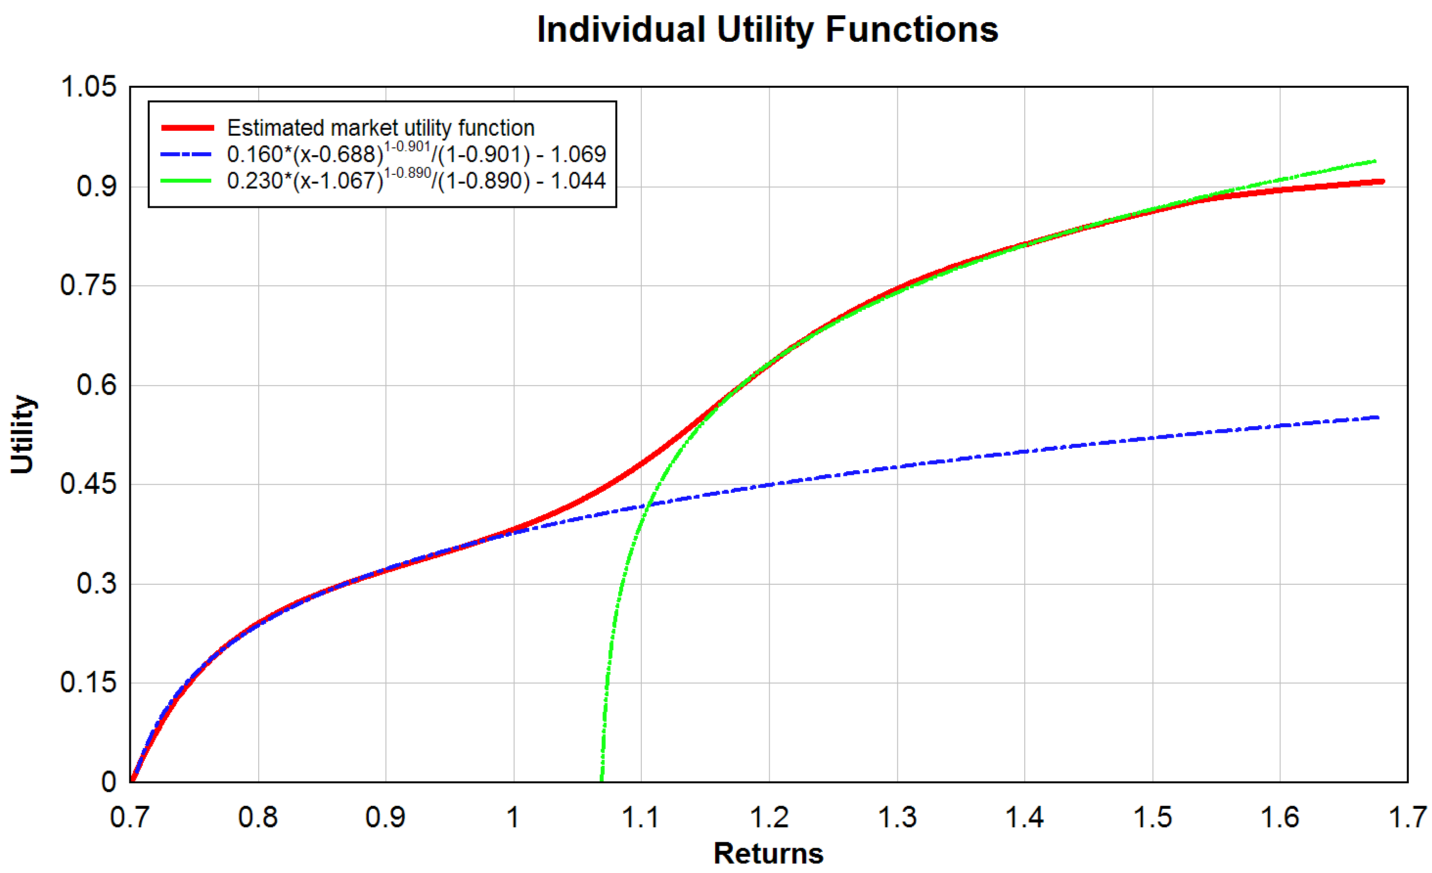

Figure 14: Decomposition of the utility function. DAX, $\tau = 0.5$ years, June 14th, 2004.

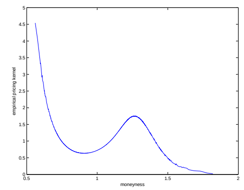

Individual Utility Functions

We abandon the hypothesis of the representative investor: there are **many investors in the market**.

