КОНСОРЦИУМ ЭКОНОМИЧЕСКИХ ИССЛЕДОВАНИЙ И ОБРАЗОВАНИЯ
ECONOMICS EDUCATION AND RESEARCH CONSORTIUM
Working Paper Series

THE LINKS BETWEEN INFLATION AND INFLATION UNCERTAINTY AT THE LONGER HORIZON

Alexander Tsyplakov

Working paper No10/09E

This project (No 09-1121) was supported
by the Economics Education and Research Consortium
and funded by GDN

JEL Classification: E31, E52, C22, C29, C32

Tsyplakov A.A. The Links between Inflation and Inflation Uncertainty at the Longer Horizon. — Kiev: EERC, 2010. — 37 p.

In this paper I examine the Okun–Friedman hypothesis of the link between inflation and inflation uncertainty using historical international data on the monthly CPI. An indicator of inflation uncertainty at the two-years-ahead horizon is derived from a time-series model of inflation with time-varying parameters by means of Monte Carlo simulations. This indicator is compared to other uncertainty measures, with the short forecast horizon and based on simpler GARCH-type models. The analysis convincingly demonstrates that both the longer horizon and changing parameters are important for the regularity. The evidence obtained strongly supports the Okun–Friedman hypothesis both in the time dimension for most countries and across countries.

Key words and phrases: inflation uncertainty, inflation forecasting, Okun–Friedman hypothesis, nonlinear state space models, scoring rules.

Acknowledgments: Financial support from EERC (grant No. R09-1121) is gratefully acknowledged. I would like to thank sincerely the panel of experts and participants of EERC workshops for fruitful discussions. My special thanks are due to Oleksandr Shepotylo, Rustam Ibragimov, Vladimir Popov and Gary Krueger for their comments on the project and expert counsel. I am also indebted to Stanislav Anatoliev and Konstantin Gluschenko for their most helpful suggestions.

Alexander Tsyplakov
Department of Economics, Novosibirsk State University
Ul. Pirogova 2, Novosibirsk, 630090
Tel.: +7 (383) 363-42-12
e-mail: `tsy@academ.org`

CONTENTS

NON-TECHNICAL SUMMARY

An important question of the monetary theory is whether the rates of inflation are positively correlated with uncertainty about the future price level and whether there is a causal link between inflation and inflation uncertainty. Inflation uncertainty is said to entail major welfare losses (including distortion in relative prices, difficulties for decisions concerning long-term contracts and planning for future in nominal terms), so if inflation cause uncertainty then these losses should be an important part of the costs of inflationary monetary policy.

A hypothesis that higher inflation is related to greater inflation uncertainty was put forward by Arthur M. Okun and became well-known when Milton Friedman in his Nobel lecture based his critique of Phillips curve on the hypothesis. The argument was that at high rates of inflation macroeconomic policy would become less predictable leading to a greater volatility of price level. The dependence between inflation and its uncertainty have since become a textbook stylized fact. No serious discussion of the costs of inflation can escape mentioning of this dependence.

This study makes an advancement in the research on Okun–Friedman hypothesis. The results obtained in this study have implications not only for hypothesis itself, but also for related research of inflation modeling and forecasting.

I examine the link between inflation and inflation uncertainty using historical international data on monthly consumer price index. The research hypothesis is that the link between inflation and uncertainty may be stronger than it was found previously. The link is stronger if a more appropriate measure of inflation uncertainty is used. I propose to use uncertainty of two-years-ahead forecasts of price level which can be considered as a more adequate measure of inflation uncertainty than the measures used habitually in the previous inflation–uncertainty research.

Policy recommendations are well known and were made many times prior to this study. Institutional arrangements encouraging responsible behavior of the government and monetary authorities in the area of fiscal and monetary policy play the key role in overcoming inflation. The problem is that existing empirical evidence on Okun–Friedman hypothesis cannot be considered as fully conclusive. Some of the previous research found weak or not very robust links between inflation and uncertainty about future prices. The goal is to provide a more sound empirical basis for popular policy recommendations.

The study suggests convincing evidence that there exists a strong positive relation between inflation and its uncertainty both *across countries* and in the time series dimension *within countries*. There is also time-series evidence that inflation precedes uncertainty.

A longer-horizon inflation uncertainty (at least one year ahead, two years in this study) is more relevant to the original concerns behind the Okun–Friedman hypothesis. Notable is the fact that empirically it is also more closely related to the level of inflation than a short-horizon inflation

uncertainty (one month in this study).

Of course, correlation is not always an evidence of causation and precedence is not an evidence of the direction of causation. Thus, too confident policy-related judgments cannot be made at the current stage of the research. However, one can argue that this study is an important step in the right direction.

Our results show that the old argument by Arthur Okun is valid. High inflation and low inflation uncertainty are largely incompatible. Using Okun's words, “the adoption of a public policy designed to yield steady, fully anticipated inflation would commit the government to an impossible goal”. Some policy implications of the result do not depend on the direction of causation. These can be stated in terms of the construction of institutions for the economic policy. We do not have a choice between institutions which encourage low inflation and low inflation uncertainty and institutions which encourage high inflation and low inflation uncertainty. High inflation always comes together with high inflation uncertainty. Thus, institutions which encourage high inflation would lead to inflation uncertainty.

1. INTRODUCTION

Okun (1971), Friedman (1977) stated a hypothesis that there is a relationship between rates of inflation and inflation uncertainty. From an economic policy perspective it is important that it is almost never possible to have high average inflation and low inflation uncertainty simultaneously (cf. Okun (1971)). If a policy-maker seeks to maintain low inflation uncertainty, but at the same time conducts an expansionary monetary policy then he follows an unachievable goal or "mirage" as Okun called it. Moreover, if high inflation uncertainty is caused by high inflation then this relationship can be a major source of the costs associated with high inflation, because inflation uncertainty is said to entail substantial welfare losses (including distortion in relative prices, difficulties for decisions concerning long-term contracts and planning for future in nominal terms).

In this study I explore the Okun–Friedman hypothesis empirically using historical international data on the monthly CPI. There are two main challenging problems here which could be attended. The first one is to verify empirically whether such regularity really exists and measure its strength. The second one is to find out what could be the mechanisms connecting the inflation level and inflation uncertainty. This includes inferring the directions of causal relations and verifying existing theories.

Although the second problem seems to be the most intriguing one and has the greatest practical significance, solving it could be an utterly difficult task as it would include disentangling consequences of multiple causes and taking apart theories which are almost observationally equivalent. At the same time the first problem is simpler and more feasible, while it is not devoid of an important policy-related content. Moreover, it implies solving methodological subproblems which can pave ground for a more sound investigation of the second problem.

This study relates to the first empirical problem. Its solution depends on the definition of inflation uncertainty. The existing literature suggests several approaches to measuring of inflation uncertainty and most of them have serious drawbacks. Cross-country studies starting from Okun (1971) use indicators which are only indirectly connected with inflation uncertainty. This raise doubts about their relevance to the problem. Time-series studies based on modeling of conditional heteroskedasticity starting from Engle (1983) use the short-run volatility from a GARCH-type model which can lead to an underestimation of the strength of inflation–uncertainty links. My contention is that the links are stronger if a more appropriate measure of inflation uncertainty is used.

Two issues are addressed here. First, a 2-years forecast horizon is chosen for the study. I think that it is more adequate than the shorter horizon measures used in much of the previous research in the time dimension. Second, in this study time-series models of inflation with time-varying parameters (AR or VAR) are utilized for the goal of measuring inflation uncertainty. Such a model can provide a reliable time-varying indicator of inflation uncertainty. The main idea is that in order to make a good forecast it is necessary to account for unobservable factors driving inflation. One should not only infer the current state of the unobservable factors, but also take into consideration possible changes in that state in the future. These changes would have only a small impact on short-term forecasting, but their contribution to the longer-horizon forecast uncertainty can be large.

I construct an indicator of inflation uncertainty, compare its time-series behavior with the behavior of inflation and analyze the role of the horizon and changing parameters. I also explore the link in the cross-country dimension and compare the average levels of inflation and inflation uncertainty in a sample of countries. It is shown that the length of the forecast horizon and changes in parameters are important considerations.

In stressing the role of the horizon and changing parameters this study follows Evans (1991). It can be viewed as a refinement and extension of Evans (1991) relying on some useful recent econometric techniques, with a more clear and universal definition of inflation uncertainty. Clearly stated definition based on an explicit forecasting framework is important as it allows to subject the chosen indicator to a series of procedures in order to verify its validity. In addition, this paper is based on much broader empirical data and introduces a cross-country dimension to the research which is absent in Evans (1991) (as well as in GARCH-based studies).

The paper is organized as follows. Section 2 reviews the inflation–uncertainty literature and discusses the problem of measuring inflation uncertainty. It also introduces the basic univariate model of inflation. In Section 3 the adequacy of the basic model is checked by assessing forecast calibration and comparing forecasts with some natural alternatives. Section 4 gives an account of the main empirical findings both in the time-series and cross-country dimensions. In Section 5 conclusions are made and some thoughts about the meaning of the results are put forward.

2. RESEARCH BACKGROUND, BASIC MODEL AND DISCUSSION OF INFLATION UNCERTAINTY

2.1. Inflation–uncertainty literature: theoretical arguments

A hypothesis that higher inflation is related to greater inflation uncertainty was advanced in Okun (1971) and became well-known when Milton Friedman in his Nobel lecture Friedman (1977) based a critique of Phillips curve on the hypothesis. The argument was that at the high rates of inflation macroeconomic policy would become more erratic leading to a greater volatility of the price level. Both Okun and Friedman provided only informal discussion.

Several theoretical models predicting links between inflation uncertainty and rates of inflation can be found in the literature. Most of the theories emphasize incentives of policy-makers and assume that the public understands these incentives (that is, the Barro and Gordon framework with a time-consistent discretionary monetary policy is employed; Barro and Gordon (1983a), Barro and Gordon (1983b)).

In a model due to Cukierman and Meltzer (1986) the public cannot directly observe policy targets and policy-makers can impair the precision of monetary control which allows to make nominal surprises a more efficient instrument for stimulating economic activity. Under greater political instability policy-makers tend to conduct both more erratic and more inflationary monetary policy. A model in Ball (1992) formalizes the arguments by Okun and Friedman. There are different types of policy-makers which randomly come into power. Nobody wants to increase inflation when it is low. However, when inflation is high the incentives to disinflate depend on policy-maker's type. In Geraats (2006) the Barro and Gordon setup is modified by assuming an asymmetric objective function for the monetary policy. As a result, when the supply shock is negative it increases both the level and variability of inflation.

Holland (1993) provides a technical explanation in the context of an econometric model in which the effect of the monetary policy on economy is uncertain. To some extent the idea applies to the basic model (2) analyzed below (cf. Holland (1995) where a similar observation is made about the model of Evans and Wachtel (1993)).

There are also theoretical explanations which are less universally applicable. In Devereux (1989) the high variance of the real shocks both reinforces incentives to produce nominal surprises (as long as the degree of wage indexation decreases) and leads to higher inflation uncertainty. Mondino, Sturzenegger, and Tommasi (1996) formalized political cycles observed in high-inflationary Latin American countries.

2.2. Inflation–uncertainty literature: empirical evidence

Two main strands dominate in the substantial literature which aims to explore the link between inflation and inflation uncertainty empirically. The older strand deals with cross-country evidence and is based on approximating inflation uncertainty by some simple inflation variability. The more recent strand of literature is based on GARCH-type time-series models and uses the corresponding conditional variance as a measure of inflation uncertainty.

Studies using cross-country data mostly confirm the existence of the relation between the level and uncertainty of inflation (for example, see Okun (1971), Logue and Willett (1976), Foster (1978), Gale (1981), Ram (1985), Hess and Morris (1996), Fang, Miller, and Yeh (2007)). In these studies simple inflation variability over some fixed period is typically used as a measure of uncertainty in a country (for example, the sample standard deviation or the mean absolute change). In general the evidence is that there exists a significant link between the average level of inflation in a country and inflation variability.

The main objection to this simple cross-country approach is that variable inflation is not necessary uncertain or unpredictable. (The point was stressed in Driffill, Mizon, and Ulph (1990) among others). There is no one-to-one correspondence between the unconditional variability of a stochastic process and its predictability. However, the correlation between the two can be quite close for a narrow class of time-series models. Within such a class the unconditional variability can be akin to uncertainty. Thus, arguably, the available cross-country evidence provides a serious support for the Okun–Friedman hypothesis. Moreover, the evidence is robust to the use of more realistic forecasts. For example, Davis and Kanago (1998), Davis and Kanago (2000) used published professional forecasts and found a significant positive cross-country relationship between average inflation rates and "standard deviations around the forecasts".

