ANNALS OF ECONOMICS AND FINANCE **14-2(A)**, 361–388 (2013)

# Forecasting Inflation via Electronic Markets

Michael Berlemann

*Helmut-Schmidt-University Hamburg*
*Chair for Political Economy & Empirical Economics, Hamburg, Germany*
E-mail: Michael.Berlemann@hsu-hh.de

and

Forrest D. Nelson

*University of Iowa, Tippie College of Business. Iowa City, IA 52242-1994, U.S.A*
E-mail: forrest-nelson@uiowa.edu

While there are various techniques of inflation forecasting in use, none of them has proved to deliver consistently more accurate forecasts than the others. That is why most users of inflation forecasts monitor a variety of inflation indicators and forecasts and check them for consistency. This paper aims at contributing to an extension of the methods in use. We show that experimental inflation forecasting markets are highly useful in uncovering market participants' inflation expectations. We also present evidence from a number of pilot experiments underlining that the proposed method might enrich the arsenal of existing forecasting techniques.

*Key Words*: Inflation forecast; Field experiments; Experimental stock markets.
*JEL Classification Numbers*: E58, F15, F33.

## 1. INTRODUCTION

Since the 1990s, interest in forecasting macroeconomic variables, and especially inflation, has increased considerably. This is at least partially due to the fact that many central banks switched to inflation targeting strategies of monetary policy. An essential part of these strategies is the usage of inflation forecasts as intermediate target of monetary policy decisions.[1] However, due to the fact that expectations about inflation are embedded in planning decisions of all kinds, macroeconomic and inflation forecasting is also an important matter outside central banks, e.g. when corporations

[1]See Svensson (1997,1999).



1529-7373/2013

and workers (or trade unions) negotiate wages or the public sector plans budgets.

The aim of this paper is to contribute to extending the conventional set of inflation forecasting techniques. We propose a new method of forecasting inflation: assessing market inflation expectations via conducting experimental stock markets. Our belief that experimental stock markets could be a useful tool is driven by the experience that experimental political stock markets have been quite successful in predicting electoral outcomes.[2] Thus, transferring the idea to forecasting macroeconomic variables like e.g. inflation might be fruitful. We therefore develop a design for inflation forecasting markets in this paper. We also report and analyze the results from a number of related pilot experiments conducted in Germany.

The outline of this paper is as follows: the second section gives a brief review of conventional forecasting techniques. In the third section we develop the design of a prototype experimental inflation forecasting market and show how the market data can be used to construct both mean and density forecasts of inflation. We also show how the uncertainty surrounding the mean forecast can be assessed. In the fourth section we report and analyze the results of 4 inflation forecasting markets conducted in Germany. Section 5 summarizes the main results.

## 2. CONVENTIONAL INFLATION FORECASTING TECHNIQUES

Conventional inflation forecasting techniques can roughly subdivided into two basic approaches: the expectations and the econometric approach.[3]

Econometric inflation forecasting models[4] use historical macroeconomic data to generate forecasts. In forward-looking models even expectations on future values of some variables can enter econometric models. Econometric forecasting models can be either theory-dominated (as e.g. Phillips curve models[5], P-star models[6] or large-scale macroeconometric models as they are used by many central banks[7]) or atheoretic (as e.g. VAR models[8]).

[2] See Berg et al. (1998), Berg, Forsythe and Rietz (1997) or Berlemann and Schmidt (2001).

[3] For a brief introduction to inflation forecasting see Tallman (1995).

[4] For a detailed description of alternative econometric inflation forecasting techniques see Bank of England (1999).

[5] See e.g. Stock and Watson (1999), Atkeson and Ohanian (2001) or Fisher, Liu and Zhou (2002).

[6] See Hallman, Porter and Small (1991), Tödter and Reimers (1994), Issing and Tödter (1995) and Gottschalk and Bröck (2000).

[7] See e.g. Bank of England (1999), Jordan and Peytrignet (2001) or Poloz, Rose and Tetlow (1994).

[8] See Doan, Litterman and Sims (1984) or Thompson and Miller (1986).

The basic idea of the expectations approach of forecasting inflation is that many people care - or at least should care - about inflation. The various forms of inflation effects[9] deliver the theoretical reasoning for this argument. The expectations approach simply suggests that people, caring about future inflation, know enough about the true determinants of price level changes to be able to predict future inflation well on average. Thus, it is sufficient to measure inflation expectations provided that individual forecasts turn out to be rational.

However, the basic problem of the expectations approaches is how to uncover market participants' inflation expectations. Direct methods of measuring inflation expectations typically rely on some sort of expectation survey in which certain subsamples of the population are asked to reveal their personal inflation expectations.[10] While in some surveys the respondents have to make some qualitative assessment of future inflation, others ask for concrete numbers. However, in both cases some appropriate way of aggregating the individual responses has to be found. Moreover, this method suffers from all problems which are extensively discussed in the literature on surveying techniques, as e.g. sampling, non-response problems, motivation of respondents etc.[11] The indirect approach to measure expectations is to derive inflation expectations from market participants' behavior on real world (financial) markets. A straightforward way to do so is to use prices of CPI futures to derive market expectations.[12] However, in most countries markets for these futures did not develop. Alternatively, several authors tried to gauge inflation expectations from the term structure of interest rates.[13] While this type of data is typically available, generating inflation forecasts from it is far from being easy. Often a number of simplifying assumptions, e.g. on the real interest rate, have to be made.

Numerous studies analyzed the relative accuracy of inflation forecasts generated by the various forecasting techniques which are in use.[14] However, when evaluating this literature no coherent picture can be drawn since the success of the techniques seems to depend very much on the studied countries, sample periods and the frequency of the available data. In consequence, most institutions do not rely on a single forecasting technique but monitor a variety of inflation indicators and/or employ various forecasting models. As an example, central banks typically use some kind of macroeconometric model in the heart of the forecasting system which is supplied with a number of additional smaller models that are used to generate input

[9] See e.g. Briault (1995).
[10] See e.g. Croushore (1996) or Thomas (1999).
[11] See e.g. Aaker and Day (1990).
[12] Lovell and Vogel (1973) and Lioui and Poncet (2002).
[13] See e.g. Fama (1970) or Mishkin (1990).
[14] See e.g. Litterman (1986), Webb (1999), Stock and Watson (1999) or Kozicki (2001).

data for the core model and to check the validity of its forecasts.[15] However, the fact that no method of forecasting inflation dominates in the practice of forecasting indicates that refining or extending the existing forecasting techniques might be both necessary and fruitful.

From a methodological point of view, the method of forecasting inflation via experimental markets, which will be explained in more detail throughout the next section, belongs to the expectations approach of forecasting. Experimental inflation forecasting markets are an indirect approach of measuring inflation expectations and thus an alternative to gauging inflation expectations from the term structure of interest rates. The major advantage of an experimental market is that it can be designed in a way which allows to derive an inflation forecast with only minimal additional assumptions. Moreover, experimental markets deliver an accurate measure of the uncertainty underlying an inflation forecast which is a highly important information for the users of the forecast. Instead of deriving this measure from historical data it is constructed from last observed market prices and can thus vary continuously.

## 3. EXPERIMENTAL STOCK MARKETS

Using experimental forecasting markets to generate inflation forecasts is some kind of combination of the two conventional methods used within the expectation approach of forecasting. Most economists argue that markets are the most efficient means of aggregating private information (see Smith (1982) and, more recently, Lioui and Poncet (2002)). However, in many countries no appropriate markets for CPI futures evolved. It thus might be useful to conduct small-scale experimental markets in which well informed individuals trade state-contingent contracts thereby revealing their inflation expectations. By using an appropriate design of these contracts we are able to extract not only the mean market expectation of inflation but also some information on the expected likelihood of different inflationary scenarios. In the following we outline the basic setup of an experimental inflation forecast market and show how the data, generated by the market, can be used to construct mean and density forecasts.

### 3.1. Market admittance

Electronic markets are typically fully computerized. To be allowed to take part in a market, participants have to register in the market software via Internet. Within the process of registration the applicants are asked to

[15] See for example the forecasting system of the Bank of England, which is described in detail in Bank of England (1999).

provide some personal information which can later be used to analyze the generated market data.