Investor i has a utility function that consists of two CRRA functions:

$$U_{i,t}(X_{t+1}) = \begin{cases} \max\left\{U_t(X_{t+1},\theta_1,c_1);U_t(X_{t+1},\theta_2,c_{2,i})\right\}, & \text{if} \quad \mathrm{X_{t+1} > c_1} \\ -\infty, \qquad \text{if} \quad \mathrm{X_{t+1} \leq c_1} \end{cases}$$

where $U_t(X_{t+1},\theta,c) = a\frac{(X_{t+1}-c)^{\gamma-1}}{\gamma-1} + b$, $\theta = (a,b,\gamma)^\top$, $c_{2,i} > c_1$. If $a_1 = a_2 = 1$, $b_1 = b_2 = 0$ and $c_1 = c_2 = 0$, we get the standard CRRA utility function.

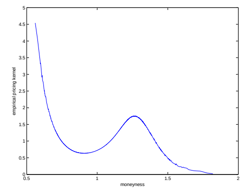

Parameters θ_1 and θ_2 and c_1 are the same for all investors. Investors differ with the shift parameter c_2.

θ_1 and c_1 are estimated on the lower 20% of observations, when, assumingly, all investors agree that the market is "bad" ("bear" market).

θ_2 is estimated on the upper 20% of observations, when all investors agree that the state of the world is "good" ("bull" market).

The distribution of c_2 that uniquely defines the distribution of switching points is computed with a "boosting" procedure.

	a_i	b_i	γ_i
$i = 1$ (bear market)	0.160	-1.069	0.901
$i = 2$ (bull market)	0.230	-1.044	0.890