The same link between inflation and inflation variability is observed to some extent in the time-series dimension of the data if a long time series of inflation is divided into smaller intervals or rolling estimation is used (for example, Gale (1981), Ram (1985), Ball and Cecchetti (1990)).

Robert Engle's seminal paper on the variance of inflation in the United States Engle (1983) initiated the steadily growing strand of literature in which inflation uncertainty is studied by means of explicit time-series models of conditional variance. He modeled the first conditional moment of inflation as an autoregression with additional lagged variables (wages, money, import deflator) and the second conditional moment as an ARCH process. Engle applied this model to quarterly U.S. inflation data and used the conditional variance as a proxy for inflation uncertainty. He was not able to find a significant relation between inflation and the inflation volatility. Although Engle (1983) obtained negative evidence, further time-series studies of the links between inflation and the short-run inflation volatility provided empirical results somewhat more favorable to the Okun–Friedman hypothesis.

Studies of the relation between inflation and its uncertainty using GARCH-type models are numerous and use data for different countries. For example, Brunner and Hess (1993) (who suggested to take into account the asymmetry in the influence of shocks on the volatility), Baillie, Chung, and Tieslau (1996) (who used the ARFIMA-GARCH model featuring long memory), Kontonikas (2004), Fountas, Karanasos, and Kim (2006), Henry, Olekalns, and Suardi (2007). Probably the most representative study of this kind in terms of the number of countries is Daal, Naka, and Sanchez (2005). The study used CPI from 23 countries to estimate a model of inflation including the asymmetric power GARCH model for conditional variance. The authors obtained results which strongly suggest that inflation influences inflation uncertainty. Thornton (2007) studied 12 emerging market economies. He also obtained a strong empirical support for the hypothesis that inflation raises uncertainty.

Two major problems with the GARCH-based approach to the study of the inflation–uncertainty links are evident. First, the implicit forecast horizon is inadequately short. Typically in a GARCH-based study $\mathsf{var}(\pi_{t+1}|\Omega_t)$ (the conditional variance or the volatility) is used as a measure of inflation uncertainty. (Here Ω_t is information available at time t.) The conditional variance $\mathsf{var}(\pi_{t+1}|\Omega_t)$ coincides with the one-step-ahead forecast variance. Since the one-step horizon is one quarter for quarterly data and only one month for monthly data this does not fit well with the original concerns behind the Okun–Friedman hypothesis. A longer horizon (at least one year) would be more appropriate. The available empirical evidence also suggests that considering long-term uncertainty is important; see Ball and Cecchetti (1990), Evans (1991), Evans and Wachtel (1993), Kim (1993).

The second serious shortcoming is that the GARCH-based approach does not fully take into account that model parameters can be time-varying. Time-varying conditional variance is captured by the GARCH component. However, other parameters can also change (trend inflation, inflation persistence, see below).

The importance of time variation in parameters was demonstrated by the studies of the inflation–uncertainty links based on time-series models with time-varying and switching parameters. For example, Evans (1991) augmented the GARCH modeling framework by assuming the parameters of his inflation model to be time-varying. Thus, he was able to distinguish several types of inflation uncertainty. He found strong correlation between inflation and the long-term uncertainty in the United States. Evans and Wachtel (1993) and Kim (1993) introduced a Markov-switching component into a time-series model of inflation to imitate regime shifts.

A time-series study can also be based on survey and/or professional forecasts measures. See Wachtel (1977), Pagan, Hall, and Trivedi (1983), Zarnowitz and Lambros (1987), Ungar and Zilberfarb (1993), Holland (1995), Davis and Kanago (1998), Davis and Kanago (2000).

An annotated bibliography on inflation–uncertainty links can be found in Davis and Kanago (2000). Golob (1993) is an older review.

2.3. Inflation forecasting and inflation uncertainty

Inflation uncertainty can be loosely defined as subjective unpredictability of the future price level. Unpredictability (forecast uncertainty) is some measure of vagueness or spread of the predictive density.

Different people have different subjective forecast distributions because of unequal possession of information and differences in forecasting skills. However, to test the Okun–Friedman hypothesis one needs an integral measure. Inflation uncertainty should be some aggregate of forecast uncertainties among economic agents. It is not obvious which method of aggregation is appropriate for the Okun–Friedman hypothesis. One approach would be to consider the average uncertainty among all economic agents who make forecasts, either skillful or not, as an ideal indicator. Another extreme is to consider the forecasts only of the most skillful forecasters as an ideal indicator.

Another point is that the price level is a subjective notion as the relevant commodity basket is individual for each economic agent. This means that the Okun–Friedman hypothesis is not only about the general price level, but also relates to price dispersion across commodity groups and space.

However, the current study concentrates on the problem of unpredictability of the general price level only. The relative price dispersion aspect of the Okun–Friedman hypothesis is an important problem in its own right and needs a separate thorough investigation.

The approach of this study is to take a reasonably good time-series model and base a measure of inflation uncertainty on predictive distributions from this model. The idea is that economic agents can produce good enough forecasts when they are to make important inter-temporal decisions and that this can be approximated by a good econometric model, however tricky required estimation and inference could be. (This is similar to the observation that a football player is able to make a complicated forecast of the ball's trajectory without consciously resorting to the physical laws of motion). If we identify inflation uncertainty with a measure of forecast uncertainty of inflation then we have to choose a model which produces adaptively varying predictive distributions. Point forecasts are unsuitable for the task as they lack associated uncertainty.

The notion of forecast uncertainty or unpredictability is not something definite.[1] One measure is the variance or standard deviation of predictive distribution. Another natural measure is the width of the forecast interval. In what follows I take the width of the central 90% forecast interval as the measure of forecast uncertainty.[2] The width of the forecast interval is defined as the distance between the suitable quantiles of the predictive distribution.

The next important point is that multistep forecasting is needed in order to study the significance of the length of the forecast horizon. Given information Ω_t available at time t the forecast is made H periods ahead for period $t + H$. Below the forecasts of the price level based on monthly data with the horizons of 1, 24 and 14 months are considered.

Moreover, it is necessary to draw attention to the fact that in an econometric model the inflation rate[3] $\pi_t = \ln P_t - \ln P_{t-1}$ would be modeled rather than the price level P_t (or $\ln P_t$) which is clearly non-stationary. In producing H-step-ahead forecasts we are not concerned with the forecasts about π_{t+H} which is the rate of change of the price level from time $t + H - 1$ to time $t + H$, but rather with the forecasts about $\sum_{i=1}^{H} \pi_{t+i} = \ln P_{t+H} - \ln P_t$ which is the rate of change of the price level H periods ahead or (equivalently) with the forecasts about $\frac{1}{H} \sum_{i=1}^{H} \pi_{t+i}$ which is the average inflation rate H periods ahead.

In measuring uncertainty parameter estimation errors and uncertainty about the form of the model (for example, see Garratt, Lee, Pesaran, and Shin (2003), Wallis (2008)) are not taken into account in this study. However, the basic model used here allows to capture structural shifts in the economy (cf. Cogley and Sargent (2001)).

[1] Various ways of reporting forecast uncertainty are discussed in Wallis (2008).

[2] The term *central interval* means that tail probabilities are taken to be equal. For an interval with the 90% confidence level this attributes 5% probability to each tail.

[3] Below actually $\ln P_t - \ln P_{t-1}$ multiplied by 100 is used as π_t.

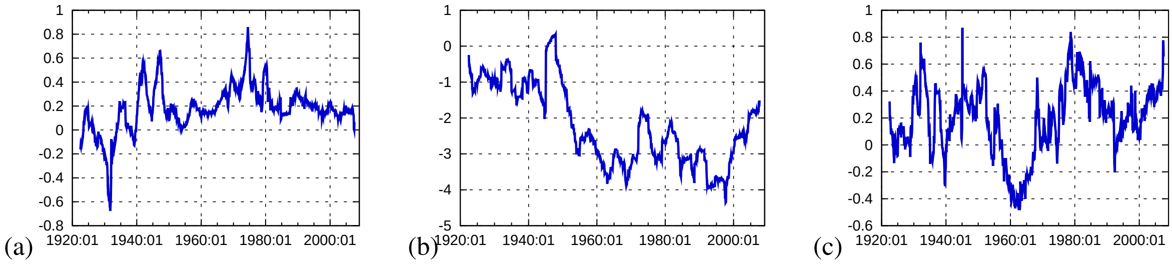

Fig. 1. Rolling AR(1) estimates for the United States, 1921:01–2009:01; (a) c, (b) $\ln \sigma$, (c) φ.

In summary, I estimate a model of inflation, assume that the estimated parameters are true parameters and measure inflation uncertainty as the width of the central 90% forecast interval for $\frac{1}{H}\sum_{i=1}^{H}\pi_{t+i}|\Omega_t$ where Ω_t is information used for forecasting at time t.

2.4. TV-AR model of inflation

Varying or switching coefficients are frequently mentioned in the literature on inflation modeling. For example, changing autoregressive coefficients is a phenomenon mentioned in Stock and Watson (2007). This can be interpreted as changing inflation persistence. Following Evans (1991) I expect that this can be an important source of long-run inflation uncertainty and that not taking this into account can lead to a distortion and underestimation of Okun–Friedman effect.

The phenomenon of time-varying parameters can be illustrated by a simple AR(1) model

$$\pi_t = c + \varphi\pi_{t-1} + \varepsilon_t, \qquad \mathrm{var}(\varepsilon_t) = \sigma^2. \tag{1}$$

fitted to inflation series over rolling intervals. Here $\pi_t = \Delta p_t = \Delta \ln P_t$ is the inflation rate obtained from price index P_t. Rolling first-order autoregressions are fitted to the monthly inflation data for the United States for the period 1921:01–2009:01 (Appendix B, Series A). The length of the rolling interval is chosen to be 35 months. That is, the autoregression (1) is estimated for the moment τ by OLS using only observations $t = \tau - 17, \ldots, \tau + 17$. The estimates $\hat{c} = \hat{c}(\tau)$, $\hat{\sigma} = \hat{\sigma}(\tau)$ and $\hat{\varphi} = \hat{\varphi}(\tau)$ are attributed to moment τ.

Figure 1 records the estimated coefficients. The estimates show that all three parameters change appreciably over time. The autoregressive coefficient is visibly negative in 1957–1967. The pre-WWII inflation data show a considerably higher volatility. Also notable are the negative constant term estimates corresponding to the disinflation episode of the early 1930s.

Ball and Cecchetti (1990) used a similar non-parametric approach to the task of estimation of changing parameters. They divided quarterly time series into five-year intervals which can be viewed as a kind of regressogram estimation (bin smoothing).

The phenomenon of parameter instability is important for studying inflation–uncertainty links. The reason for this is that not only changing variance (σ^2), but also time variations of other underlying parameters can contribute to the behavior of the long-horizon uncertainty (cf. Evans (1991), Lendvai (2006)). Cogley and Sargent (2001) note that "the well-known positive correlation between the mean and variance of inflation reflects an even stronger correlation between the mean and degree of persistence".

These observations motivate the choice of the basic model for this study. The simple first-order autoregression (1) can be modified by assuming that each of its three parameters is governed by its own AR(1) process. Positivity of "variance" variable σ^2 is ensured by taking logarithms. This leads to the following AR(1) model with time-varying parameters (TV-AR):

$$\begin{aligned} &\pi_t = c_t + \varphi_t \pi_{t-1} + \varepsilon_t, \qquad \varepsilon_t = e^{h_t/2}\xi_t, \qquad \xi_t \sim skt_{\nu,\lambda}, \\ &\Lambda_{l,t} = (1-\delta_l)\omega_l + \delta_l \Lambda_{l,t-1} + \sigma_l \sqrt{1-\delta_l^2}\,\eta_{l,t}, \qquad \eta_{l,t} \sim N(0,1), \\ &\quad \text{where } \Lambda_{l,t} = c_t, h_t, \varphi_t \text{ for } l = 1,2,3. \end{aligned} \tag{2}$$

Here c_t can be interpreted as the trend contribution to inflation, φ_t as the inflation persistence, and h_t as the short-run volatility. The disturbances ξ_t, $\eta_{l,t}$ are assumed to be independent white noise series. I employ asymmetric and fat-tailed disturbances ξ_t distributed according to the skewed Student's t distribution with ν degrees of freedom and asymmetry parameter $\lambda \in (0,1)$. The distribution is parametrized in such a way that $\mathsf{E}\,\xi_t = 0$ and $\mathsf{var}\,\xi_t = 1$ (see Hansen (1994)).