In general, electronic markets are organized as real-money markets, i.e. all transactions in the market are based on real money. In these markets participants initially have to decide on their personal investments[16] which have to be covered by the traders. Thus, each participant can win or lose money in the market, depending on his or her success in trading. As soon as the initial investment has been transferred to the market organizer (typically this is done via cash or bank transfer to a market account) the participant gets a trader-ID and a password to login the market. In addition to that a trader account for the participant is created and his initial investment is transferred to the account.

Technical precondition for taking part in an electronic market is an Internet connection. In general, there are no formal restrictions for participation in electronic markets.

### 3.2. Market design

The type of market we propose to use is called "winner-takes-all-market". In such a market a set of binary lock-in options is traded. The underlying of these options is some measure of inflation, e.g. CPI inflation as typically measured and announced by national statistical institutes. A binary lock-in option[17] has a fixed, predetermined payoff if the underlying is inside the strike range at expiration. In experimental forecasting markets this payoff is typically normalized to one currency unit (e.g. 1 Euro). The payoff function of such an option is visualized in figure 1. Thus, the type of lock-in options traded in an inflation forecasting market is formally identical to what is called pure, Arrow or Arrow-Debreu securities in financial markets literature.[18]

The set of binary lock-in options which is traded in the market consists of $n$ different options. The strike ranges of these options do not overlap and cover the whole range of possible outcomes of the underlying, i.e. inflation. Since the number of unique linearly independent securities is equal to the total number of alternative states of nature we deal with a complete market.[19] Regardless of the initial distribution of securities it is thus possible to reduce the uncertainty about the value of future wealth to zero. A set of options defining a complete market and the related payoff rules are shown in table 1. Such a complete set of options is also called "unit portfolio" or "bundle".

[16] Typically investments are restricted due to legislative restrictions.

[17] In financial literature this type of option is also called digital, simplex, all-or-nothing, bet or lottery option.

[18] See e.g. Copeland and Weston (1992) or Eichberger and Harper (1997).

[19] Copeland and Weston (1993), p. 112.

**FIG. 1.** Payoff function of a binary lock-in option on the CPI.

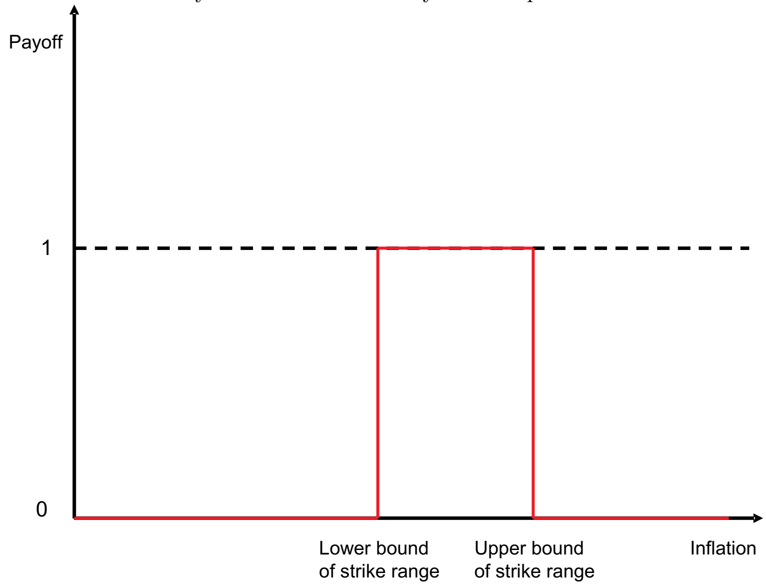


**TABLE 1.**

Complete set of $n$ options.

| Contract name | Pays off 1 currency unit, if |
|---|---|
| $\pi(k_1-)$ | $\pi < k_1$ |
| $\pi(k_1, k_2)$ | $k_1 \leq \pi < k_2$ |
| $\vdots$ | $\vdots$ |
| $\pi(k_{n-1}, k_n)$ | $k_{n-1} \leq \pi < k_n$ |
| $\pi(k_n+)$ | $\pi \geq k_n$ |

There is no general rule how many contracts should be traded and how large the strike ranges should be. In order not to induce some kind of bias by arbitrary contract design it seems to be reasonable to center the contracts symmetrically around the last announced inflation rate. Experiences from political stock markets show that the number of traded contracts $n$ should not be too large ($\leq$ 10) and the strike ranges should not be too small. Otherwise it will be quite hard for the participants to guess how likely it is that inflation will fall into a certain strike range. However, in order to get a forecast as precise as possible the market organizer might be interested in having the possibility to vary the set of traded options during the market period. This can easily be done by splitting options into two (or more) contracts with smaller strike ranges (given that the options do not overlap and all possible outcomes of the underlying are still covered). In order not to influence the values of the participants' portfolios by contract splits, every holder of a split contract gets one of every newly issued contract in exchange.[20]

### 3.3. Trading in the market

Upon entering the market and any time thereafter participants can buy unit portfolios from the market organizer for the price of 1 currency unit (e.g. 1 Euro) until the market closes. Complete unit portfolios can also be sold back to the market organizer during the market period for the price of 1 currency unit each. Selling and buying unit portfolios from or to the market organizer are primary market transactions. Together with the earlier described payoff structure of the binary lock-in options the pricing of the unit portfolios guarantees that the market is a zero-sum-game for the market organizer. All initial investments get paid back to the participants. However, the market is typically no zero sum game for the individual participant since he can win or loose money, depending on his success in trading within the secondary market.

On the secondary market participants can buy or sell contracts from or to other participants. The secondary market is organized as a so-called "double auction market". Market participants can issue offers to buy (bids) or offers to sell (asks) contracts. When using a first type of transactions, so-called "limit orders", traders have to choose the order type (bid or ask), the contract type, the number of contracts he wants to trade, the transaction price and finally the order's expiration date. Limit orders are maintained in separate bid and ask queues ordered first by offer price and then by the time of issuance. Whenever an offer enters one of the queues it remains there until the offer turns out to be unfeasible (e.g. because of a lack of liquidity to realize a buying transaction), is withdrawn by the trader, reaches its expiration date or is carried out. Orders are carried out whenever bid- and ask-prices overlap. The second type of transactions, the so-called "market orders", are orders to buy or sell at current market prices which are carried out immediately.

All primary and secondary market transactions are organized via a market software.[21] Besides serving as a market platform the software provides several facilities for the traders to obtain information on the market. A trader can access personal information on his market account, current portfolio or already submitted orders. The software also delivers information about the highest bids to buy and lowest asks to sell for each traded contract type or the last prices for which a certain share was traded.

[20] We illustrate contract splits by a concrete example later.

[21] The 4 markets we report on in this paper were organized using two different software packages. However, their basic features were quite similar.

Different from real world stock markets, short sales and purchases on margin are typically disallowed to secure the zero-sum game character of experimental forecasting markets. In addition, there are typically no transaction costs levied by the market organizer for both, primary and secondary market transactions. The market participants only have to bear those transaction costs resulting from getting Internet access and opportunity costs from spending time on trading in the market.

### 3.4. Market liquidation

The forecasting markets get liquidated as soon as the realization of the underlying is known, i.e. the inflation rate is announced by the responsible institution. The individual payoff of each participant consists of (i) the money the trader held on his market account when the market closed and (ii) the liquidation value of the portfolio of contracts the trader held at the end of the market. We illustrate the liquidation procedure using an example as shown in table 2.

**TABLE 2.**

Portfolio liquidation for an imaginary participant $j$ under the assumption that the inflation rate turns out to be in between $k_2$ and $k_3$

| Contract/asset | Number of contracts in portfolio of participant $j$ | Liquidation value per contract unit | Total value in Euro |
|---|---|---|---|
| $\pi(k_1-)$ | 76 | 0 | 0 |
| $\pi(k_1, k_2)$ | 4 | 0 | 0 |
| $\pi(k_2, k_3)$ | 13 | 1 | 13 |
| $\pi(k_3, k_4)$ | 2 | 0 | 0 |
| $\pi(k_4, k_5)$ | 5 | 0 | 0 |
| $\pi(k_5+)$ | 0 | 0 | 0 |
| Cash on account | 3 | - | 3 |
| Total payoff | - | - | 16 |

Therefore, we assume that the inflation rate turns out to be in between $k_2$ and $k_3$ percent. The second column in table 2 shows the individual portfolio of participant $j$ and the third column the liquidation values of the contracts. The total value of participant $j$'s portfolio of contracts is thus 13 Euro. Adding the 3 Euro participant $j$ is assumed to hold on his market account at the end of the market period leads to a total payoff of 16 Euro.