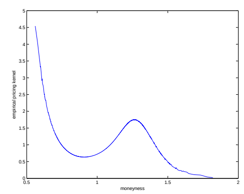

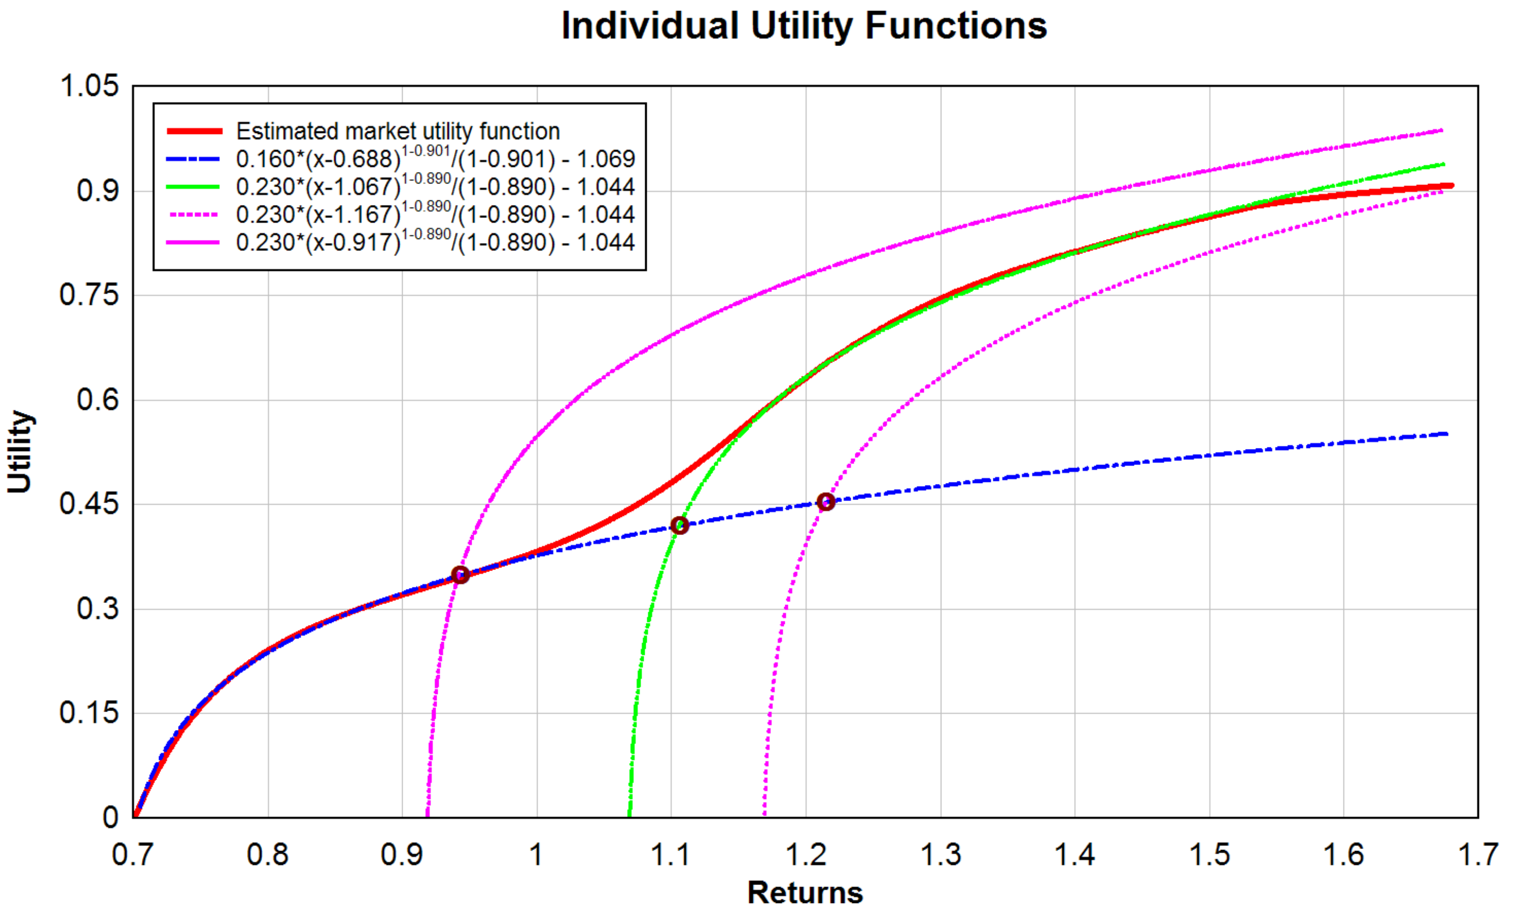

Figure 15: Individual utility functions, DAX, $\tau = 0.5$ years, June 14th, 2004

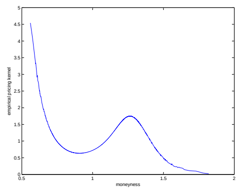

Investor Types

- A change of behaviour from bearish to bullish happens at a switching point
- Different investors have different perceptional outlooks concerning the future state of economy, i.e. have different boundary between "good" and "bad" states
- Most of investors have switching points in the interval $[0.95; 1.1]$, i.e. in the area that corresponds to present unit returns times half-year risk free interest rates

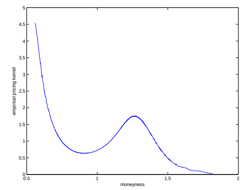

The individual utility function can be conveniently denoted as:

$$U_i(X_{t+1}) = \begin{cases} \max\left\{U_{bear}(X_{t+1}); U_{bull}(X_{t+1}, c_i)\right\}, & \text{if} \quad \mathrm{X_{t+1} > c_1} \\ -\infty, \qquad \text{if} \quad \mathrm{X_{t+1} \leq c_1} \end{cases}$$

Switching between U_{bear} and U_{bull} happens at the *switching point* Z_{t+1}, where $U_{bear}(Z_{t+1}) = U_{bull}(Z_{t+1}, c_i)$. The switching point is determined by $c_i \equiv c_{2,i}$

The notations *bear* and *bull* have been chosen because U_{bear} is activated when returns are low ("bear" market) and U_{bull} when returns are high ("bull" market)

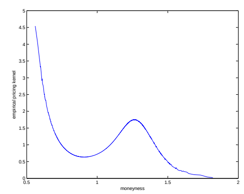

Market Conditions and the Switching Point

Each investor is characterised with a switching point Z_{t+1}

The smoothness of the market utility function is the result of the aggregation of different attitudes

U_{bear} characterises more cautious attitudes when returns are low
U_{bull} describes the attitudes when the market is booming

Both U_{bear} and U_{bull} are concave. However, due to switching the total utility function can be locally convex