Two papers which used similar approaches to the inflation–uncertainty links were sources of inspiration for the basic model. First, the approach of this study is close to that of Ball and Cecchetti (1990). However, Ball and Cecchetti used a very basic time-series model of inflation, that is, IMA(1) model. By using a simple IMA(1) model Ball and Cecchetti were able to readily decompose inflation into two components, a permanent (random-walk) component and a transitory (white-noise) component. The variance of random-walk error can be viewed as a measure of long-term uncertainty as it is the dominant component of long-term forecast error variance.

An important difference with the current study is that Ball and Cecchetti were interested in the error variance of long-term forecasts of one-period (quarterly) inflation while the unpredictability of average inflation H periods ahead is a more relevant indicator.

Ball and Cecchetti (1990) only partially took into account the fact that model parameters are time-varying. For the cross-country regressions they used the IMA(1) with fixed parameters for the whole data period. Then they divided inflation series into 5-year periods, fitted the IMA(1) with fixed parameters for these shorter series and used the results in panel regressions. However, gradually changing parameters can provide a very important contribution to the long-term inflation uncertainty. Most importantly, variances can change over time. Thus, Ball and Cecchetti (1990) provide very rough approximations to the desired uncertainty measure.

Another essential example of a previous research which considers the difference between the short- and long-run inflation–uncertainty links is Evans (1991). Evans allowed for both time-varying error variance (changing short-run volatility) and time-varying model coefficients. His model is an autoregression in which the autoregression coefficients are modeled as random walks and the disturbances are represented as an ARCH process. The study was confined to the United States data, so it is not clear whether the findings relate to other countries.

The long-horizon measure of inflation uncertainty in Evans (1991) assumes the infinite horizon and is based on heuristic considerations. In terms of the model (2) it would correspond to $\mathsf{var}(c_t/(1-\varphi_t)|\Omega_t)$.

Other sources of ideas for this study belong to a more general literature on inflation modeling and forecasting. Most important are the papers by Stock and Watson on inflation forecasting and Cogley and Sargent (2001). In particular, Stock and Watson (2007) used an unobserved components model with stochastic volatility to forecast inflation. Cogley and Sargent (2001) modeled the joint movement of inflation and other macroeconomic series by means of time-varying VAR.

The model (2) is a kind of nonlinear state space model. Such a model includes observed series $\boldsymbol{Y} = (\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_T)$ and unobserved state series $\boldsymbol{\Lambda} = (\boldsymbol{\Lambda}_1, \ldots, \boldsymbol{\Lambda}_T)$. For a given vector of parameters $\boldsymbol{\theta}$ a state space model determines the transition density $f(\boldsymbol{\Lambda}_t|\boldsymbol{\Lambda}_{t-1}, \boldsymbol{\theta})$ and the measurement (observation) density $f(\boldsymbol{Y}_t|\boldsymbol{Y}_{t-1}, \boldsymbol{\Lambda}_t, \boldsymbol{\theta})$. These components make up the joint density of $\boldsymbol{Y}$ and $\boldsymbol{\Lambda}$, the density of the complete data. However, for the maximum likelihood estimation one needs the marginal density $f(\boldsymbol{Y}|\boldsymbol{\theta})$ rather than the density of the complete data $f(\boldsymbol{Y}, \boldsymbol{\Lambda}|\boldsymbol{\theta})$. In general the former density is not known in a closed form. Consequently both estimation of a nonlinear state space model and obtaining density forecasts from it require computer-intensive algorithms like Monte Carlo integration. Most popular methods for nonlinear state space models are based on approximating the unknown conditional ("posterior") distribution with density

$$f(\boldsymbol{\Lambda}|\boldsymbol{Y}, \boldsymbol{\theta}) = f(\boldsymbol{Y}, \boldsymbol{\Lambda}|\boldsymbol{\theta})/f(\boldsymbol{Y}|\boldsymbol{\theta})$$

by a simpler Gaussian distribution $g(\boldsymbol{\Lambda}|\boldsymbol{\theta})$.

There are many ways to estimate the model (2). The results in this study were obtained by means of efficient importance sampling (EIS) algorithm with 100 simulations and 10 iterations (see details in Appendix A).

Table 1 reports the estimates for the United States for the period 1921:01–2009:01. Also the estimates were obtained for the main data-set. The source of the data is International Financial Statistics database from IMF (Appendix B, Series B).[4] The inflation rates were calculated from the

[4] The sample does not include Brazil, Argentina and Peru. CPI series for these countries exhibit an abrupt behavior due to outbursts of very high inflation rates and monetary reforms. The algorithms used failed to converge for the countries. A more complicated model with jumps could be needed to describe inflation in these countries. Moreover,

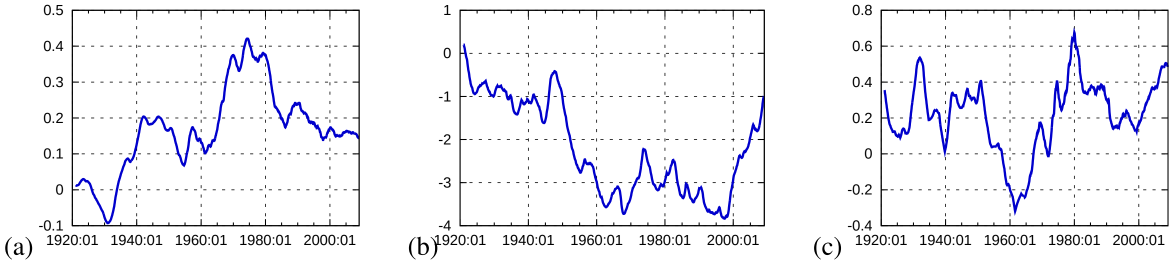

Fig. 2. Model estimates for the United States, 1921:01–2009:01, posterior mode; (a) $\hat{c}_t$, (b) $\hat{h}_t$, (c) $\hat{\varphi}_t$

monthly CPI series for the period 1957:02–2009:01. (For some countries the series are shorter.). The series were seasonally adjusted by means of Census X12 automatic adjustment. Table 1 reports the estimates for the U.S., Japan, Greece and Egypt. The estimates of deltas for the model are close to unity, which explains the wide use of random walk in inflation modeling. For a sizable proportion of the countries some of the three deltas degenerate, tending from 1 towards 0 during iterations, while the corresponding sigmas can tend towards 0. For such factors δ_l was fixed at 0.99 (and some σ_l at a negligibly small level 10^{-5}) so that the factor remained practically constant. In Table 1 this is the case for Japan (fixed parameters are marked with F), for which the autoregressive coefficient is almost constant around 0.021.

Table 1. Estimates of the TV-AR model (2) for CPI inflation in four countries

Parameter	U.S. 1921:01–2009:01	U.S. 1957:02–2009:01	Japan 1957:02–2009:01	Greece 1957:02–2009:01	Egypt 1957:02–2009:01
ω_c	0.149	0.200	0.183	0.416	0.494
δ_c	0.993	0.983	0.994	0.995	0.995
σ_c	0.131	0.071	0.237	0.412	0.377
ω_h	-1.660	-3.590	-1.895	-1.990	-0.443
δ_h	0.995	0.933	0.996	0.925	0.991
σ_h	1.260	0.947	1.000	0.909	1.524
ω_φ	0.233	0.292	0.021	0.060	0.013
δ_φ	0.981	0.986	0.99^{F}	0.981	0.728
σ_φ	0.241	0.184	$10^{-5\,F}$	0.178	0.206
ν	7.892	10.28	4.730	4.462	3.089
λ	0.135	0.058	0.156	0.125	0.134

The three graphs in Figure 2 show estimates of the three unobserved factors (time-varying parameters) from the model (2) for the United States for the period 1921:01–2009:01. The estimates were obtained as the mode of the posterior distribution. These are similar to the rolling estimates above.

The estimates of uncertainty used below are based on forecasts. All the forecasts here (including those used in forecast comparisons and diagnostics) are quasi out-of-sample ones because future observations are used in the estimation of the model parameters. Truly out-of-sample forecasts calculated on the basis of available observations only are not used in this study because obtaining

it is necessary to take into account anticipated monetary reforms to estimate inflation uncertainty adequately.

them is very time-consuming.

For complicated enough models such as (2) it is impossible to find closed-form expressions for characteristics of the predictive distribution. Non-linearity of the model, time-varying parameters and non-Gaussian disturbances all add to this difficulty. Additionally, forecasts of the average inflation rate during H periods rather than of the monthly inflation H periods ahead are needed. Thus, it is necessary to resort to a numerical approximation by some Monte Carlo method.

The forecast distribution is approximated by a discrete distribution of future trajectories with corresponding importance weights. Below the EIS method was used with 1000 simulations and 10 iterations to fit a proposal distribution. For filtering 100000 initial Monte Carlo trajectories were used, which were then reduced to 1000 trajectories with equal importance weights via systematic resampling. The methods are described in Appendix A.

3. FORECAST VERIFICATION

3.1. Discussion of forecast verification

It is important that the model used for evaluating uncertainty provides high-quality density forecasts. There are two principal ways to verify forecast quality. One way is to conduct diagnostics and assess forecast calibration. Another way is to compare the forecast procedure with alternative procedures in terms of forecast accuracy. Discussion of these quality verification methods can be found in Diebold, Gunther, and Tay (1998), Wallis (2008) and Gneiting, Balabdaoui, and Raftery (2007).

As discussed in Gneiting, Balabdaoui, and Raftery (2007) a high-quality density forecast must possess two important features: calibration and sharpness. Calibration refers to the property of a forecast to adequately capture the behavior of the data under consideration. Sharpness refers to a high concentration of the forecast density. A good density forecast should have small uncertainty associated with it (be sharp). On the other hand, it should be in agreement with the observed data (be properly calibrated), in particular, it should not be misleading with respect to the associated uncertainty and should not be biased.

One way to judge the adequacy of calibration is to use the cumulative distribution function of the predictive distribution to transform the actual observations to $U[0,1]$ distributed variables. This is called the probability integral transform (PIT) and is discussed below.

Sharpness can be measured roughly by the average forecast uncertainty, for example, by the average width of forecast intervals. However, one should be cautious as too sharp a forecast can be badly calibrated due to under-reported uncertainty. Thus, it is better to have indicators which combine both sharpness and calibration when comparing density forecasts. Such indicators are called scoring rules. A comprehensive discussion of scoring rules can be found in Gneiting and Raftery (2007).

One such integral measure of forecast quality for density forecasts is the continuous ranked probability score (CRPS). It is known to be a strictly proper scoring rule, which means that it stimulates truthful forecasts statement (or penalizes any tweaking in the statement of one's forecast density). Among other things, it penalizes biased forecasts, or too sharp or too fuzzy forecasts. CRPS is not the only strictly proper scoring rule, but it is one of the most appealing. It is not that popular, the reason for which can be that it is not readily available for a complicated forecast distribution. However, if one has a Monte Carlo sample from the forecast distribution one can obtain a sample estimate of CRPS by replacing theoretical expectations by the corresponding sample means. Suppose there is a random sample q^s, $s = 1, \dots, S$ from predictive density $\tilde{f}$ with associated normalized importance weights w^s. In terms of the ordered sample $q^{(1)} \leq q^{(2)} \leq \cdots \leq q^{(S)}$ CRPS can be estimated as

$$CRPS(\tilde{f}, q) \approx \sum_{s=1}^{S} w^{(s)} |q^{(s)} - q| - \sum_{s=1}^{S} w^{(s)} \Big(2 \sum_{r=1}^{s} w^{(r)} - 1 - w^{(s)} \Big) q^{(s)}.$$

CRPS is used here in the negatively oriented form; this is convenient, as negatively oriented CRPS is positive and specializes to the popular absolute error criterion function for a degenerate single-point forecast distribution. To compare different forecasting techniques one should use average values of CRPS for some period $t = T_1, \dots, T_2$:

$$\overline{CRPS} = \frac{1}{T_2 - T_1 + 1} \sum_{t=T_1}^{T_2} CRPS(\tilde{f}_t, q_t).$$

Below the following procedures are conducted to verify forecasts from the basic model. First, the behavior of the TV-AR model is compared with the Livingston survey for the U.S. case. Second, it is checked whether addition of other macroeconomic indicators can seriously influence the proposed measure of uncertainty by comparing forecasts from the basic model with forecasts from a more complicated time-varying vector autoregressive model. Third, the basic model is compared with a simpler time-series model (ARMA-GARCH). Forth, PIT-based diagnostics for the density forecasts from the basic model are inspected.