### 3.5. Density forecast, mean forecast and forecast uncertainty

In order to show how and why the described design of an experimental stock market should enable us to obtain a reasonable inflation forecast we argue on the basis of Arbitrage Pricing Theory.[22]

[22] Arbitrage Pricing Theory goes back to Ross (1976).

According to Arbitrage Pricing Theory the equilibrium price of a pure security principally depends on three factors: the risk-free rate of return, individuals' attitudes towards risk and expectations as to the probability that a particular state will occur.[23] More precisely we can express the price at time $t$ of any pure security as

$$p_{s,t} = \frac{E[p_{s,T}]}{(1 + r^f + r^r)^{(T-t)}}$$

with $r^f$ being the risk-free rate of return, $E[p_{s,T}]$ the expected payoff of the contract at time $T$ and $r^r$ a risk premium for taking over unsystematic risk (which can not be diversified). In order to determine the equilibrium price of the binary lock-in option traded in a forecasting market we consider these three determinants sequentially.

Let us first focus on the risk-free rate of return. There are two risk-free portfolios in an inflation forecasting market. The first one consists of not holding any of the pure securities at all. Obviously, the return on this portfolio is zero. The second risk-free portfolio consists of one of each of the pure securities. Such a portfolio would deliver 1 unit of currency in every state of the nature. Such a unit portfolio can be purchased from or sold back to the market organizer at any time during the market period for the price of one currency unit. Thus, the return on a unit-portfolio is zero by construction. One might be tempted to argue that the price of a unit-portfolio could be smaller than 1 currency unit when the sum of asks to sell prices adds up to a smaller value than 1. However, such a situation can not be an equilibrium since it is not arbitrage-free.[24] One might also argue that there is a significantly positive risk-free return outside the inflation market, e.g. government bonds. However, once the decision to transfer money to the market account has been made (a necessary precondition to take part in an inflation market), the money cannot be withdrawn during the market period. Thus, the risk-free rate of return is zero ($r^f = 0$) in an inflation forecasting market.

The second determinant of a pure security's price lies in individual beliefs concerning the relative likelihood of different states $s$ occurring, the so-called state probabilities. For simplicity, let us assume that in equilibrium individuals agree on the probabilities $h_{\pi(k_{n-1},k_n),t}$ of the states of nature.[25]

[23] Copeland and Weston (1993), p. 116.

[24] In the described situation one could easily make sure profits by buying unit-portfolios on the secondary market and selling them for the price of 1 currency unit to the market organizer. Realizing these transactions will quickly eliminate all arbitrage opportunities.

[25] This assumption is without effect on the line of argument. Principally, subjective beliefs concerning state probabilities can also differ (see e.g. Copeland and Weston (1993), p. 117).

In that case we can decompose the expected price of a pure security in state $\pi(k_{n-1}, k_n)$ at time $t$ into the probability of the state occurring and the price of an expected currency unit contingent on the state occurring. Thus, we have

$$E[p_{\pi(k_{n-1},k_n),T}] = h_{\pi(k_{n-1},k_n),t} \cdot 1.$$

In consequence, expected prices of pure securities differ to the same degree as market participants expect different states to occur with different probabilities.

The third determinant of pure securities' prices is market participants' attitude toward risk. While there are obviously no risk premia for the case of risk-neutral investors, one might argue that risk-averse individuals will demand for such a premium in order to take over risk. However, this is only true for the case that aggregate wealth differs between the different states of nature.[26] In the described market setting aggregate wealth is the same regardless of which state is realized (this is due to the zero-sum game character of the market). Thus, there is no non-diversifiable risk and, consequently, equilibrium prices include no risk premia, i.e. $r^r = 0$.

Altogether, we end up with the following pricing formula for the binary lock-in options traded in inflation forecasting markets:

$$\begin{aligned} p_{\pi(k_{n-1},k_n),t} &= \frac{E[p_{\pi(k_{n-1},k_n),T}]}{(1 + r^f + r^r)^{(T-t)}} \\ &= \frac{h_{\pi(k_{n-1},k_n),t} \cdot 1}{(1 + 0 + 0)^{(T-t)}} \\ &= h_{\pi(k_{n-1},k_n),t}. \end{aligned}$$

Thus, the prices of the pure securities traded in an inflation market are perfect predictors of the probabilities, market participants attach to the different states of nature.[27] However, this is true only if the market is in equilibrium. It is then when all available information is reflected in the current market prices. However, there is still a lack of a commonly accepted dynamic model how the market participants learn from the observed market prices, i.e. how exactly the process of aggregating disseminated information works. Although one might be able to build such a model it will be hard to test it empirically since the necessary data on individual beliefs are hard to obtain. One might also be somewhat sceptic whether a formal behavioral model will be able to capture the diversities of individual learning.

[26] See e.g. Copeland and Weston (1993), p. 118.

[27] The same results can principally be derived from Capital Asset Pricing Theory (CAPM). Since the assumptions of CAPM are somewhat more restrictive than those of Arbitrage Theory we base our expositions on the latter.

While an experimental inflation forecasting market directly generates a density forecast of inflation it does not automatically deliver some form of mean inflation forecast. Whenever we are in need of mean forecasts we have to make some simplifying assumptions on the distribution of inflation expectations within the intervals as marked by the strike ranges of the option contracts. For sufficiently small intervals it seems to be reasonable to assume that the market participants expect all realizations of inflation within this interval to be equally likely. In this case the interval can be represented by its class middle. However, a complete set of options includes two options with infinitely large strike ranges ($\pi(k_1-)$ and $\pi(k_5+)$ in the example in table 1). To deal with this problem one might use the (upper respectively the lower) bounds of these infinitely large intervals instead of the class middles.[28] We can then calculate the mean market inflation forecast $\pi_t^f$ at time $t$ by multiplying the last observed market prices with the class middles (respective the bounds of the lowest and the highest interval) and adding up for all traded contracts

$$\begin{aligned}\pi_t^f \;=\; & p_{\pi(k_1-),t} \cdot k_1 + p_{\pi(k_1,k_2),t} \cdot \frac{k_2 - k_1}{2} + \ldots \\ & + p_{\pi(k_{n-1},k_n),t} \cdot \frac{k_n - k_{n-1}}{2} + p_{\pi(k_n+),t} \cdot k_n\end{aligned}$$

Experiences from previous electronic markets research showed that last traded prices do not always add up to one. One possible reason for this observation is that traders in most markets were cash-restrained due to the maximum initial investment prescription. Thus, even if some traders would have perfect information on the fair prices of all contracts they will typically not have enough funds to fix the prices to their fair values. While this problem should diminish in markets with a high number of traders it can hardly be ignored in smaller markets. To deal with this problem most market organizers normalize the sum of last traded prices to unity before calculating the mean market forecast. We follow this procedure and calculate the mean market inflation forecast as

$$\begin{aligned}\pi_t^f \;=\; & \frac{p_{\pi(k_1-),t}}{P_t} \cdot k_1 + \frac{p_{\pi(k_1,k_2),t}}{P_t} \cdot \frac{k_2 - k_1}{2} + \ldots \\ & + \frac{p_{\pi(k_{n-1},k_n),t}}{P_t} \cdot \frac{k_n - k_{n-1}}{2} + \frac{p_{\pi(k_n+),t}}{P_t} \cdot k_n.\end{aligned}$$

[28] Doing so is obviously problematic when the observed prices for the options covering the infinitely large intervals are quite high. However the market administrator can lower the market prices of these intervals by making use of the earlier described split option of contracts.

with

$$P_t := p_{\pi(k_1-),t} + \cdots + p_{\pi(k_n+),t}.$$

In the literature on political stock markets often daily volume-weighted prices are used instead of last observed prices for generating forecasts. We therefore also report forecasts which are based on normalized weighted prices. However, from a theoretical point of view last observed prices should be superior since they belong to the most actual, marginal transactions.