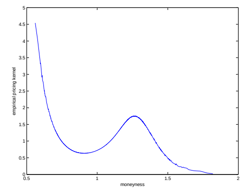

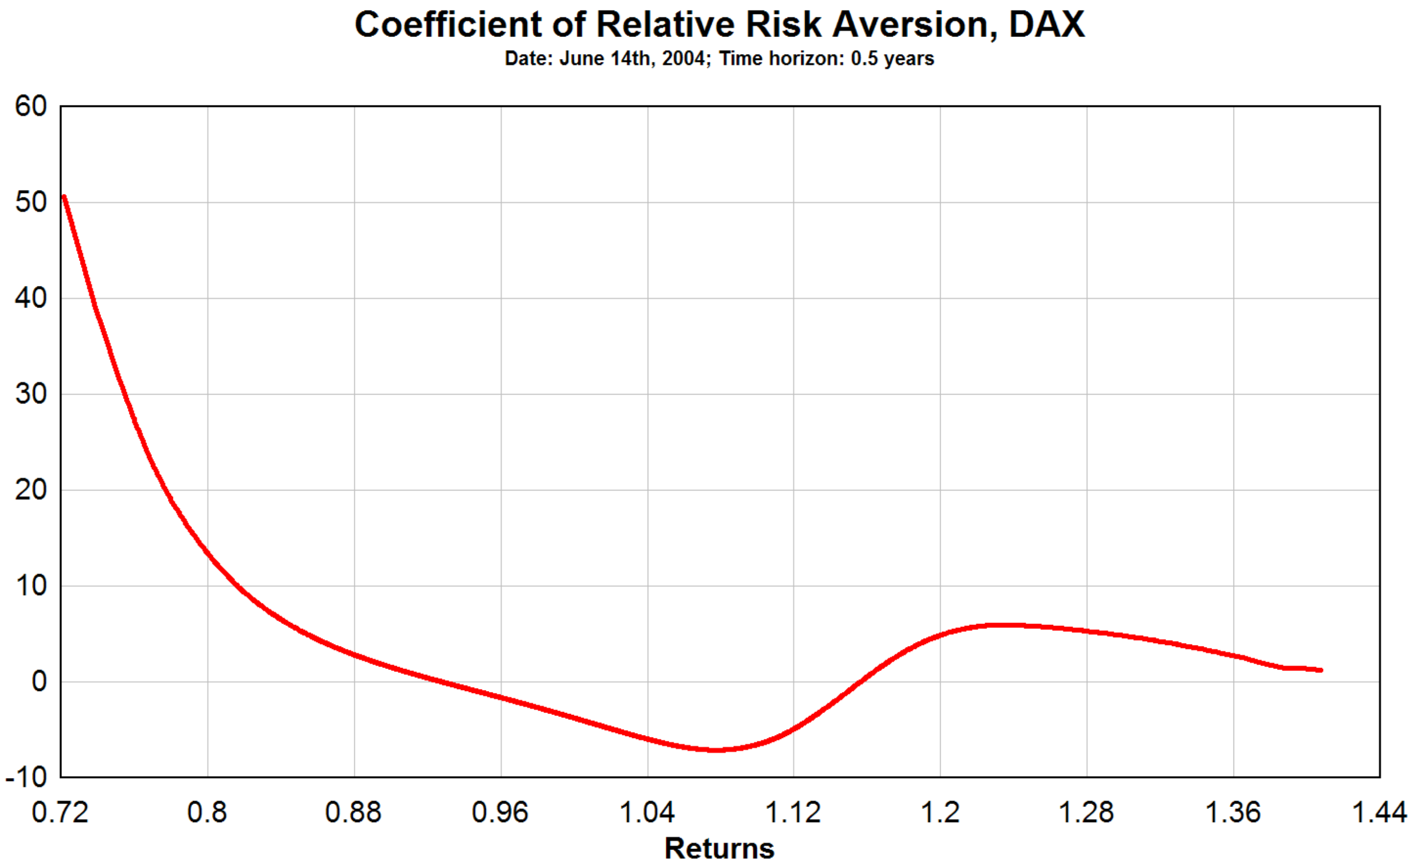

Figure 16: Market Relative Risk Aversion Coefficient, DAX, $\tau = 0.5$ years, June 14th, 2004

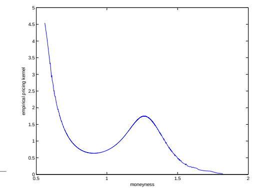

The coefficient of relative risk aversion is:

$$a_R(X_{t+1}) = -\frac{U''(X_{t+1})X_{t+1}}{U'(X_{t+1})}.$$

We compute it non-parametrically from the estimated pricing kernel, which equals $const_t \cdot U'(X_{t+1})$

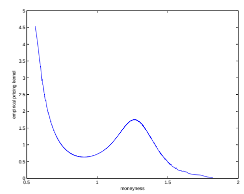

Naive Utility Aggregation

- Specify the **observable** states of the world in the future by returns X_{t+1}
- Find a weighted average of the utility functions for each state. If the importance of the investors is the same, then the weights are equal
- **Problem**: utility functions of different investors cannot be summed up since they are incomparable

$$U_t(X_{t+1}) = \frac{1}{N} \sum_{i=1}^{N} U_t^{(i)}(X_{t+1})$$

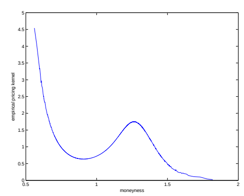

Investor's Attitude Aggregation

- Specify **perceived** states of the world given by utility levels $\tilde{u}$
- Aggregate the outlooks concerning the **returns** in the future X_{t+1} for each perceived state

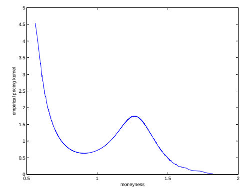

For a **subjective** state described with the utility level $\tilde{u}$, such that

$$\tilde{u} = U^{(1)}(X_{t+1}^{(1)}) = U^{(2)}(X_{t+1}^{(2)}) = \ldots = U^{(N)}(X_{t+1}^{(N)})$$

the aggregate estimate of the resulting return is

$$X_{t+1}^{A}(\tilde{u}) = \frac{1}{N} \sum_{i=1}^{N} X_{t+1}^{(i)}(\tilde{u})$$

if all investors have the same market power.