3.2. Comparison with Livingston survey for the U.S.

It is important to compare the proposed measure of inflation uncertainty with alternatives. One rough check can be performed by comparing it with surveys of inflation forecasts. Figure 3 compares the uncertainty estimates for the United States from the model (2) with the 1-year-ahead CPI predictions from the semi-annual Livingston survey for the period 1948:06–2009:01 (Appendix B,

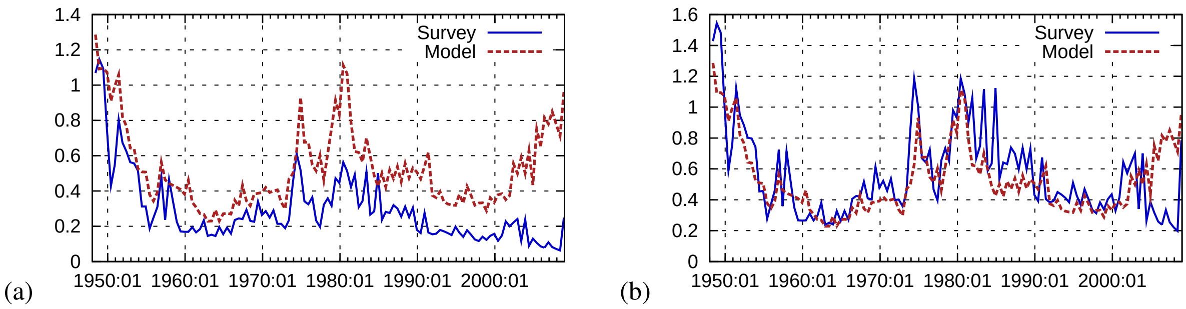

Fig. 3. Measures of uncertainty based on the TV-AR model (2) and on Livingston Survey, $H = 14$, U.S., 1948:06–2009:01: (a) with initial survey measure, (b) with recalibrated survey measure.

Series C). The two uncertainty indicators were made comparable. The fourteen-month horizon ($H = 14$) was used for the model-based indicator. This length of the forecast horizon is an approximation; see Carlson (1977) on the horizon implicit in the 1-year-ahead CPI forecasts and on other problems of the Livingston survey data. For the Livingston Survey data I considered samples of survey forecasts at different times periods as forecast distributions and computed the differences between sample 95 % and 5 % quantiles. (On average there were about 47 participants). The survey-based indicator is sharper (except for the starting years) as evident from Figure 3(a). The similarity between the two indicators is noticeable, although there is a huge discrepancy at the end of the period.

Forecasters' disagreement is not the same as inflation uncertainty. It is clear that such an indicator does not reflect the full forecast uncertainty as each forecaster provides only some plausible point value without attaching his intrinsic subjective ambiguity to the value. Although logically inflation uncertainty is not the same thing as disagreement of inflation forecasts, empirically these indicators can behave similarly. For more in-depth discussion of the issues see Zarnowitz and Lambros (1987) and Bomberger (1996).

It can be conjectured that there was a gradual decline in forecasters' disagreement relative to inflation uncertainty. To test the conjecture a recalibrated version of the survey measure of uncertainty was computed by minimizing CRPS. For each time period the distribution of forecasts was scaled around the median. The logarithm of the scaling factor was assumed to change according to a linear trend. The coefficients of the trend were chosen by minimizing average (negatively oriented) CRPS. Figure 3(b) shows the result. The greatest difference is at the end of the period. The survey measure is much lower since 2004 or 2005. During the last years of the period inflation rates show a high short-run volatility with a not so high average rate ($\approx 0.2\,\%$ monthly inflation) and an increased agreement between the forecasters. However, in general it is striking that two uncertainty measures obtained by very different methods can be so similar.

Long enough surveys of inflation forecasts are not available for most of the countries so it is not possible to use them systematically. Moreover, this subject deserves a separate study. For the current study a rough check is enough. From the comparison it can be concluded that the proposed measure of inflation uncertainty is not completely amiss and that it captures some important aspects of inflation uncertainty in the U.S.

3.3. Comparison with TV-VAR models

The main concern about the basic model (2) is not that it is purely technical and unrelated to the macroeconomic theory. Actually, it is able to capture regime shifts in monetary policy which are considered as an important feature in the recent theoretical and empirical literature on inflation. The main concern is that it is univariate. It does not take into account all information which is relevant to inflation forecasting and can be available to a forecaster at the time he makes his forecast. Of course, some additional information other than historical records of CPI can improve forecasting of the future CPI. Thus, a forecast interval from a univariate model is a rough approximation to some imaginary ideal measure of unpredictability.

In view of potential distortions it is necessary to check whether my measure based on a univariate time-series model is appropriate and whether additional macro variables can change the behavior of the model-based forecast uncertainty considerably. It can be that multivariate models of inflation are only marginally better than a good univariate model; see Atkeson and Ohanian (2001), Stock and Watson (2007), Stock and Watson (2008), Ang, Bekaert, and Wei (2005), Doyle (2006). Below I examine whether this is indeed the case. A forecast uncertainty measure which takes into account information on several variables is obtained from a vector autoregression with time-varying parameters (TV-VAR). It is similar to the model of inflation in Cogley and Sargent (2005) which extends the model in Cogley and Sargent (2001) by including stochastic volatility.

Consider the following TV-VAR model:

$$\boldsymbol{Y}_t = \boldsymbol{M}_t + \sum_{j=1}^{p} \boldsymbol{Y}_{t-j}\boldsymbol{\Pi}_{j,t} + \boldsymbol{U}_t\boldsymbol{D}_t\boldsymbol{C}_t, \qquad \boldsymbol{U}_t \sim \mathcal{N}(\boldsymbol{0}_k, \boldsymbol{I}_k). \tag{3}$$

Here $\boldsymbol{Y}_t$ is a row vector of modeled variables at time t, $\boldsymbol{M}_t$ is a time-varying vector of constant terms, $\boldsymbol{\Pi}_{jt}$ are matrices of autoregressive coefficients,

$$\boldsymbol{C}_t = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ c_{21,t} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ c_{k1,t} & \cdots & c_{k,k-1,t} & 1 \end{pmatrix}$$

is a lower triangular matrix which determines correlations between disturbances,

$$\boldsymbol{D}_t = \operatorname{diag}(\exp(h_{1,t}/2), \ldots, \exp(h_{k,t}/2))$$

Table 2. Comparison of the basic TV-AR model (2) with bivariate TV-VAR models, $H = 24$

	TV-AR	Ump	TBR	IP	IP gap	M2
U.S., 1959:01–2008:12						
CRPS	1.79	1.77	1.76	1.78	1.67	1.82
r_s		0.94	0.86	0.95	0.94	0.92
Japan, 1959:01–2008:11						
CRPS	2.55	2.24		2.49	2.66	2.44
r_s		0.93		0.95	0.9	0.93
Korea, 1983:02–2009:01						
CRPS	1.95	1.58		1.90	1.87	1.9
r_s		0.88		0.96	0.95	0.95

is a diagonal matrix of volatilities. It is assumed that all of the time-varying parameters follow independent stationary Gaussian AR(1) processes. Denote $\mathbf{\Lambda}_t$ a row vector of all time-varying parameters: $M_{i,t}$, $\mathbf{\Pi}_{is,j,t}$, $h_{i,t}$, $c_{is,t}$. It is assumed that each of the elements of $\mathbf{\Lambda}_t$ is described by

$$\Lambda_{l,t} = (1-\delta_l)\omega_l + \delta_l \Lambda_{l,t-1} + \sigma_l \sqrt{1-\delta_l^2}\, \eta_{l,t}, \tag{4}$$

where the disturbances $\eta_{l,t}$ are independent $\mathcal{N}(0,1)$. The parameters vector for the TV-VAR model is $\boldsymbol{\theta} = (\boldsymbol{\omega}, \boldsymbol{\delta}, \boldsymbol{\sigma})$.

Similarly to the basic model (2) this is a nonlinear state-space model. Estimation and forecasting methods are similar (Appendix A). TV-VAR was estimated using the Laplace's method to approximate the likelihood function (simulated maximum likelihood methods proved to be too time-consuming).[5] As was done before for the basic TV-AR model for degenerating factors $\Lambda_{l,t}$ coefficients δ_l were fixed at the level 0.99 (while some coefficients σ_l were fixed at the level 10^{-5}). Bivariate models for U.S., Japan and Korea were fitted with monthly series. Unemployment ("Ump"), Treasury Bills rate ("TBR"), growth rate of industrial production index ("IP"), output gap based on industrial production ("IP gap") and growth rate of M2 monetary aggregate were used in addition to inflation.[6]

The 24-months-ahead dynamic quasi out-of sample forecasts of monthly inflation from bivariate TV-VAR models (3)–(4) and from the univariate model (2) were compared on the basis of the continuous ranked probability score. Table 2 presents the average values of CRPS.

It should be taken into account that the CRPS measure contains a component which is inherently unpredictable and cannot be captured by any forecasting technique. That is why the levels of

[5] Estimation based on the EIS approximation for likelihood function with 300 simulations and 10 iterations was about 700 times slower.

[6] All of the series except Korean and Japanese unemployment are from International Financial Statistics database (Appendix B, Series B, E, F). Seasonally unadjusted series were adjusted prior to estimation. Kalman filtering procedure in the state-space model $y_t = y_t^* + \varepsilon_t$, $\Delta^2 y_t^* = \zeta_t$ with $\operatorname{var}\varepsilon_t / \operatorname{var}\zeta_t = \lambda = 1.5 \cdot 10^6$ was employed to obtain the output gap series. This is the one-sided counterpart of the Hodrick–Prescott filter.

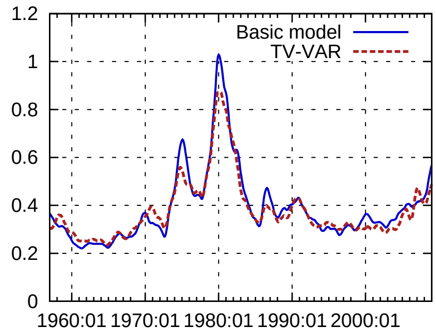

Fig. 4. Measures of uncertainty based on the basic TV-AR model (2) and the TV-VAR with output gap, $H = 24$, U.S., 1957:01–2008:12.

average CRPS do not differ markedly.

One conclusion from Table 2 is that indicators other than inflation itself (unemployment, output gap) can considerably improve inflation forecasts. Table 2 additionally shows Spearman's rank correlations between uncertainty measures based on (2) and uncertainty measures based on TV-VAR models. All uncertainty series were smoothed prior to computation of rank correlations (details are in subsection 4.2). The correlations are quite high and demonstrate that although unemployment and output gap are useful for predicting inflation, their inclusion would not lead to substantial change in the uncertainty measure.

Figure 4 compares uncertainty as measured by the basic model and uncertainty as measured by the TV-VAR with unemployment for the United States. It can be concluded that although unemployment proved to be important for predicting inflation it produces a very similar measure of uncertainty. The discrepancy is small relative to the ambiguity in deciding what is the relevant measure of uncertainty for the Okun–Friedman hypothesis.

A bivariate TV-VAR is much harder to estimate and unemployment or industrial production series are not available, too short or unreliable for most of the countries. So I choose to use the basic univariate TV-AR model (2) for the current study.

3.4. Comparison with ARMA-GARCH

The basic model (2) needs a lot of computer time for estimation, obtaining uncertainty series and diagnostics. Therefore it is desirable to compare it with less computer-intensive models. If it does not provide better forecasts then the additional effort is not worth paying. The TV-AR model (2) was compared with ARMA(2,1)-GARCH(1,1) with skewed t disturbances. The comparison was made for 38 countries from the IFS data-set (Series B) for the period 1957:02–1998:12 and for the $H = 24$ forecast horizon. For 32 countries TV-AR gives lower CRPS. On average CRPS for ARMA-GARCH is 6% higher. The mean value of Spearman's rank correlation between the two uncertainty series is 0.71. (The CRPS values and uncertainty series for the two models were obtained by simulation. Both uncertainty series were smoothed prior to computation of the rank correlations; details on smoothing are in subsection 4.2). The conclusion is that there is a good reason to use TV-AR rather than the less cumbersome ARMA-GARCH.