By far most published inflation forecasts are mean forecasts. Typically these forecasts do not provide any information on the underlying probabilities of different inflation scenarios. Since one and the same mean forecast can principally result from many different probability distributions information about the uncertainty surrounding a mean inflation forecast is important in addition to the forecast itself. A measure of forecast uncertainty helps to qualify a forecast and is useful to give a richer picture of the expected range of likely outcomes.[29] Inflation forecasting markets allow to assess the mean inflation forecast's uncertainty directly. Since the normalized market prices $p^n_{t,j}$ (either last observed or weighted prices) can be interpreted as the market's aggregated evaluation of the probabilities of different inflationary scenarios, these probabilities can be used to calculate the empirical variance of the daily mean inflation forecast as

$$\begin{aligned}\sigma^2_{\pi^f,t} = & \frac{p_{\pi(k_1-),t}}{P_t} \cdot (k_1 - \pi^f_t)^2 + \frac{p_{\pi(k_1,k_2),t}}{P_t} \cdot \left(\frac{k_2 - k_1}{2} - \pi^f_t\right)^2 + \ldots \\ & + \frac{p_{\pi(k_{n-1},k_n),t}}{P_t} \cdot \left(\frac{k_n - k_{n-1}}{2} - \pi^f_t\right)^2 + \frac{p_{\pi(k_n+),t}}{P_t} \cdot (k_n - \pi^f_t)^2.\end{aligned}$$

It should be underlined that electronic markets allow to provide forecasts and to assess their empirical variance at any point in time during the market period. Thus, electronic forecasting markets deliver time-series of fixed-event forecasts.

## 4. RESULTS FROM INFLATION FORECASTING MARKETS

In this section we report and analyze the results from a series of inflation forecasting markets conducted in Germany. While the four inflation forecasting markets differed in some respects, they all made use of the basic design developed in the previous section. We start out with a description of the basic setup and major properties of the four inflation forecasting markets and then turn to a presentation of the time-series of (fixed-event)

[29]See e.g. Ericsson (2001), p. 88-89.

mean inflation forecasts constructed from the market data. In the next subsection we assess the uncertainty surrounding the mean forecasts and analyze how this uncertainty behaved over time. We then evaluate the accuracy of the mean inflation forecasts by applying the concept of weak rationality of fixed-event forecasts developed by Nordhaus (1987). We also show that it is easily possible to construct forecast confidence intervals from the market data when there is information on the underlying distributional form of the forecast.

## 4.1. Market descriptions

### *4.1.1. The February 2001 inflation market*

The first experimental inflation forecasting market was organized at Dresden University of Technology (Germany) in close cooperation with the Iowa Electronic Markets (United States). The market was designed to forecast the German February 2001 CPI inflation rate and was conducted using the IEM software. The market opened on 20th October 2000 and was closed on 14th March 2001, soon after February 2001 CPI inflation was announced on 10th March. All transactions were based on real money. Altogether, 44 traders participated in the market, most of which were students of economics and business administration at Dresden University of Technology. While the market was principally open to all interested people, the market was advertised primarily in economics courses at Dresden University of Technology. The total amount of money invested was 1021.62 Euro (23.22 Euro per trader). In the following we will refer to this market as "market I". The complete set of contracts traded in market I and the contracts' payoff structure is shown in table 3.

**TABLE 3.**

Traded contracts in market I.

| Contract number | Contract name | Interval middle/limit | Pays off 1 Euro, if |
|---|---|---|---|
| 1 | $\pi(0.0-)$ | 0.000 | $\pi < 0.0$ |
| 2 | $\pi(0.0-1.5)$ | 0.750 | $0.0 \leq \pi < 1.5$ |
| 3 | $\pi(1.5-2.0)$ | 1.750 | $1.5 \leq \pi < 2.0$ |
| 4 | $\pi(2.0-2.5)$ | 2.250 | $2.0 \leq \pi < 2.5$ |
| 5 | $\pi(2.5-3.0)$ | 2.750 | $2.5 \leq \pi < 3.0$ |
| 6 | $\pi(3.0-3.5)$ | 3.250 | $3.0 \leq \pi < 3.5$ |
| 7 | $\pi(3.5-4.0)$ | 3.750 | $3.5 \leq \pi < 4.0$ |
| 8 | $\pi(4.0+)$ | 4.000 | $4.0 \leq \pi$ |

As it is shown in table 3 there were initially eight binary lock-in options in the market. On November 22, the contract "$\pi(2.0-2.5)$" was split into two contracts "$\pi(2.0 - 2.25)$" and "$\pi(2.25 - 2.5)$". Therefore, each participant who held former "$\pi(2.0-2.5)$"-contracts in his portfolio was endowed with the same number of the two new contracts. Thus, the expected value of the participants' portfolios was not influenced by the contract split. The contract split was done because it was observed that the former "$\pi(2.0 - 2.5)$"-contract had been traded for quite high prices, thus indicating that the participants attached a high probability to the event that the inflation rate would have been in between 2.0 and 2.5 percent.

#### *4.1.2. The December 2001 inflation market*

The second experimental inflation forecasting market was conducted by Dresden Electronic Markets (DEM) at Dresden University of Technology. The CPI inflation rate to be forecasted was the one of December 2001 in Germany. The market opened on 17th October 2001 and closed on 15th January 2002. In the following we refer to this market as "market II". The complete set of contracts traded in market II and the contracts' payoff structure is shown in table 4.

**TABLE 4.**

Traded contracts in market II.

| Contract number | Contract name | Interval middle/limit | Pays off 1 Euro, if |
|---|---|---|---|
| 1 | $\pi(1.0-)$ | 1.000 | $\pi < 1.0$ |
| 2 | $\pi(1.0 - 1.5)$ | 1.250 | $1.0 \leq \pi < 1.5$ |
| 3 | $\pi(1.5 - 1.75)$ | 1.625 | $1.5 \leq \pi < 1.75$ |
| 4 | $\pi(1.75 - 2.0)$ | 1.875 | $1.75 \leq \pi < 2.0$ |
| 5 | $\pi(2.0 - 2.5)$ | 2.250 | $2.0 \leq \pi < 2.5$ |
| 6 | $\pi(2.5 - 3.0)$ | 2.750 | $2.5 \leq \pi < 3.0$ |
| 7 | $\pi(3.0 - 3.5)$ | 3.250 | $3.0 \leq \pi < 3.5$ |
| 8 | $\pi(3.5+)$ | 3.500 | $3.5 \leq \pi$ |

Different from the earlier described pioneer market the December 2001 market was conducted with a newly designed software, providing some additional features. However, the basic trading system is almost identical to the one of the IEM software. Altogether, 32 traders took part in the market most of which were again students of economics and business administration at Dresden University of Technology. Again the market was principally open to all interested people. The sum of investments was 414 Euro (12.94 Euro per trader).

### *4.1.3. The June 2002 inflation market*

Market organizer of the third market was again Dresden University of Technology and again the DEM software was used to conduct the market. The market was designed to forecast the June 2002 CPI inflation rate in Germany. The trading period begun on 23rd April and ended on 11th July 2002. A total number of 47 traders took part in the market most of which were again students of economics and business administration at Dresden University of Technology. However, a considerable number of people outside the university took part in the market. The sum of investments was 841 Euro (on average 17.87 Euro per trader). In the following we refer to this market as "market III". The complete set of contracts traded in the market and the contracts' payoff structure is shown in table 5.

**TABLE 5.**

Traded contracts in market III.

| Contract number | Contract name | Interval middle/limit | Pays off 1 Euro, if |
|---|---|---|---|
| 1 | $\pi(1.0-)$ | 1.000 | $\pi < 1.0$ |
| 2 | $\pi(1.0 - 1.25)$ | 1.125 | $1.0 \leq \pi < 1.25$ |
| 3 | $\pi(1.25 - 1.5)$ | 1.375 | $1.25 \leq \pi < 1.5$ |
| 4 | $\pi(1.5 - 1.75)$ | 1.625 | $1.5 \leq \pi < 1.75$ |
| 5 | $\pi(1.75 - 2.0)$ | 1.875 | $1.75 \leq \pi < 2.0$ |
| 6 | $\pi(2.0 - 2.25)$ | 2.125 | $2.0 \leq \pi < 2.25$ |
| 7 | $\pi(2.25 - 2.5)$ | 2.375 | $2.25 \leq \pi < 2.5$ |
| 8 | $\pi(2.5 - 3.0)$ | 2.750 | $2.5 \leq \pi < 3.0$ |
| 9 | $\pi(3.0 - 4.0)$ | 3.500 | $3.0 \leq \pi < 4.0$ |
| 10 | $\pi(4.0+)$ | 4.000 | $4.0 \leq \pi$ |

### *4.1.4. The October 2002 inflation market*

The last German inflation forecasting market we report on was again organized by Dresden Electronic Markets with the DEM software. The market was designed to foreast the October 2002 CPI inflation rate in Germany. The trading period begun on 26th July and ended on 12th November 2002. A total number of 36 traders took part in the market most of which were identical to the traders in market III. The sum of investments was 288 Euro (on average 8.00 Euro per trader). In the following we refer to this market as "market IV". The complete set of contracts traded in the market and the contracts' payoff structure is shown in table 6.