N is the number of investors

Important property: the return aggregation procedue is invariant of *any* monotonic transformation

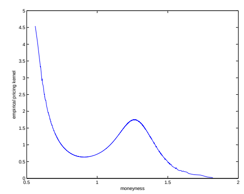

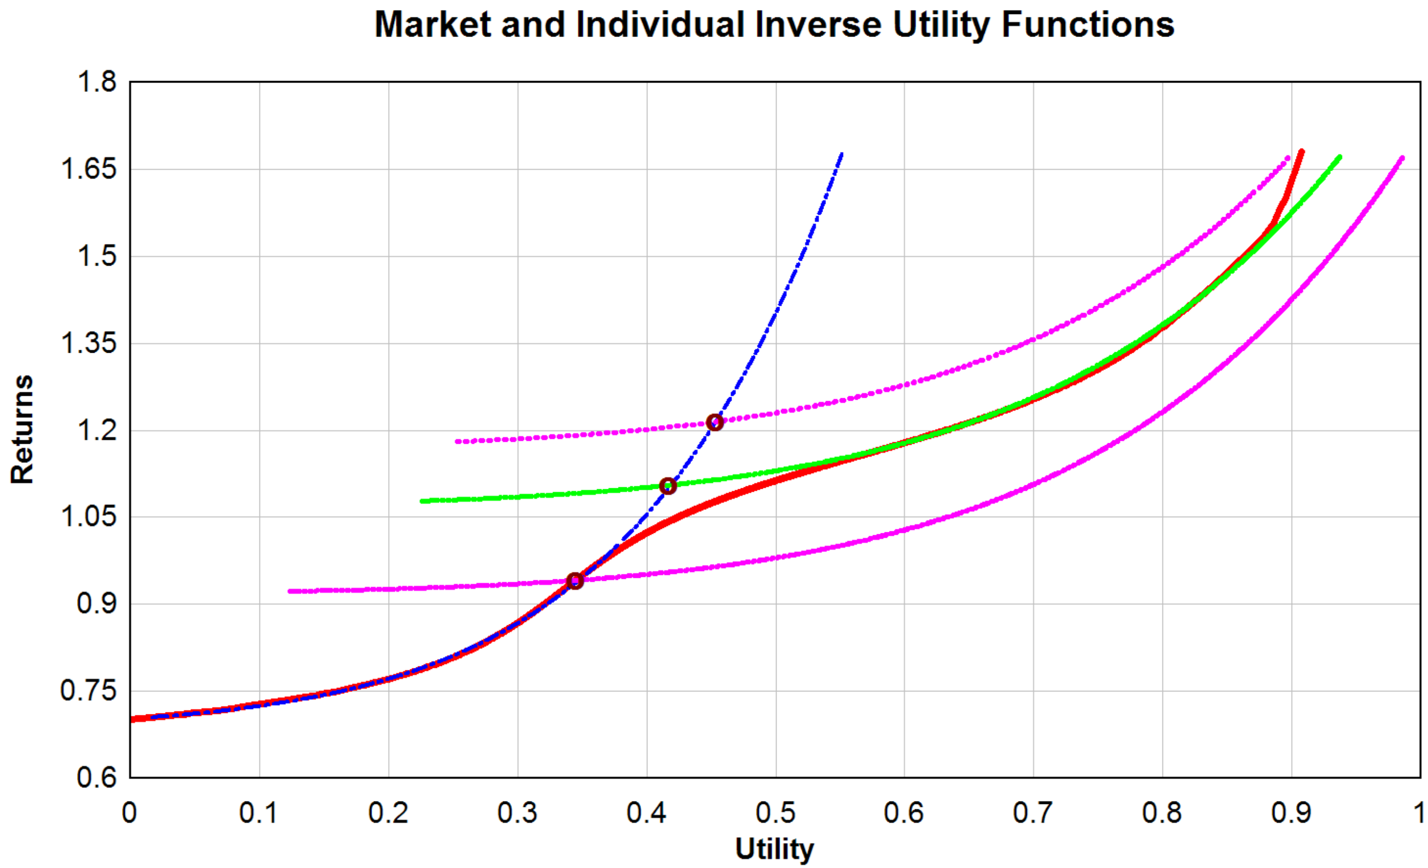

Figure 17: Inverse market and individual utility functions, DAX, $\tau = 0.5$ years, June 14th, 2004

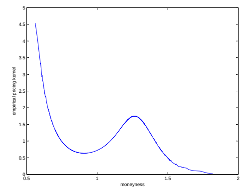

Estimating the Distribution of Switching Points with a Boosting Algorithm

1. Generate N realisations of individual utility functions with switching points $Z^{(i)}$, $i = 1, \dots, N$ with any prior distribution with a compact support