3.5. PIT-based assessment of forecast calibration

One way to verify the underlying model is to check forecast calibration. Some aspects of calibration can be tested on the basis of so-called probability integral transform (PIT), e.g. Diebold, Gunther, and Tay (1998), Gneiting, Balabdaoui, and Raftery (2007), Wallis (2008). It is a very popular diagnostic tool in the context of time-series modeling. If $y_1, y_2, \ldots$ are observed data and forecast is calibrated so that the c.d.f. of the forecast distribution $F(y_{t+H}|y_1, \ldots, y_t)$ (which is called PIT) adequately represents the data generating process then $v_{t,H} = F(y_{t+H}|y_1, \ldots, y_t)$ is uniformly distributed $U[0, 1]$. For $H = 1$ the series of $v_{t,1}$ should be i.i.d. For $H > 1$ a series of $v_{t,H}$ would in general be dependent. It is useful to convert $v_{t,H}$ to a standard normal form $z_{t,H} = \Phi^{-1}(v_{t,H})$ were $\Phi(\cdot)$ is the standard normal c.d.f. Also useful is "folded" PIT $v'_{t,H} = \left|2v_{t,H} - 1\right|$ and corresponding $z'_{t,H} = \Phi^{-1}(v'_{t,H})$ which should be distributed as $U[0, 1]$ and $\mathcal{N}(0, 1)$ respectively.

In order to verify the inflation forecasts from the model (2) the series of $z_{t,1}$, $z'_{t,1}$, $z_{t,24}$ and $z'_{t,24}$ were calculated for 38 countries for the period 1957:02–1998:12. Inspection of these series reveals several problems with the model. There are signs of seasonal autocorrelation (the twelfth-order autocorrelation of $z_{t,1}$ series is -0.156 on average) and unmodeled volatility clustering (the first-order autocorrelation of $z'_{t,1}$ series is 0.025 on average). The forecast distributions are often biased upwards (the mean of $z_{t,1}$ is 0.029 and the mean of $z_{t,24}$ is 0.135). Also notable is the fact that in most of the countries outliers are more frequent than the model suggests, especially positive outliers corresponding to outbursts of higher inflation (skewness of $z_{t,24}$ is 0.220 on average). More detailed statistics are available from the author.

These diagnostics for forecast calibration should not be considered as ordinary tests for model specification. Even a "wrong" model can provide good forecasts. The purpose is to detect departures from model properties which suggest potential sizable improvements in the accuracy of forecasts. The conclusion from the PIT-based diagnostics is that although the departures are statistically significant, there is no serious indication that the basic model is unusable for an approximate measurement of inflation uncertainty. (Here the well-known disparity between statistical significance and practical importance comes into play). It is not clear what could be a better model which would account for the detected problems. My belief is that the task of developing such a model is excessively difficult. So the decision here is to be satisfied with the current not very ideal model until the arrival of a better one.

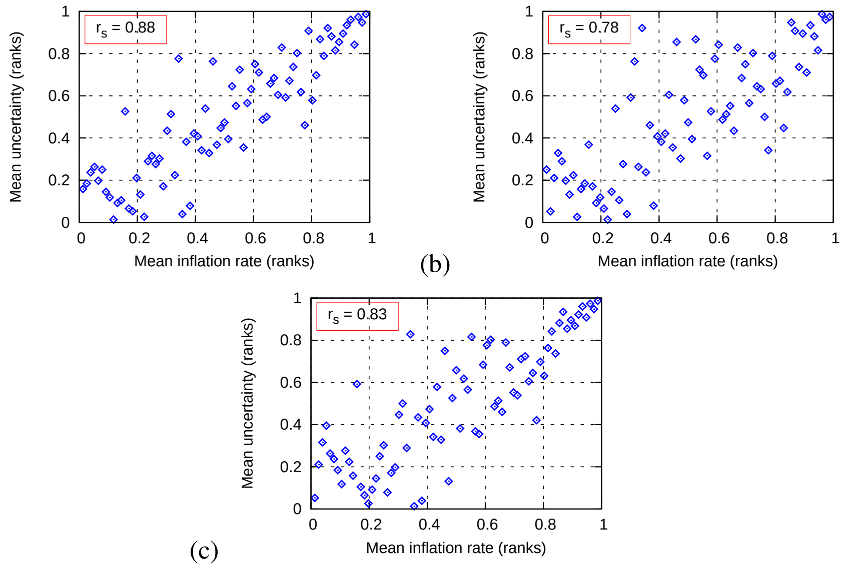

(a) (b) (c)

Fig. 5. The mean inflation and mean inflation uncertainty, 75 countries: (a) TV-AR, $H = 24$, (b) TV-AR, $H = 1$, 1994:01–1998:12 (c) ARMA-GARCH, $H = 24$

4. EMPIRICAL FINDINGS FOR INFLATION–UNCERTAINTY LINKS

4.1. Cross-country links

First I explore the cross-country links between inflation and its uncertainty. The inflation uncertainty indicator was obtained from the basic TV-AR model (2). The cross-country evidence is presented in Figure 5(a), (b), where the average inflation uncertainty is plotted against the average inflation for the period 1994:01–1998:12 for 75 countries. Ranks rather than the original variables are used to make the evidence more robust to fat tails and non-linearity. That inflation rates are characterized by the distributions with fat tails is evident from the estimates of both the TV-AR and ARMA-GARCH models. Robustness to non-linearity is important because the functional form of the relationship is not known.

The cross-country links are rather strong. For the longer-horizon uncertainty ($H = 24$) Spearman's rank correlation is $r_s = 0.88$ and is higher than $r_s = 0.78$ for the short-horizon uncertainty ($H = 1$). Figure 5(c) presents a similar scatterplot for the ARMA-GARCH model discussed above for $H = 24$. Rank correlation is $r_s = 0.83$. For $H = 1$ and ARMA-GARCH the scatterplot (not shown) is very similar to Figure 5(b) and rank correlation is $r_s = 0.78$. These correlation coefficients demonstrate the effect of a longer horizon and the effect of using TV-AR rather than ARMA-GARCH.

With this ranks-based approach information on the quantitative aspect of the inflation–uncertainty link is lost. However, the scatterplots for ranks such as those shown in Figure 5 are tightly related to the copula modeling. (An introductory exposition of copulas can be found in Trivedi and Zimmer (2006)). Assuming some parametric marginal distributions and some specific copula function allows to make quantitative statements about the link. Moreover, it can be observed from the scatterplots that for high-inflation countries the dependence is more pronounced. A quantitative measure of this phenomenon is provided by the tail dependence coefficient known from the copula analysis literature.

Here we have two variables, mean inflation x_i and mean uncertainty y_i. In order to model their joint bivariate distribution using a copula first the marginal distributions of the two variables are modeled by continuous distributions. The corresponding marginal cumulative distribution functions are used to transform the initial variables to uniformly distributed $U[0; 1]$ variables. Second, a joint distribution on the unit square with uniform marginals is fit (the distribution is called copula).

Suppose that the variables x_i and y_i can be transformed to (approximately) normal variates by taking logarithms with shift: $\ln(x - x_0) \sim \mathcal{N}(\mu_x, \sigma_x^2)$ and $\ln(y - y_0) \sim \mathcal{N}(\mu_y, \sigma_y^2)$. Then the two variables after normalization and applying the standard normal c.d.f.

$$u = \Phi\left(\frac{\ln(x - x_0) - \mu_x}{\sigma_x}\right), \qquad v = \Phi\left(\frac{\ln(y - y_0) - \mu_y}{\sigma_y}\right)$$

are approximately uniformly distributed as $U[0; 1]$.

A popular simple copula which allows for upper tail dependence, the rotated Clayton copula, is used. The copula function is

$$C(u, v) = u + v - 1 + (1 - u)(1 - v)D(u, v)^{-1/\delta},$$

where

$$D(u, v) = (1 - u)^{\delta} + (1 - v)^{\delta} - (1 - u)^{\delta}(1 - v)^{\delta},$$

and the corresponding density function is

$$c(u, v) = (1 + \delta)(1 - u)^{\delta}(1 - v)^{\delta}D(u, v)^{-2-1/\delta}.$$

The copula does not exhibit lower tail dependence ($\lambda_L = 0$), while the upper tail dependence coefficient is $\lambda_U = 2^{-1/\delta}$.

The model was fitted to the data on the mean inflation and mean uncertainty ($H = 24$) from the TV-AR model. The estimation was done by the maximum likelihood method in one stage (the copula parameter δ and parameters of the marginal distributions μ_x, σ_x, μ_y, σ_y were estimated jointly). The shift parameters were fixed at $x_0 = 0.01$ and $y_0 = 0.15$ levels which ensure approximate normality after taking logarithms. The ML estimate of the copula parameter is $\hat{\delta} = 2.83$, which corresponds to the upper tail dependence coefficient of $\hat{\lambda}_U = 0.78$. The coefficient is defined as

$$\lambda_U = \lim_{\varepsilon \downarrow 0} \Pr(v > 1 - \varepsilon | u > 1 - \varepsilon) = \lim_{\varepsilon \downarrow 0} \frac{\Pr(u > 1 - \varepsilon; v > 1 - \varepsilon)}{1 - \varepsilon}.$$

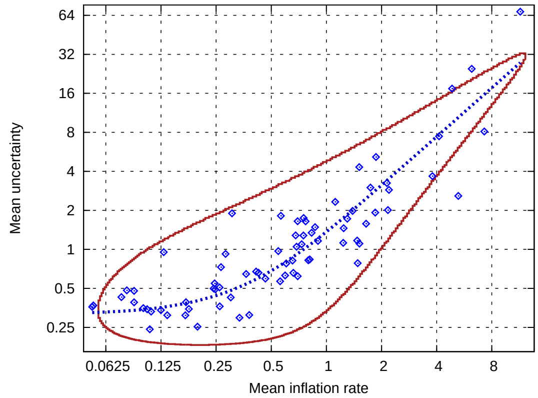

Fig. 6. The 90% region for the mean inflation and mean inflation uncertainty, $H = 24$, 1994:01–1998:12, 75 countries.

It can be interpreted as follows: for a country with mean inflation in the upper 10% tail it shows (approximately) what is the conditional probability that mean uncertainty is also in the 10% tail. The estimate shows that the probability is around 78%, which is rather high. Of course, the sample is too small for the estimate to be precise enough. A reliable inference on tail events requires a much larger sample.

Once the estimates of the copula and the marginals are obtained, some other useful indicators can be inferred.

First, one can obtain a confidence region for the joint distribution of the mean inflation and mean uncertainty. The purpose is to visualize this joint distribution. A confidence region is not something definite. I have chosen the 90% confidence region bordered by the contour of constant density for the joint distribution with standard normal marginals $\Phi^{-1}(u)$, $\Phi^{-1}(v)$ linked with the estimated copula function $\hat{C}(u, v)$. The result in terms of the initial variables

$$x = \exp(\mu_x + \sigma_x \Phi^{-1}(u)) + x_0, \qquad y = \exp(\mu_y + \sigma_y \Phi^{-1}(v)) + y_0$$

is shown in Figure 6.

Second, the estimated copula can be used to produce a regression line. Note that thus far I avoided causal interpretations. The copula analysis does not imply a definite direction of causality. However, if one is willing to assume that inflation causes uncertainty then a regression where inflation uncertainty is the dependent variable can be of interest.

It is natural to use median regression based on estimated copula. It is readily obtained from copula function due to invariance of median under monotone transformations. A formula of the theoretical median regression function for the rotated Clayton copula is

$$v_{med}(u) = 1 - (1 - u)(2^{\delta/(1+\delta)} + (1 - u)^{\delta} - 1)^{-1/\delta}.$$

Pairs $(x, y_{med}(x))$ are obtained from pairs $(u, v_{med}(u))$. The median regression is drawn in Figure 6 as a dotted line. According to the estimated regression function an increase in average monthly inflation rate from 0.5% to 1% is associated with widening of the inflation forecast interval by a factor of 2.[7]

As argued in Ram (1985), Driffill, Mizon, and Ulph (1990) and Davis and Kanago (1996) a cross-country analysis may be not very relevant to the Okun–Friedman hypothesis. Time-series evidence is more important for assessing potential welfare losses due to high inflation as it represents within-country point of view. However, there is a difference between everyday macroeconomic policy and construction of institutions for such policy. For the later cross-country evidence can be more relevant.