**TABLE 6.**

Traded contracts in market IV.

| Contract number | Contract name | Interval middle/limit | Pays off 1 Euro, if |
|---|---|---|---|
| 1 | $\pi(0.0-)$ | 0.000 | $\pi < 0.0$ |
| 2 | $\pi(0.0-0.5)$ | 0.250 | $0.0 \leq \pi < 0.5$ |
| 3 | $\pi(0.5-0.75)$ | 0.625 | $0.5 \leq \pi < 0.75$ |
| 4 | $\pi(0.75-1.00)$ | 0.875 | $0.75 \leq \pi < 1.00$ |
| 5 | $\pi(1.00-1.25)$ | 1.125 | $1.00 \leq \pi < 1.25$ |
| 6 | $\pi(1.25-1.75)$ | 1.500 | $1.25 \leq \pi < 1.75$ |
| 7 | $\pi(1.75-2.25)$ | 2.000 | $1.75 \leq \pi < 2.25$ |
| 8 | $\pi(2.25+)$ | 2.250 | $2.25 \leq \pi$ |

## 4.2. Mean inflation forecasts

As discussed earlier electronic inflation forecasting markets allow to obtain actual inflation forecasts at any point in time during the market period. Thus, for every of the four prototype markets we end up with a time series of inflation forecasts $\pi^f_{T,T-t}$ of the same event, i.e. the inflation rate $\pi_T$ at time $T$. Thus, we deal with so-called fixed-event forecasts.[30]

While we principally can construct forecasts of any frequency we decided to work with daily forecasts since new information on inflation is typically not arriving more often than once a day, if at all. For every market we report two types of forecasts. The first forecast, which we will call "last traded prices forecast" (LTP) is based on the prices of the last observed transaction for every contract type at midnight. The second forecast, the so-called "average traded prices forecast" (ATP) is calculated on the basis of the volume-weighted daily average price of every type of contract. When there were no transactions on a certain day, the forecast remains on the previous day's value.

Experiences with previous markets showed that it typically takes several days before trading in the markets begins. On the one hand this is due to the time lag between applying for admittance for a market and transfer of the initial investment on the market account. On the other hand the number of traders in the market is quite low in the beginning of every market, leading to relatively few and unrealistic offers placed in the market queues. We therefore report forecasts not before every contract type has been traded in the market at least once.

The mean inflation forecasts generated by the four markets are shown in figures 2 to 5 . Besides the forecasts the figures also show the actual

[30] Clements and Hendry (1998).

inflation rate as it had to be predicted and the last announced inflation rate.

**FIG. 2.** Mean inflation forecast for February 2001.

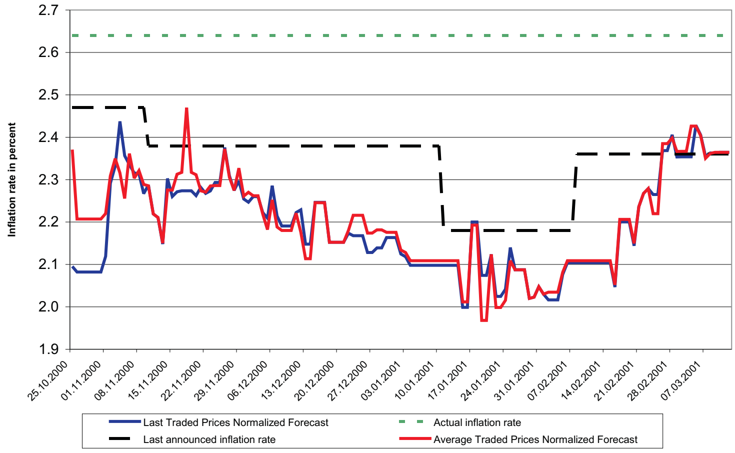


**FIG. 3.** Mean inflation forecast for December 2001.

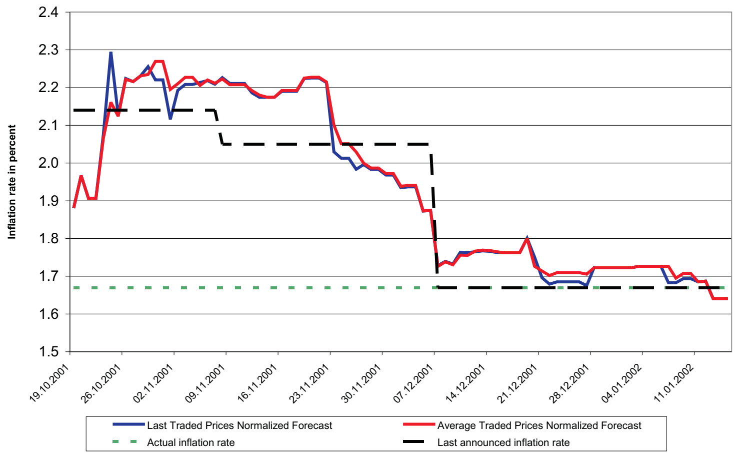


### 4.3. Uncertainty of mean inflation forecasts

Whenever inflation has some predictable component we should expect empirical variances of the forecasts to decrease in the course of time since the market participants are getting more information over the market peri-

**FIG. 4.** Mean inflation forecast for June 2002.

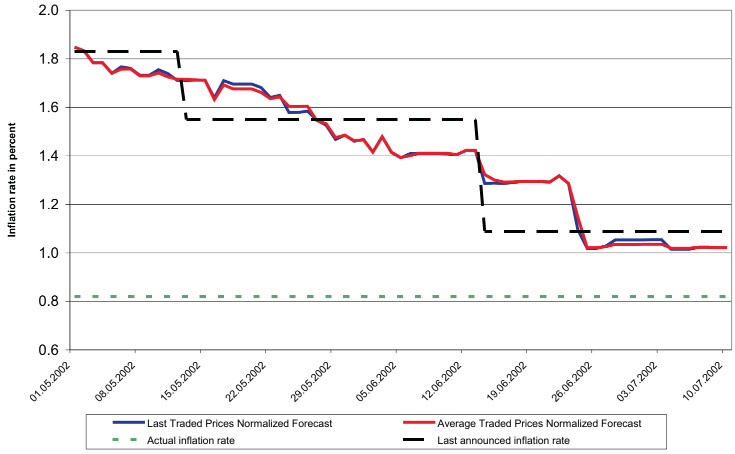


**FIG. 5.** Mean inflation forecast for October 2002.

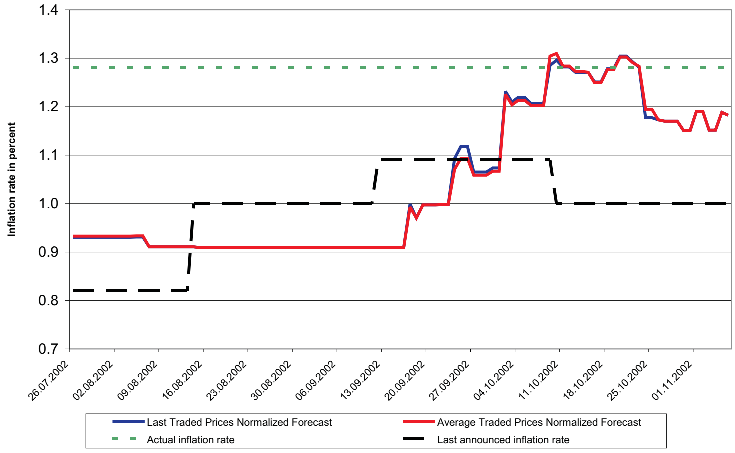


od. In figure 6 we show the variances of the mean inflation forecasts (LTP) during the market periods.