2. Add one individual utility function $U^{(i)}$ and delete another $U^{(j)}$ with random switching points $Z^{(i)}$ and $Z^{(j)}$ respectively

3. Aggregate individual utility functions using subjective state aggregation. If the proximity to estimated market utility function has increased, retain the new swithching point, otherwise do nothing

4. Repeat steps 2 and 3 until the estimated market and fitted utility functions become close

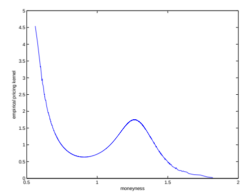

The aggregate return in the *perceptional* state $\tilde{u}$ is given by:

$$X_f^A(\tilde{u}) = \frac{1}{N} \sum_{i=1}^{N} U_{Z_i}^{-1}(\tilde{u})$$

where $U_{Z_i}^{-1}(\tilde{u})$ is the inverse individual utility function or:

$$X_f^A(\tilde{u}) = \int U_Z^{-1}(\tilde{u}) f(Z) dZ$$

where $f(Z)$ is the distribution of switching points, which is derived as the solution of the minimisation problem:

$$\min_{f(Z)} \int \left\{ U_M^{-1}(\tilde{u}) - X_f^A(\tilde{u}) \right\}^2 dP(\tilde{u}),$$

where $U_M^{-1}(\tilde{u})$ is the inverse of the estimated market utility function.