4.2. Time-series links

To explore time-series links smoothed series of both inflation and estimated inflation uncertainty were used for 38 countries for the period 1957:02–1998:12. Full series (including data beyond 1998) were used for estimation of the model parameters. The series were smoothed using standard kernel regression on the time trend

$$\hat{y}_t = \frac{\sum_{\tau=1}^{T} K((t-\tau)/h) y_\tau}{\sum_{\tau=1}^{T} K((t-\tau)/h)}$$

with Epanechnikov canonical kernel defined as

$$K(z) = 0.75(1 - z^2/\sigma^2) I(|z| \le \sigma)/\sigma, \qquad \sigma = 15^{1/5}$$

and bandwidth parameter $h = 4$. Such smoothing implements a low-pass filter which filters out high-frequency noise in inflation and uncertainty and allows to estimate the long-run dependence. Of course, smoothing leads to some distortion and a loss of degrees of freedom. Hence, the correlations between the series should not be assessed at the face value.

The uncertainty indicator was computed from 24-months-ahead ($H = 24$) and 1-month-ahead ($H = 1$) forecasts to explore the dependence of the inflation–uncertainty links on the length of the forecast horizon. The uncertainty for $H = 1$ is similar to the conditional variance from the ARMA-GARCH model for monthly inflation series, but it additionally contains the effect of changing persistence. Figure 7 shows the graphs of inflation and uncertainty for 4 countries. The uncertainty series were scaled to be comparable with inflation.

As is evident from the graphs, there is a pronounced similarity in the behavior of inflation and the longer-horizon uncertainty in all four countries. This visual impression of a strong connection is confirmed by the rank correlation coefficients in Table 3. Links between inflation and the

[7] When interpreting this number one should remember that it relates to forecasting of the logarithm of the price level.

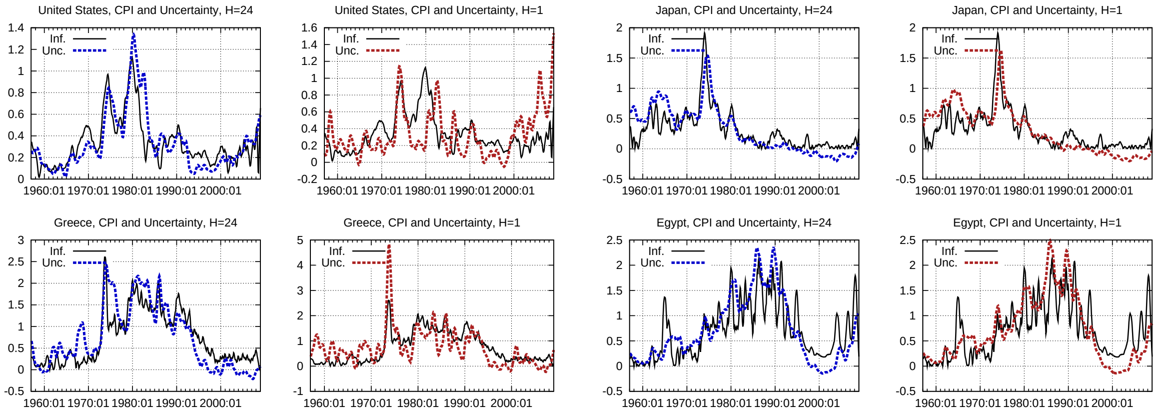

Fig. 7. Inflation and uncertainty for the United States, Japan, Greece and Egypt, 1957:02–1998:12 (longer-horizon, $H = 24$, and short-horizon, $H = 1$).

Table 3. Spearman's rank correlations of inflation with uncertainty and ARMA-GARCH volatility, 1957:02–1998:12

	U.S.	U.S., core	Japan	Greece	Egypt
TV-AR uncertainty, $H = 24$	0.81	0.89	0.73	0.79	0.81
TV-AR uncertainty, $H = 1$	0.32	0.27	0.76	0.56	0.81
ARMA-GARCH volatility	0.37	0.25	0.77	0.54	0.76
ARMA-GARCH uncertainty, $H = 24$	0.39	0.25	0.78	0.56	0.79
ARMA-GARCH uncertainty, $H = 1$	0.37	0.26	0.77	0.56	0.79

GARCH volatility (as well as the short-horizon uncertainty) are weak for the U.S which explains the negative results in Engle (1983).

An additional column for the U.S. in Table 3 was computed from the "core inflation" index which does not include food and energy (Appendix B, Series D). The link between inflation and the longer-horizon uncertainty is stronger for this series.

Figures 8(a) and (b) present the cross-country distributions of Spearman's rank correlation coefficients of inflation and its uncertainty based on the TV-AR model (2). For the longer-horizon uncertainty ($H = 24$) the links are much stronger with median correlation 0.60 for $H = 24$ and 0.34 for $H = 1$. Figure 8(c) presents a similar distribution for the inflation uncertainty indicator from the ARMA-GARCH model for $H = 24$. The median correlation is 0.37. For the short-horizon ARMA-GARCH uncertainty ($H = 1$) and for the ARMA-GARCH conditional variance the distributions are very similar.

Another aspect of the inflation–uncertainty link is the question of the temporal order. The concept of Granger causality is not quite applicable without formulation of a joint model of the links between inflation and its uncertainty. However a simpler technique can still be used. It is based on examination of empirical cross-correlations between ranks of two smoothed series (with the $H = 24$ horizon for uncertainty). Table 4 shows the number of countries for which the given lag corresponds to the largest value of rank correlation. "Lag" corresponds to the lag of inflation (for positive values inflation precedes uncertainty). No country has the maximum in the range $-20, \dots, 0$

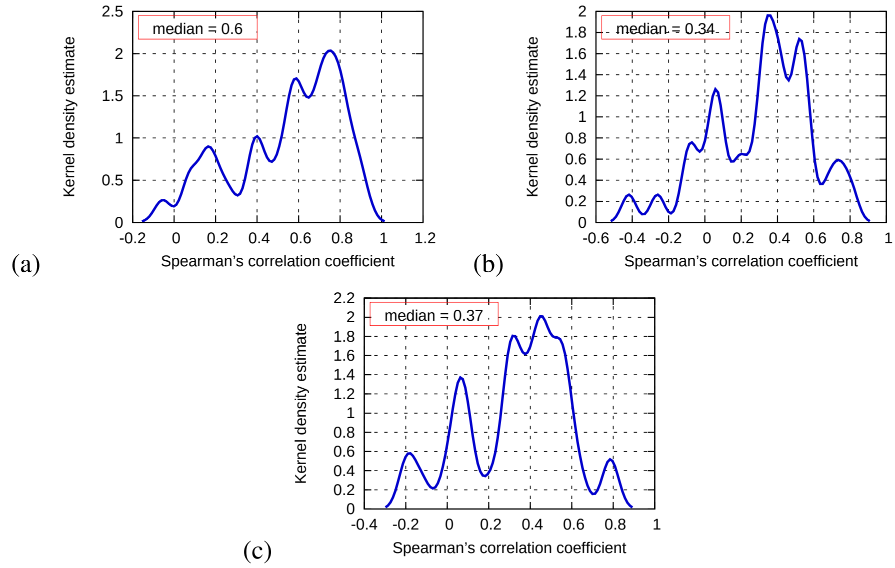

Fig. 8. Time-series rank correlations between inflation and uncertainty, 1957:02–1998:12, 38 countries, kernel density estimates: (a) TV-AR, $H = 24$, (b) TV-AR, $H = 1$, (c) ARMA-GARCH, $H = 24$.

Table 4. Lag for which cross-correlation of inflation and uncertainty is at maximum, 1957:02–1998:12, 38 countries

Lag	1	2	3	4	5	6	7	8	9	10	11–15	16–20
# of countries	1	5	5	4	3	3	2	2	1	1	2	2

(not shown in the table). Seven countries have the maximum outside the range $-20, \ldots, 20$. The median corresponds to the lag of 4 months. In addition, for 6 of the 7 countries with maximum outside the range $-20, \ldots, 20$ there exists also a local maximum in the range $1, \ldots, 20$.

The comparison of cross-correlations provides evidence that inflation precedes uncertainty. (Of course, one should be extremely cautious with causal interpretation of such evidence). Note too that in all of the countries zero lag does not correspond to the highest cross-correlation. Hence, the links between inflation and inflation uncertainty are in fact stronger than can be deduced from Figure 8(a).

5. DISCUSSION AND CONCLUSIONS

This study makes advancement in the research on the Okun–Friedman hypothesis. The results obtained have implications not only for the hypothesis itself, but also for a related research on inflation modeling and forecasting. The significance of time variation in the processes which drive inflation and of the length of the forecast horizon was explored. Both aspects are shown to be essential for the Okun–Friedman hypothesis.

One important finding of this study is that a TV-AR model with three latent factors can be a useful instrument of modeling and forecasting inflation. The measure of uncertainty derived from the TV-AR model proved to be comparable to a measure based on forecasters' disagreement in the Livingston Survey. Also it was shown for the U.S., Japan and Korea that adding other macroeconomic indicators can improve forecast quality compared to the univariate TV-AR model, but does not considerably change the measure on inflation uncertainty. At the same time a simpler ARMA-GARCH model of inflation cannot fully capture the behavior of uncertainty in many countries.

The method of obtaining time-varying estimates of inflation uncertainty developed in this paper can be used in other relevant contexts. The possibilities in the area of analyzing the effects of inflation uncertainty are numerous. One example is examining the links between inflation uncertainty and labor contract durations (see Rich and Tracy (2004)).

The study suggests convincing evidence that there exists a strong positive relation between inflation and its uncertainty both *across countries* and in the time series dimension *within countries*. There is also time-series evidence that inflation precedes uncertainty.

A longer-horizon inflation uncertainty (at least one year ahead, two years in this study) is more relevant to the original concerns behind the Okun–Friedman hypothesis. Notable is the fact that empirically it is also more closely related to the level of inflation than a short-horizon inflation uncertainty (one month in this study).

The results obtained conform to the theoretical model by Ball (1992). A more technical argument by Holland (1995) stressing the role of uncertainty about functioning of inflation process is also compatible with the evidence as long as the TV-AR model provides a good enough description of inflation. For the models by Cukierman and Meltzer (1986), Devereux (1989) and Geraats (2006) it is not clear how to explain the observed patterns of cross-correlations between inflation and inflation uncertainty. The models by Devereux (1989) and Mondino, Sturzenegger, and Tommasi (1996) use very specific mechanisms (based on wage indexation and political cycles respectively) and are not universally applicable. They cannot explain the links for a substantial portion of the countries in which wage indexation and monetary reforms are not common. Models by Devereux (1989) and Geraats (2006) are based on shocks which can be of short-run nature, so the observed strong links for the long-horizon uncertainty are not fully explained. Moreover, for the U.S. the evidence is that when some of the shocks are removed by using a core inflation series the link is even stronger.

Of course, correlation is not always an evidence of causation and precedence is not an evidence of the direction of causation. Thus, too confident policy-related judgments cannot be made at the current stage of the research. However, one can argue that this study is an important step in the right direction.

To conclude: it was demonstrated by applying modern non-linear state-space modeling to extensive empirical data that the old argument by Okun (1971) is valid. High inflation and low inflation uncertainty are largely incompatible. Using Okun's words, "the adoption of a public policy designed to yield steady, fully anticipated inflation would commit the government to an impossible goal" (p. 488). Some policy implications of the result do not depend on the direction of causation. These can be stated in terms of the construction of institutions for the economic policy. We do not have a choice between institutions which encourage low inflation and low inflation uncertainty and institutions which encourage high inflation and low inflation uncertainty. High inflation always comes together with high inflation uncertainty. Thus, institutions which encourage high inflation would lead to inflation uncertainty.

REFERENCES

Ang, A., G. Bekaert, and M. Wei (2005): "Do Macro Variables, Asset Markets or Surveys Forecast Inflation Better?," Working Paper 11538, National Bureau of Economic Research.

Atkeson, A., and L. E. Ohanian (2001): "Are Phillips Curves Useful For Forecasting Inflation?," *Federal Reserve Bank of Minneapolis Quarterly Review*, 25(1), 2–11.

Baillie, R. T., C. Chung, and M. A. Tieslau (1996): "Analysing Inflation by the Fractionally Integrated ARFIMA–GARCH Model," *Journal of Applied Econometrics*, 11(1), 23–40.