In fact, there seems to be a decreasing tendency of the empirical variance during the market periods. In order to test for time trends in the forecast variances formally we run OLS-regressions of the type

$$\sigma_t^2 = c + \alpha \cdot t + \epsilon_t$$

**FIG. 6.** Variances of daily mean inflation forecast (LTP).

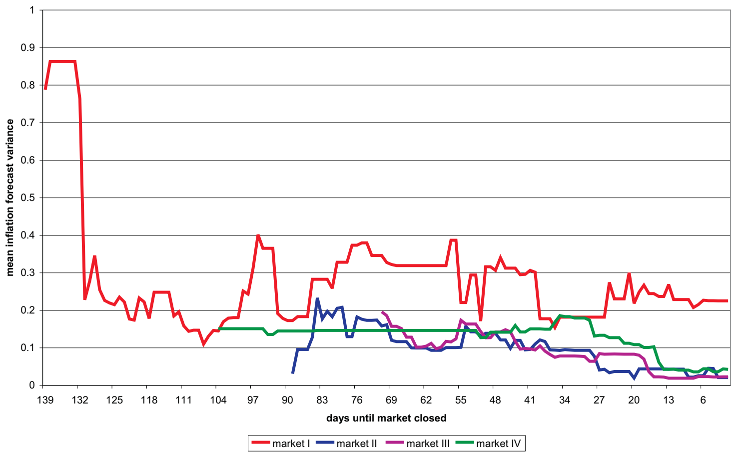


for every market ($c$ is the regression constant and $\epsilon_t$ the unexplained residual). The existence of a negative time trend implies the coefficient $\alpha$ of the time variable $t$ to be negative. The regression results are reported in table 7. We find highly significant negative time trends of the forecast variances for all four markets.[31]

**TABLE 7.**

Time trends of inflation forecast variances (LTP).

| Market | Constant | T-statistic | Probability | Trend | T-statistic | Probability | Obs. |
|---|---|---|---|---|---|---|---|
| market I | 0.377 | 15.28 | 0.00 | -0.001 | -4.29 | 0.00 | 138 |
| market II | 0.169 | 25.25 | 0.00 | -0.002 | -12.35 | 0.00 | 83 |
| market III | 0.165 | 31.17 | 0.00 | -0.002 | -16.32 | 0.00 | 72 |
| market IV | 0.173 | 28.75 | 0.00 | -0.001 | -8.28 | 0.00 | 111 |

### 4.4. Forecast accuracy and efficiency

In the end the value of a forecasting technique depends very much on the accuracy of the forecasts. In the following we therefore spend some effort on analyzing the accuracy of the forecasts generated by the 4 pilot markets.

[31] We should note that the results for market I are primarily driven by the excessively high variance in the first days of the active trading period. When excluding the first week of market activity the time trend is still negative but insignificant.

Several ways of assessing a forecast's accuracy have been discussed in the literature. However, the fact that we only have four forecasting markets at hand is obviously limiting our possibilities to study the predictive quality of inflation forecasts generated by experimental markets. For example, we can hardly run any test on forecast rationality since we have no time series of rolling event forecasts. We also have no comparable forecasts published by other institutions. Nevertheless, there are some possibilities to get a an impression on the accuracy and efficiency of the inflation forecasts generated by electronic markets. First, we can study in how far the fixed event forecasts from the markets perform in comparison to naive forecasts[32] which are often used as a kind of benchmark of forecasting models. Second, we can study in how far the markets show signs of weak efficiency.

We start out with an analysis of the relative performance of the markets' forecasts in comparison with naive forecasts. We therefore study on a daily basis whether an agent could improve on the market forecast when simply expecting that inflation will stay the same as it was when last announced by Statistisches Bundesamt Wiesbaden. An inspection of the figures 2 to 5 already gives some impression of the results of a more formal analysis. The February 2001 inflation market did obviously perform much worse than the naive forecasts. While all forecasts underestimate the actual inflation rate during the whole market period, the naive forecast is closest to the final inflation rate with only a few exceptions. In the December 2001 market all forecasts in general overestimate the actual rate of inflation. While the picture is not as clear as in the February 2001 market, the naive forecast seems to perform better than the market forecasts. In the June 2002 market again all forecasts overestimated the actual inflation rate during the whole market period. However, the market forecast was closer to the final result for a considerable part of time. Similarly the market seems to have performed quite well in the October 2002 market.

In table 8 we show the mean absolute forecast errors of both market forecasts and naive predictions. The results confirm the findings of the inspection of the graphs. The naive forecast outperforms both market forecasts quite clearly in the February 2001 market. In the remaining three markets the differences are considerably lower. While the naive forecast performed better in the December 2001 market, the opposite is true for the June 2002 and the October 2002 markets.

[32]A naive forecast predicts that the variable to be forecasted stays as it was when last observed.

**TABLE 8.**

Comparison between market and naive forecasts.

| Market | Average LTP forecast error | Average ATP forecast error | Average naive forecast error | Observations |
|---|---|---|---|---|
| Market I | 0.45 | 0.44 | 0.30 | 138 |
| Market II | 0.28 | 0.29 | 0.25 | 83 |
| Market III | 0.60 | 0.60 | 0.62 | 72 |
| Market IV | 0.27 | 0.27 | 0.30 | 111 |

Let us now turn to an analysis of the efficiency of the time series of fixed event market inflation forecasts. In order to do so we make use of the concept of testing for weak efficiency proposed by Nordhaus (1987). Up to now, the concept has rarely been applied to inflation forecasts. One might suggest this to be due to the fact that most existing time series of inflation forecast are not fixed but rolling event forecasts. One of the rare tests for efficiency of a fixed event time series of price forecasts was conducted by Nordhaus (1987) himself. When analyzing a time series of oil price forecasts collected by Data Resources Inc. he finds significant autocorrelation of forecast revisions. Nordhaus also reports the results of weak efficiency tests of 3 additional time series of fixed event forecasts (forecasts of nuclear capacity, energy forecasts and real GNP forecasts). Similarly, the forecast revisions turn out to be autocorrelated. Nordhaus interprets his finding that most of the significant autocorrelations are positive as an indication that the hypothesis of social psychologists that people tend to hold on to prior views too long (see Tversky and Kahneman (1981)) might be correct.

Recently Bakhshi, Kapetanios and Yates (2003) studied the efficiency of 7 time series of fixed event inflation forecasts. The forecasts were constructed by Meryll Lynch from a survey of about 70 fund managers. Respondents had to predict the annual increase in the U.K. Retail Price Index at December 1994, 1995, 1996 and 1997 and the annual increase in the U.K. RPIX index at December 1998, 1999 and 2000. The surveys were conducted monthly providing a database of 23 forecasts per event. Bakhshi, Kapetanios and Yates reject the hypothesis that the forecast errors are uncorrelated with past revisions for 5 out of 7 time series of fixed event forecasts. Similarly 2 of the time series exhibit autocorrelation of forecast revisions.

The major problem of the studies by Nordhaus (1987) and Bakhshi, Kapetanios and Yates (2003) is the relatively low number of observations per time series of fixed event forecasts. The time series generated by electronic forecasting markets are considerably longer since they principally allow to generate continuous inflation forecasts.[33] Thus, these time series provide an excellent database to study the efficiency of the forecasts.

We start out with analyzing in how far the forecast errors constructed by market prices are correlated with past forecast revisions. We therefore run the OLS regression

$$\pi_\tau - \pi_t^f = \alpha_0 + \alpha_1 \cdot (\pi_{t-1}^f - \pi_{t-2}^f) + \alpha_1 \cdot (\pi_{t-2}^f - \pi_{t-3}^f) + \epsilon_t. \qquad (1)$$

The results are shown in table 9. For all markets and all forecasts we find a highly significant constant indicating that the forecasts are biased. Different from tests on rationality of rolling event forecasts, such a bias is no sign of inefficiency of fixed event forecasts. However, both forecasts constructed from market I show significant first-order correlation with past forecast revisions. The forecasts from the remaining three markets show no significant correlation with past forecast revisions and thus pass the first test on efficiency.