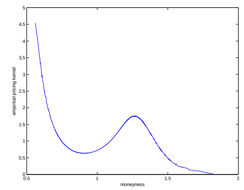

Distribution of Switching Points

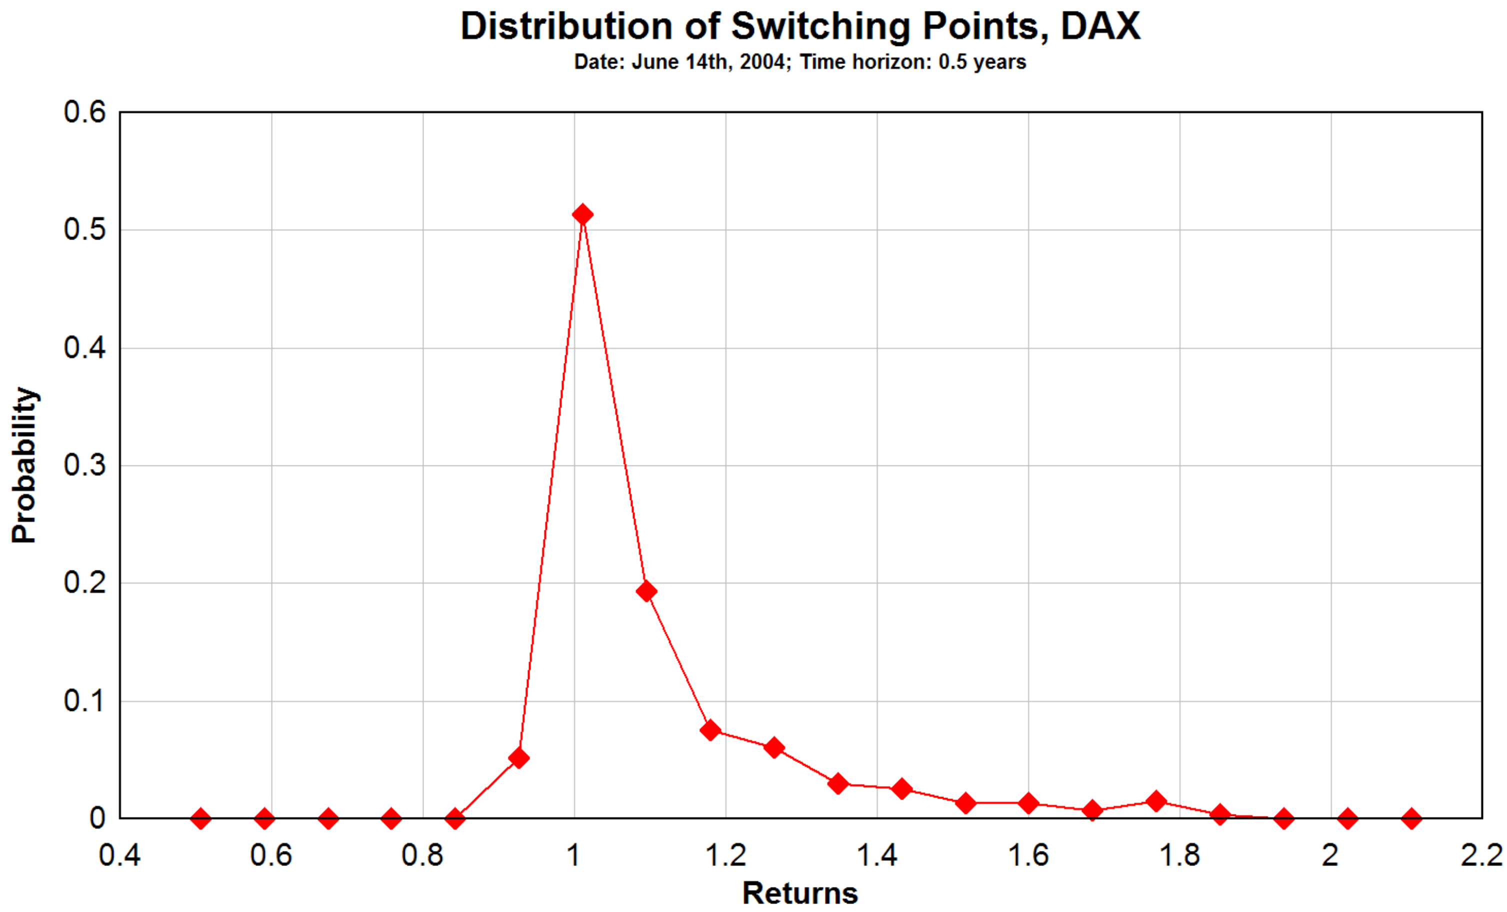

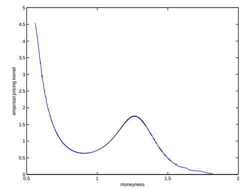

Behavioural Experiment Design

- There are several states of the world ranging from "bad" (low returns) to "good" (high returns)
- There are three groups of participants that are told that the world is more likely to be in the "bad", "good" or approximately the same state in the future, respectively. In this way we expect participants in the three groups to operate with U_{bear}, U_{bull} or in the switching regime
- Each participant is asked to place 100 markers denoting desired outcomes into the future states, thus building a subjective distribution

- The prices of putting a marker into a state are given by a risk-neutral distribution estimated from real option market data with the Heston model. Several other distributions of state-prices, such as the log-normal distribution, can also be tested

- Each participant has an endowment of 100 EUR that he must completely spend building the distribution with markers. In this way the budget constraint is implemented

- The pricing kernel is computed as the ratio of the risk-neutral density and experimentally derived subjective density times the discount factor, i.e.

$$\hat{m}_t(X_{t+1}) = \beta \frac{\hat{q}_{risk-neutral}(X_{t+1})}{\hat{p}_{experimental}(X_{t+1})}$$

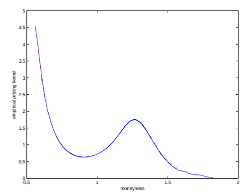

Distribution Builder (Sharpe, 2006)