Ball, L. (1992): "Why Does High Inflation Raise Inflation Uncertainty?," *Journal of Monetary Economics*, 29(3), 371–388.

Ball, L., and S. G. Cecchetti (1990): "Inflation and Uncertainty at Short and Long Horizons," *Brookings Papers on Economic Activity*, (1), 215–254.

Barro, R. J., and D. B. Gordon (1983a): "A Positive Theory of Monetary Policy in a Natural Rate Model," *Journal of Political Economy*, 91(4), 589–610.

——— (1983b): "Rules, Discretion and Reputation in a Model of Monetary Policy," *Journal of Monetary Economics*, 12(1), 101–121.

Bomberger, W. A. (1996): "Disagreement as a Measure of Uncertainty," *Journal of Money, Credit and Banking*, 28(3), 381–392.

Brunner, A. D., and G. D. Hess (1993): "Are Higher Levels of Inflation Less Predictable? A State-Dependent Conditional Heteroscedasticity Approach," *Journal of Business & Economic Statistics*, 11(2), 187–197.

Carlson, J. A. (1977): "A Study of Price Forecasts," in *Annals of Economic and Social Measurement*, vol. 6 of *NBER Chapters*, pp. 33–63. National Bureau of Economic Research.

Cogley, T., and T. J. Sargent (2001): "Evolving Post-World War II U.S. Inflation Dynamics," *NBER Macroeconomics Annual*, 16, 331–388.

——— (2005): "Drifts and Volatilities: Monetary Policies and Outcomes in the Post WWII US," *Review of Economic Dynamics*, 8(2), 262–302.

Cukierman, A., and A. H. Meltzer (1986): "A Theory of Ambiguity, Credibility, and Inflation under Discretion and Asymmetric Information," *Econometrica*, 54(5), 1099–1128.

DAAL, E., A. NAKA, AND B. SANCHEZ (2005): "Re-Examining Inflation and Inflation Uncertainty in Developed and Emerging Countries," *Economics Letters*, 89(2), 180–186.

DAVIS, G. K., AND B. KANAGO (1996): "The Missing Link: Intra-Country Evidence on the Relationship between High and Uncertain Inflation from High-Inflation Countries," *Southern Economic Journal*, 63(1), 205–222.

——— (1998): "High and Uncertain Inflation: Results from a New Data Set," *Journal of Money, Credit and Banking*, 30(2), 218–230.

DAVIS, G. K., AND B. KANAGO (2000): "The Level and Uncertainty of Inflation: Results from OECD Forecasts," *Economic Inquiry*, 38(1), 58–72.

DAVIS, R. A., AND G. RODRIGUEZ-YAM (2005): "Estimation for State-Space Models Based on a Likelihood Approximation," *Statistica Sinica*, 15, 381–406.

DEVEREUX, M. (1989): "A Positive Theory of Inflation and Inflation Variance," *Economic Inquiry*, 27(1), 105–116.

DIEBOLD, F. X., T. A. GUNTHER, AND A. S. TAY (1998): "Evaluating Density Forecasts with Applications to Financial Risk Management," *International Economic Review*, 39(4), 863–883.

DOYLE, M. (2006): "Empirical Phillips Curves in OECD Countries: Has There Been a Common Breakdown?," Staff General Research Paper 12684, Iowa State University, Department of Economics.

DRIFFILL, J., G. E. MIZON, AND A. ULPH (1990): "Costs of inflation," in *Handbook of Monetary Economics*, ed. by B. M. Friedman, and F. H. Hahn, vol. 2, chap. 19, pp. 1013–1066. Elsevier.

ENGLE, R. F. (1983): "Estimates of the Variance of U. S. Inflation Based upon the ARCH Model," *Journal of Money, Credit and Banking*, 15(3), 286–301.

EVANS, M. (1991): "Discovering the Link Between Inflation Rates and Inflation Uncertainty," *Journal of Money, Credit and Banking*, 23(2), 169–184.

EVANS, M., AND P. WACHTEL (1993): "Inflation Regimes and the Sources of Inflation Uncertainty," *Journal of Money, Credit and Banking*, 25(3, Part 2), 475–511.

FANG, W., S. M. MILLER, AND C.-C. YEH (2007): "Does a Threshold Inflation Rate Exist? Quantile Inferences for Inflation and Its Variability," Working papers 2007-45, University of Connecticut, Department of Economics, (revised: June 2009).

FOSTER, E. (1978): "The Variability of Inflation," *The Review of Economics and Statistics*, 60(3), 346–350.

FOUNTAS, S., M. KARANASOS, AND J. KIM (2006): "Inflation Uncertainty, Output Growth Uncertainty and Macroeconomic Performance," *Oxford Bulletin of Economics and Statistics*, 68(3), 319–343.

FRIEDMAN, M. (1977): "Nobel Lecture: Inflation and Unemployment," *The Journal of Political Economy*, 85(3), 451–472.

GALE, W. A. (1981): "Temporal Variability of United States Consumer Price Index," *Journal of Money, Credit and Banking*, 13(3), 273–297.

GARRATT, A., K. LEE, M. H. PESARAN, AND Y. SHIN (2003): "Forecast Uncertainties in Macroeconomic Modeling: An Application to the U.K. Economy," *Journal of the American Statistical Association*, 98(464), 829–838.

GERAATS, P. M. (2006): "Inflation and Its Variation: An Alternative Explanation," Discussion paper, University of Cambridge.

GNEITING, T., F. BALABDAOUI, AND A. E. RAFTERY (2007): "Probabilistic Forecasts, Calibration and Sharpness," *Journal* *of the Royal Statistical Society: Series B*, 69, 243–268.

Gneiting, T., and A. E. Raftery (2007): "Strictly Proper Scoring Rules, Prediction, and Estimation," *Journal of the American Statistical Association*, 102, 359–378.

Golob, J. E. (1993): "Inflation, Inflation Uncertainty, and Relative Price Variability: A Survey," Research Working Paper 93-15, Federal Reserve Bank of Kansas City.

Hansen, B. E. (1994): "Autoregressive Conditional Density Estimation," *International Economic Review*, 35(3), 705–730.

Henry, O. T., N. Olekalns, and S. Suardi (2007): "Testing for Rate Dependence and Symmetry in Inflation Uncertainty: Evidence from the G7 Economies," *Economics Letters*, 94(3), 383–388.

Hess, G. D., and C. S. Morris (1996): "The Long-Run Costs of Moderate Inflation," *Federal Reserve Bank of Kansas City Economic Review*, 81(Q II), 71–88.

Holland, A. S. (1993): "Uncertain Effects of Money and the Link between the Inflation Rate and Inflation Uncertainty," *Economic Inquiry*, 31(1), 39–51.

Holland, A. S. (1995): "Inflation and Uncertainty: Tests for Temporal Ordering," *Journal of Money, Credit and Banking*, 27(3), 827–837.

Jungbacker, B., and S. J. Koopman (2007): "Monte Carlo Estimation for Nonlinear Non-Gaussian State Space Models," *Biometrika*, 94(4), 827–839.

Kim, C.-J. (1993): "Unobserved-Component Time Series Models with Markov-Switching Heteroscedasticity: Changes in Regime and the Link between Inflation Rates and Inflation Uncertainty," *Journal of Business & Economic Statistics*, 11(3), 341–349.

Kontonikas, A. (2004): "Inflation and Inflation Uncertainty in the United Kingdom, Evidence from GARCH Modelling," *Economic Modelling*, 21(3), 525–543.

Lendvai, J. (2006): "Inflation Dynamics and Regime Shifts," ECB Working Paper 684, European Central Bank.

Logue, D. E., and T. D. Willett (1976): "A Note on the Relation between the Rate and Variability of Inflation," *Economica*, 43(170), 151–158.

Mondino, G., F. Sturzenegger, and M. Tommasi (1996): "Recurrent High Inflation and Stabilization: A Dynamic Game," *International Economic Review*, 37(4), 981–996.

Okun, A. M. (1971): "The Mirage of Steady Inflation," *Brookings Papers on Economic Activity*, (2), 485–498.

Pagan, A. R., A. D. Hall, and P. K. Trivedi (1983): "Assessing the Variability of Inflation," *The Review of Economic Studies*, 50(4), 585–596.

Ram, R. (1985): "Level and Variability of Inflation: Time-Series and Cross-Section Evidence from 117 Countries," *Economica*, 52(206), 209–223.

Rich, R., and J. Tracy (2004): "Uncertainty and Labor Contract Durations," *The Review of Economics and Statistics*, 86(1), 270–287.

Richard, J.-F., and W. Zhang (2007): "Efficient High-Dimensional Importance Sampling," *Journal of Econometrics*, 141(2), 1385–1411.

Shimada, J., and Y. Tsukuda (2005): "Estimation of Stochastic Volatility Models: An Approximation to the Nonlinear State Space Representation," *Communications in Statistics — Simulation and Computation*, 34(2), 429–450.

Stock, J. H., and M. W. Watson (2007): "Why Has U.S. Inflation Become Harder to Forecast?," *Journal of Money, Credit and Banking*, 39(s1), 3–33.

——— (2008): "Phillips Curve Inflation Forecasts," Working Paper 14322, National Bureau of Economic Research.

Thornton, J. (2007): "The Relationship between Inflation and Inflation Uncertainty in Emerging Market Economies," *Southern Economic Journal*, 73(4), 858–870.

Trivedi, P. K., and D. M. Zimmer (2006): "Copula Modeling: An Introduction for Practitioners," *Foundations and Trends® in Econometrics*, 1(1), 1–111.

Ungar, M., and B. Zilberfarb (1993): "Inflation and Its Unpredictability —Theory and Empirical Evidence," *Journal of Money, Credit and Banking*, 25(4), 709–720.

Wachtel, P. (1977): "Survey Measures of Expected Inflation and Their Potential Usefulness," in *Analysis of Inflation: 1965–1974*, ed. by J. Popkin, chap. 9, pp. 361–402. National Bureau of Economic Research.

Wallis, K. F. (2008): "Forecast Uncertainty, Its Representation and Evaluation," in *Econometric Forecasting and High-Frequency Data Analysis*, ed. by R. S. Mariano, and Y.-K. Tse, vol. 13 of *Lecture Notes Series, Institute for Mathematical Sciences, National University of Singapore*. World Scientific Publishing Company.

Zarnowitz, V., and L. A. Lambros (1987): "Consensus and Uncertainty in Economic Prediction," *The Journal of Political Economy*, 95(3), 591–621.

A. ESTIMATION AND FORECASTING FOR NON-LINEAR STATE-SPACE MODELS

A.1. Estimation using the Laplace's approximation and simulated maximum likelihood methods

Let $\boldsymbol{Y} = (\boldsymbol{Y}_1, \ldots, \boldsymbol{Y}_T)$ be an observed series. A typical observation $\boldsymbol{Y}_t$ is a $1 \times k$ state vector at time t. The model for the series is stated in terms of the state variables $\boldsymbol{\Lambda} = (\boldsymbol{\Lambda}_1, \ldots, \boldsymbol{\Lambda}_T)$. A typical observation $\boldsymbol{\Lambda}_t$ is a $1 \times m$ vector of state variables. The joint (complete data) density of $\boldsymbol{Y}$ and $\boldsymbol{\Lambda}$ is known up to some vector of parameters $\boldsymbol{\theta}$: $f(\boldsymbol{Y}, \boldsymbol{\Lambda}) = f(\boldsymbol{Y}, \boldsymbol{\Lambda}|\boldsymbol{\theta})$. The maximum likelihood estimates of $\boldsymbol{\theta}$ are given by

$$\hat{\boldsymbol{\theta}} = \arg\max_{\boldsymbol{\theta}} f(\boldsymbol{Y}|\boldsymbol{\theta}),$$

where $f(\boldsymbol{Y}|\boldsymbol{\theta})$ is the marginal distribution of the observed data, which is called the likelihood function in this context. The marginal distribution of $\boldsymbol{Y}$ is given by a multidimensional integral $f(\boldsymbol{Y}|\boldsymbol{\theta}) = \int f(\boldsymbol{Y}, \boldsymbol{\Lambda}|\boldsymbol{\theta}) d\boldsymbol{\Lambda}$ and is not known in a closed form. Below the dependence on the vector of parameters $\boldsymbol{\theta}$ is ignored for notation simplicity.