**TABLE 9.**

Correlation between forecast errors and past forecast revisions.

| Market | Forecast ype | Constant (t-value) | $\pi_{t-2}^f - \pi_{t-1}^f$ (t-value) | $\pi_{t-3}^f - \pi_{t-2}^f$ (t-value) |
|---|---|---|---|---|
| Market I | LTP | -0.45 | 0.45 | 0.31 |
| | | $(-48.62)^{***}$ | $(2.32)^{*}$ | (1.63) |
| Market I | ATP | -0.44 | 0.38 | 0.14 |
| | | $(-46.20)^{***}$ | $(2.13)^{*}$ | (0.83) |
| Market II | LTP | 0.29 | 0.69 | 0.64 |
| | | $(11.20)^{***}$ | (1.44) | (1.36) |
| Market II | ATP | 0.28 | 0.95 | 0.91 |
| | | $(11.62)^{***}$ | (1.48) | (1.46) |
| Market III | LTP | 0.59 | 0.42 | 0.40 |
| | | $(16.95)^{***}$ | (0.50) | (0.47) |
| Market III | ATP | 0.60 | 0.45 | 0.49 |
| | | $(16.76)^{***}$ | (0.48) | (0.52) |
| Market IV | LTP | -0.24 | 0.54 | 0.50 |
| | | $(-16.18)^{***}$ | (0.98) | (0.90) |
| Market IV | ATP | -0.24 | 0.55 | 0.49 |
| | | $(-16.30)^{***}$ | (0.96) | (0.86) |

[33]Since new information on future inflation typically not occurs more than once a day (if ever) it is somewhat questionable whether it makes sense to construct forecasts with higher than daily frequencies.

In a second test of efficiency we study in how far the forecast revisions are autocorrelated. We therefore run the OLS regression

$$\pi_t^f - \pi_{t-1}^f = \beta \cdot (\pi_{t-1}^f - \pi_{t-2}^f) + \epsilon_t. \tag{2}$$

The results are reported in table 10. Again we run the regression for both types of forecasts (LTP,ATP). For none of the forecasts we find significant positive first-order correlation, indicating that the market forecasts incorporate newly arriving information in an efficient manner.[34]

**TABLE 10.**

First-order autocorrelation of forecast revisions.

| Market | Forecast type | $\pi_{t-2}^f - \pi_{t-1}^f$ | T-value | Significance |
|---|---|---|---|---|
| Market I | LTP | 0.12 | 1.39 | 0.17 |
| Market I | ATP | -0.18 | -2.22 | $0.03^{**}$ |
| Market II | LTP | 0.12 | 1.10 | 0.27 |
| Market II | ATP | -0.03 | -0.32 | 0.75 |
| Market III | LTP | -0.06 | -0.49 | 0.63 |
| Market III | ATP | -0.14 | -1.18 | 0.24 |
| Market IV | LTP | 0.04 | 0.40 | 0.69 |
| Market IV | ATP | 0.05 | 0.51 | 0.61 |

Altogether, the results indicate that while the markets were quite efficient in disseminating information the relative performance of the market inflation forecasts was not overwhelming since the naive forecast outperformed the market forecast in two out of four cases. However, naive forecasts are known to perform quite well in short-term forecasting when dealing with hysteretic variables like inflation. Thus, the benchmark of naive forecasts is quite restrictive when dealing with short-term forecasts as we did in the four pilot markets (on average the forecasts we deal with are 50-days-ahead forecasts). Therefore, the fact that the market forecast outperformed the naive forecasts in at least two out of four cases is quite encouraging. Market inflation forecasts might perform even better when forecasting at larger time horizons.

However, even though the markets did not systematically outperform naive forecasts in our pilot experiments they deliver a considerably higher degree of information on future inflation. This is due to the fact that experimental inflation markets not only allow to construct mean forecasts

[34]We should note that we found a significantly negative coefficient for the ATP forecast in market 1. However, since the coefficient is negative the reason for this can hardly be that traders stick to their expectations for too long.

of inflation but also deliver important information on the likelihood of alternative inflationary scenarios. We will extend the discussion on this argument throughout the subsequent subsection.

### 4.5. Normally distributed forecasts and further applications

It was already shown that electronic inflation forecasting markets allow to calculate a mean forecast and its variance (or standard deviation) at any point in time during the market period. In addition, for any point in time we can visualize the market's actual evaluation of the probability of different inflation realizations in a histogram. We might illustrate this at the example of data from the February 2001 inflation market. In figure 7 we show the empirical distribution of the February 2001 inflation forecast of 12th December 2000 (LTP). For every contract we have one observation which refers to the last observed transactions (again we could use weighted prices instead).

**FI** ).

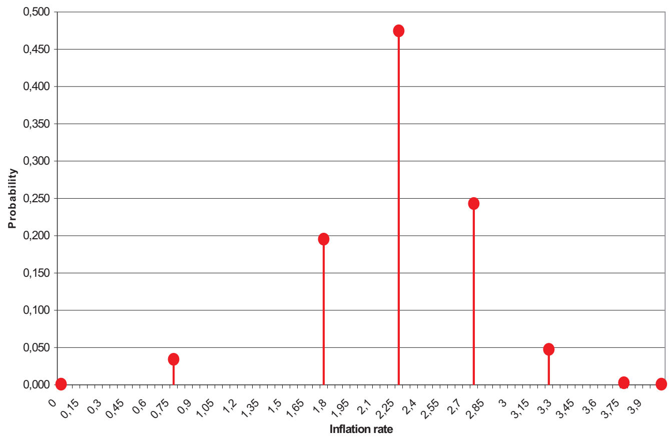


An inspection of the histogram suggests that inflation expectations might be normally distributed. When extending the inspection to a larger number of days at different points in time during the market period and to the other 3 markets this suggestion is substantiated. However, because of the relatively low number of observations per forecast we are not able to test for this hypothesis formally.

Obviously, knowing about the distributional form of the market inflation forecasts would be highly valuable since we would be able to make more precise statements about the probability of different inflationary scenarios. To show this, we will assume in the following that the market forecasts are normally distributed - an assumption which has often been made for individual forecasts. However, the applications we present in the following do not crucially depend on the assumption of a certain distribution. Principally they can easily be adapted to other distributional forms. The assumption on the distributional form of the forecast together with our empirical observations of its mean and variance enables us (i) to calculate the probability for every inflationary scenario that might be of interest, even if no contract in the market has been traded for the referring interval and (ii) to construct forecast confidence intervals.

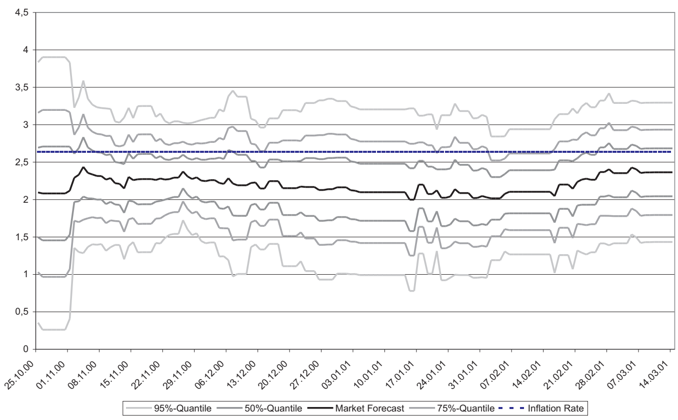


Soon after the Bank of England started to use fan charts to present their inflation forecasts this type of graphical illustration has become standard. Fan charts include a graphical representation of forecast confidence intervals for inflation at different times in the future. Since our prototype market generates a time series of fixed event forecasts, we would need several markets to reproduce such a fan chart. To produce a modified version of a fan chart we calculate the confidence intervals for different $\alpha$-levels and for different forecasting dates (see figure 8). Thus, we receive some graphical representation of how different forecast confidence intervals developed in the course of time.

## 5. SUMMARY AND OUTLOOK

In this paper we showed that conducting electronic markets can fill a gap in existing forecasting techniques. Properly designed experimental markets seem to be quite efficient in information aggregation and deliver time series of (fixed event) interval forecasts which can easily be transformed into mean forecasts and be supplied by an accurate measure of forecast uncertainty.[35] Altogether, the results indicate that building up a regular forecasting system for inflation (and possibly for additional macroeconomic variables) might be a fruitful task. Doing so would allow to evaluate the forecasts constructed from experimental forecasting markets in a more systematic way.