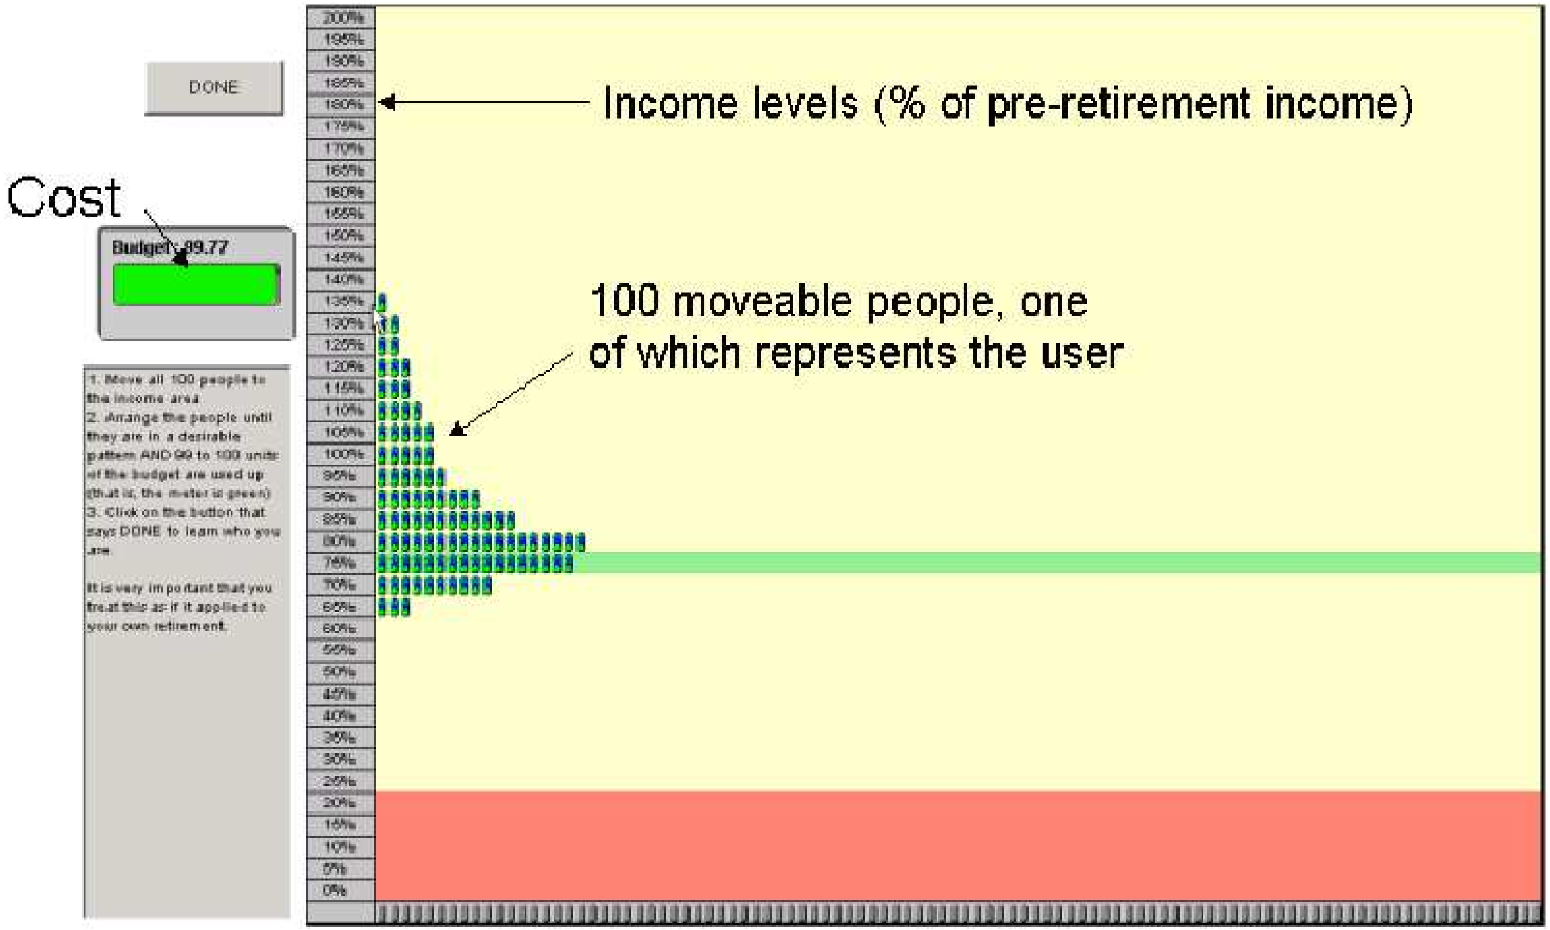

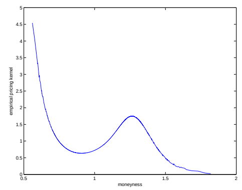

Claims

- Representation of individual utility functions as consisting of two parts, activated during perceptionally "good" and "bad" states of the world. The perceptional change happens at the swithing point. Investors behave as risk averse individuals in "good" and "bad" states but become risk seeking when switching occurs
- Utility function aggregation procedure based on subjective states of the world
- Use of DAX data and the Heston model to estimate the market pricing kernel
- Introduction of a "boosting" procedure for the estimation of the distribution of switching points

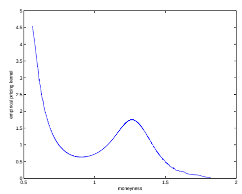

Outlooks

- Extension of the experiment with a trading simulator, so that prices are determined by the participants
- Testing alternative utility function designs
- Refining the technique for estimating the distribution of switching points as an inverse problem
- Study of the dynamics of pricing kernels and individual utility functions

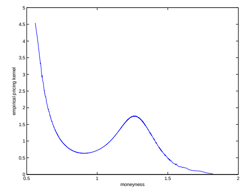

References

Sharpe, W. F. (2006) Investors and Markets: Portfolio Choices, Asset Prices and Investment Advice, Princeton University Press, Princeton, NJ (in print)

Rosenberg, J. V. and Engle, R. F. (2002) Empirical Pricing Kernels, *Journal of Financial Economics*, **64**, June, 341-372.

Jackwerth, J. C. (2002) Recovering Risk Aversion from Option Prices and Realized Returns, *Review of Financial Studies*, **13**, 433-451.

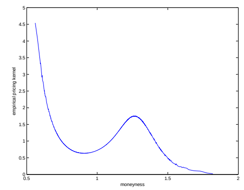

Cochrane, J. H. (2001) Asset Pricing, Princeton University Press, Princeton, NJ.

Aït-Sahalia, Y. and Lo, A. W. (2000) Nonparametric Risk Management and Implied Risk Aversion, *Journal of Econometrics*, **94**, 9-51.

Kahneman, D and Tversky, A. (1979) Prospect Theory: An Analysis of Decision under Risk, *Econometrica*, **47**, March, 263-291.