Of course, it would be desirable to run also experimental forecasting markets with longer time-horizons, such as one or even several years. However, to organize and conduct medium- or long-term inflation forecasting markets is not too easy since the markets can not be liquidated before the event, the market is conducted on, has realized. Thus, when running a forecasting market on the two-year ahead inflation rate, what is technically possible, the market participants receive their payoffs after the same period of time. While the market organizer could invest the initial investments in some interest bearing asset and pay some interest on the payoffs it is nevertheless not easy to motivate participants to take part in a market with such a long time horizon. In order to motivate traders to engage even in medium-term markets one could combine short-term markets with medium-term ones. First experiences with such a staggered system of forecasting markets have been made in a series of markets conducted throughout 2002 in Bulgaria.[36] The fact that almost all traders engaged in the short-term markets, also took part in the medium-term markets is quite promising in this respect.

[35] Of course, even in experimental forecast markets the quality of the mean and the interval forecast depends crucially on the quality of infiormation avaiable to the traders. However, since information can spread easily through the market, markets will likely perform better than e.g. surveys.

[36] See Berlemann, Dimitrova and Nenovsky (2005).

## REFERENCES

Aaker, David A., and George S. Day, 1990. *Marketing Research.* 4th Edition, New York: John Wiley.

Atkeson, Andrew, and Lee E. Ohanian, 2001. Are Phillips Curves Useful for Forecasting Inflation? *Federal Reserve Bank of Minneapolis Quarterly Review* **25(1)**, 2-11.

Bakshi, Hasan, George Kapetanios, and Anthony Yates, 2003. Rational Expectations and Fixed-event Forecasts: An Application to UK Inflation. Bank of England Working Paper 176, Bank of England, London.

Bank of England, 1999. Economic Models at the Bank of England. London.

Berg, Joyce, Robert Forsythe, Forrest Nelson, and Thomas Rietz, 1998. Results from a Decade of Election Futures Markets Research. Working Paper, College of Business Administration, University of Iowa.

Berg, Joyce, Robert Forsythe, and Thomas Rietz, 1997. What Makes Markets Predict Well? Evidence from the Iowa Electronic Markets. In Understanding Strategic Interaction. Essays in Honor of Reinhard Selten. Edited by W. Albers et al., Berlin: Springer, 444-463.

Berlemann, Michael, Kalina Dimitrova, and Nikolay Nenovsky, 2005. Assessing Market Expectations on Exchange Rates and Inflation: A Pilot Forecasting System for Bulgaria. William Davidson Institute Working Paper 759.

Berlemann, Michael, and Carsten Schmidt, 2001. Predictive Accuracy of Political Stock Markets. Empirical Evidence from a European Perspective. Discussion Paper, Sonderforschungsbereich 373, Humboldt-University Berlin.

Briault, Clive B., 1995. The Costs of Inflation. *Bank of England Quarterly Review* **35(1)**, 33-45.

Clements, Michael P., and David F. Hendry, 1998. *Forecasting Economic Time Series.* Cambridge: Cambridge University Press.

Copeland, Thomas E., and John F. Weston, 1992. *Financial Theory and Corporate Policy.* 3rd Edition, Reading/Mass.: Addison Wesley.

Croushore, Dean, 1996. Inflation Forecasts: How Good Are They? *Federal Reserve Bank of Philadelphia Business Review* **May/June**, 1-11.

Doan, Thomas, Ronald Litterman, and Christopher Sims, 1984. Forecasting and Conditional Projection Using Realistic Prior Distributions. *Econometric Reviews* **3(1)**, 1-100.

Eichberger, Jrgen, and Ian Harper, 1997. *Financial Economics.* Oxford: Oxford University Press

Ericsson, Neill R., 2001. Forecast Uncertainty in Economic Modeling. In Understanding Economic Forecasts Edited by D. H. Hendry and N. R. Ericsson, Cambridge/Mass: Cambridge University Press, 68-92.

Fama, Eugene F., 1975. Short-term Interest Rates as Predictors of Inflation. *American Economic Review* **65(3)**, 269-282.

Fisher, Jonas D. M., Chin T. Liu and Ruilin Zhou, 2002. When Can We Forecast Inflation? Federal Reserve Bank of Chicago Economic Perspectives 1/02, Federal Reserve Bank of Chicago.

Gottschalk, Jan, and Susanne Brck, 2000. Inflationsprognosen fr den Euro-Raum: Wie gut sind P*-Modelle? *Vierteljahreshefte zur Wirtschaftsforschung* **69(1)**, 69-89.

Hallman, Jeffrey D., Richard D. Porter, and David H. Small, 1991. Is the Price Level Tied to the M2 Monetary Aggregate in the Long Run? *American Economic Review* **81(4)**, 841-858.

Issing, Otmar, and Karl-Heimz Tdter, 1995. Geldmenge und Preise im vereinten Deutschland. In Neuere Entwicklungen in der Geldtheorie und Whrungspolitik. Edited by R. Duwendag. Berlin: Duncker & Humblodt, 97-123.

Jordan, Thomas J. and Michel Peytrignet, 2001. Die Inflationsprognose der Schweizerischen Nationalbank. *Quartalsheft der Schweizerischen Zentralbank* **2**, 55-61.

Kozicki, Sharon, 2001. Why Do Central Banks Monitor So Many Inflation Indicators? *Federal Reserve Bank of Kansas City Economic Review* **3**, 5-42.

Lioui, Abraham, and Patrice Poncet, 2002. Revealing Inflation Expectations : Let the Market Do It. Working Paper, Bar Ilan University, Ramat Gan.

Litterman, Ronald B., 1986. Forecasting With Bayesian Vector Autoregressions - Five Years of Experience. *Journal of Business & Economic Statistics* **4(1)**, 25-38.

Lovell, Michael C., and Richard C. Vogel, 1973. A CPI-Futures Market. *Journal of Political Economy* **81(2)**, 1009-1012.

Mishkin, Frederic S., 1990. What Does the Term Structure Tell us About Future Inflation? *Journal of Monetary Economics* **25**, 77-95.

Nordhaus, William D., 1987. Forecasting Efficiency: Concepts and Applications. *Review of Economics and Statistics* **69**, 667-674.

Poloz, Stephen, David Rose, and Richard Tetlow, 1994. The Bank of Canada's new Quarterly Projection Model (QPM): An introduction. Bank of Canada Review, Autumn.

Ross, Stephen A., 1976. The Arbitrage Theory of Capital Asset Pricing. *Journal of Economic Theory* **13**, 341-360.

Smith, Vernon, 1982. Markets as Economizers of Information: Experimental Evidence on the "Hayek Hypothesis". *Economic Inquiry* **20**, 165-179.

Stock, James H., and Mark W. Watson, 1999. Forecasting Inflation. *Journal of Monetary Economics* **44**, 293-335.

Svensson, Lars E. O., 1997. Inflation Forecast Targeting: Implementing and Monitoring Inflation Targets. *European Economic Review* **41**, 1111-1146.

Svensson, Lars E. O., 1999. Inflation Targeting: Some Extensions. *Scandinavian Journal of Economics* **101(3)**, 337-361.

Tallman, Ellis W., 1995, Inflation and Inflation Forecasting: An Introduction. *Federal Reserve Bank of Atlanta Economic Review* **80(1)**, 13-27.

Thomas, Lloyd B., 1999. Survey Measures of Expected U.S. Inflation. *Journal of Economic Perspectives* **13(4)**, 125-144.

Thompson, Patrick A., and Robert B. Miller, 1986. Sampling the Future: A Bayesian Approach to Forecasting From Univariate Time Series Models. *Journal of Business & Economic Statistics* **4(4)**, 427-436.

Tdter, Karl-Heinz, and Hans-Eggert Reimers, 1994. P-star as a Link Between Money and Prices in Germany. *Weltwirtschaftliches Archiv* **130(2)**, 273-289.

Tversky, Amos, and Daniel Kahneman, 1981. The framing of decisions and psychology of choice. *Science* **211**, 453-458.

Webb, Roy H., 1999. Two Approaches to Macroeconomic Forecasting. *Federal Reserve Bank of Richmond Economic Quarterly* **85(3)**, 23-40.