In popular methods "posterior density" $f(\boldsymbol{\Lambda}|\boldsymbol{Y})$ is approximated with some appropriate density $g(\boldsymbol{\Lambda})$ to get an estimate of this integral. The approximating density $g(\boldsymbol{\Lambda})$ should be close enough to $f(\boldsymbol{\Lambda}|\boldsymbol{Y})$ and should be chosen adaptively given the vector of parameters $\boldsymbol{\theta}$. (So actually this is a family of distributions $g(\boldsymbol{\Lambda}|\boldsymbol{Y};\boldsymbol{\theta})$ indexed by $\boldsymbol{\theta}$ and depending on $\boldsymbol{Y}$.)

One crude estimate to the log-likelihood is given by applying the *Laplace's approximation* to the integral (e.g. Shimada and Tsukuda (2005); in Davis and Rodriguez-Yam (2005) it is called the approximate likelihood method). It is based on the quadratic expansion of $\ln f(\boldsymbol{Y},\boldsymbol{\Lambda})$ around $\hat{\boldsymbol{\Lambda}}(\boldsymbol{Y})$ defined as the argument of the maximum of $f(\boldsymbol{Y},\boldsymbol{\Lambda})$ with respect to $\boldsymbol{\Lambda}$. (The point $\hat{\boldsymbol{\Lambda}} = \hat{\boldsymbol{\Lambda}}(\boldsymbol{Y})$ is also the mode of the posterior density $f(\boldsymbol{\Lambda}|\boldsymbol{Y})$). Suppose $\ln f_a(\boldsymbol{Y},\boldsymbol{\Lambda})$ is such a quadratic approximation. Then $f_a(\boldsymbol{Y},\boldsymbol{\Lambda})$ can be integrated analytically with respect to $\boldsymbol{\Lambda}$ to give an approximate likelihood function

$$f_{LA}(\boldsymbol{Y}) = f_a(\boldsymbol{Y},\boldsymbol{\Lambda})/g(\boldsymbol{\Lambda}).$$

Here $g(\boldsymbol{\Lambda})$ corresponds to the multivariate normal distribution and is obtained by matching the first and second derivatives of $\ln g(\boldsymbol{\Lambda})$ at $\hat{\boldsymbol{\Lambda}}$ to those of $\ln f(\boldsymbol{Y},\boldsymbol{\Lambda})$. The approximation $f_{LA}(\boldsymbol{Y})$ is maximized with respect to $\boldsymbol{\theta}$ to obtain an estimate of $\boldsymbol{\theta}$.

A better approximation (e.g. Jungbacker and Koopman (2007)) can be obtained (at a price of more extensive computations) by Monte Carlo integration. The integral for $f(\boldsymbol{Y})$ is represented as

$$f(\boldsymbol{Y}) = \int \frac{f(\boldsymbol{Y},\boldsymbol{\Lambda})}{g(\boldsymbol{\Lambda})} g(\boldsymbol{\Lambda}) d\boldsymbol{\Lambda}$$

which can be approximated by a sample mean given a Monte Carlo sample $\boldsymbol{\Lambda}^s$, $s = 1,\ldots,S$ from the proposal distribution $g(\boldsymbol{\Lambda})$:

$$f(\boldsymbol{Y}) \approx f_{\mathrm{MC}}(\boldsymbol{Y}) = \frac{1}{S}\sum_{s=1}^{S} \frac{f(\boldsymbol{Y},\boldsymbol{\Lambda}^s)}{g(\boldsymbol{\Lambda}^s)}.$$

A proposal distribution can, for example, be obtained from the same quadratic approximation at the mode as above. The so called *simulated maximum likelihood* method (also known as the Monte Carlo maximum likelihood method) maximizes $f_{\mathrm{MC}}(\boldsymbol{Y}) = f_{\mathrm{MC}}(\boldsymbol{Y}|\boldsymbol{\theta})$ with respect to $\boldsymbol{\theta}$ to obtain an estimate of $\boldsymbol{\theta}$. The estimator is asymptotically equivalent to the exact maximum likelihood one provided S tends to infinity at suitable pace as T is increased.

When $f(\boldsymbol{Y},\boldsymbol{\Lambda})$ can be represented as a product of T components with t-th component depending only on $\boldsymbol{\Lambda}_t$ and $\boldsymbol{\Lambda}_{t-1}$ it is convenient to specify the multi-period normal distribution corresponding to $g(\cdot)$ recursively as a chain of conditional single-period normal distributions:

$$\boldsymbol{\Lambda}_t|\boldsymbol{\Lambda}_{t-1} \sim N(\boldsymbol{K}_t + \boldsymbol{\Lambda}_{t-1}\boldsymbol{L}_t, \boldsymbol{M}_t). \tag{5}$$

This factorization allows to take advantage of the block tridiagonal structure of the Hessian of $\ln f(\boldsymbol{Y},\boldsymbol{\Lambda})$ with respect to $\boldsymbol{\Lambda}$ and thus to devise efficient algorithms for finding the mode $\hat{\boldsymbol{\Lambda}}$ and for Monte Carlo simulation. The details are available from the author upon request.

A.2. Efficient importance sampling

Monte Carlo integration can give unreliable results if the proposal distribution is poor. A good proposal distribution can be very important for obtaining good Monte Carlo approximations to filtering and forecasting distributions (see below). Richard and Zhang (2007) propose a piecemeal approach to optimal fitting of a proposal distribution to be used in Monte Carlo integration in high-dimensional models. The method is called the *efficient importance sampling* (EIS). The goal is to obtain an "efficient" Gaussian proposal distribution $g(\boldsymbol{\Lambda})$ represented as (5).

Define a "stopgap" function $\tilde{\mu}_t$ which for period t captures dependence on $\boldsymbol{\Lambda}_{t-1}$ such that logarithm of stopgap, $\ln \tilde{\mu}_t$, is a quadratic function of $\boldsymbol{\Lambda}_{t-1}$. Then the basic EIS regression can be rewritten for $t = 2, \ldots, T$ as

$$\ln \phi_t + \ln \tilde{\mu}_{t+1} = D_t + \boldsymbol{\Lambda}_t \boldsymbol{D}_t^0 + \boldsymbol{\Lambda}_{t-1} \boldsymbol{D}_t^1 + \boldsymbol{\Lambda}_t \boldsymbol{D}_t^{00} \boldsymbol{\Lambda}_t^\top + \boldsymbol{\Lambda}_t \boldsymbol{D}_t^{01} \boldsymbol{\Lambda}_{t-1}^\top + \boldsymbol{\Lambda}_{t-1} \boldsymbol{D}_t^{11} \boldsymbol{\Lambda}_{t-1}^\top + R_t. \quad (6)$$

For $t = 1$ the terms with $\boldsymbol{\Lambda}_{t-1}$ are missing. Matrices $\boldsymbol{D}_t^{00}$ and $\boldsymbol{D}_t^{11}$ are assumed to be symmetric here. The EIS regressions are estimated from S observations $\boldsymbol{\Lambda}^s$, $s = 1, \ldots, S$ generated from a proposal distribution $g(\boldsymbol{\Lambda})$ based on a current approximation for $\boldsymbol{K}_t$, $\boldsymbol{L}_t$, $\boldsymbol{M}_t$. New parameters $\boldsymbol{K}_t$, $\boldsymbol{L}_t$ and $\boldsymbol{M}_t$ can be recovered from coefficients of the EIS regressions:

$$\boldsymbol{M}_t = -\frac{1}{2}(\boldsymbol{D}_t^{00})^{-1}, \qquad \boldsymbol{K}_t = (\boldsymbol{D}_t^0)^\top \boldsymbol{M}_t, \qquad \boldsymbol{L}_t = (\boldsymbol{D}_t^{01})^\top \boldsymbol{M}_t.$$

The stopgap function is obtained after estimation of period t regression as

$$\ln \tilde{\mu}_t = \ln \phi_t + \ln \tilde{\mu}_{t+1} - \ln g(\boldsymbol{\Lambda}_t, \boldsymbol{\Lambda}_{<t}) - R_t,$$

where R_t are the residuals from the regression. Several iterations of the method are made. New $\boldsymbol{K}_t$, $\boldsymbol{L}_t$ and $\boldsymbol{M}_t$ give proposal distribution, from which new $\boldsymbol{\Lambda}_t^s$ are taken. New $\boldsymbol{\Lambda}_t^s$ are used as data in the EIS regressions leading to new $\boldsymbol{K}_t$, $\boldsymbol{L}_t$ and $\boldsymbol{M}_t$ and so on. Finally, Monte Carlo likelihood $f_{\mathrm{MC}}(\boldsymbol{Y}) = f_{\mathrm{MC}}(\boldsymbol{Y}|\boldsymbol{\theta})$ for given $\boldsymbol{\theta}$ is obtained as described above. The problem of finding a good approximation to $f(\boldsymbol{\Lambda}|\boldsymbol{Y}, \boldsymbol{\theta})$ should be solved anew for each value of parameters vector $\boldsymbol{\theta}$.

A.3. Filtering and on-line forecasting

Filtering refers to exploring characteristics of a series of conditional distribution $\boldsymbol{\Lambda}_{\leq t}|\boldsymbol{Y}_{\leq t}$, where $t = 1, 2, \ldots$, given some vector of parameters $\boldsymbol{\theta}$. (The following shortcut notation is used here: $\boldsymbol{X}_{\leq t} = (\boldsymbol{X}_1, \ldots, \boldsymbol{X}_t)$). Filtering imitates inference in the situation of a sequential flow of information. If we know observable variable up to time t, $\boldsymbol{Y}_{\leq t}$, we can explore $\boldsymbol{\Lambda}_{\leq t}|\boldsymbol{Y}_{\leq t}$. With the arrival of the next observation $\boldsymbol{Y}_{t+1}$ we can explore $\boldsymbol{\Lambda}_{\leq t+1}|\boldsymbol{Y}_{\leq t+1}$, and so on. The results of filtering can be used to implement quasi out-of-sample forecasting, that is, to (partially) imitate on-line forecasting.

Filtering distributions $\mathbf{\Lambda}_{\le t}|\mathbf{Y}_{\le t}$ are not known, but their densities $f(\mathbf{\Lambda}_{\le t}|\boldsymbol{y}_{\le t})$ are proportional to the known densities $f(\boldsymbol{y}_{\le t}, \mathbf{\Lambda}_{\le t})$. Monte Carlo simulations can be used to approximate them with discrete distributions based on a set of generated trajectories. A single proposal distribution $g(\mathbf{\Lambda}|\mathbf{Y}) = g(\mathbf{\Lambda}_{\le T}|\mathbf{Y}_{\le T})$ is used in the paper. It approximates posterior (smoothing) distribution $f(\mathbf{\Lambda}|\mathbf{Y})$ for the full time series available at time T. The quality of proposal distribution is crucial for good Monte Carlo so EIS is used to fit $g(\mathbf{\Lambda}|\mathbf{Y})$. A set of trajectories $\mathbf{\Lambda}^s$, $s = 1, \ldots, S$ is generated according to $g(\mathbf{\Lambda}|\mathbf{Y})$.

For time t there is a set of S trajectories $\mathbf{\Lambda}^s_{\le t}$ obtained by truncating trajectories $\mathbf{\Lambda}^s$. Approximate filtering distribution is a discrete distribution assigning to each trajectory a probability w^s_t. This probability is the normalized importance weight

$$w^s_t = \frac{v^s_t}{\sum_{k=1}^S v^k_t}$$

which is calculated from unnormalized weights

$$v^s_t = f(\boldsymbol{y}_{\le t}, \mathbf{\Lambda}^s_{\le t})/g(\mathbf{\Lambda}^s_{\le t}|\mathbf{Y}).$$

The weights can be computed recursively following obvious factorization of the densities $f(\boldsymbol{y}_{\le t}, \mathbf{\Lambda}_{\le t})$ and $g(\mathbf{\Lambda}_{\le t}|\mathbf{Y})$.

On-line forecasting for the model can also be implemented by means of Monte Carlo simulation. Assume that distributions $\mathbf{Y}_t|\mathbf{Y}_{t-1}, \mathbf{\Lambda}_t$ and $\mathbf{\Lambda}_t|\mathbf{\Lambda}_{t-1}$ are determined by the model and there is an algorithm for generating random variables from these distributions. Then given $\mathbf{Y}_t$ and $\mathbf{\Lambda}^s_t$ one can generate future values $\mathbf{Y}^s_{t+1}$, $\mathbf{\Lambda}^s_{t+1}$, $\mathbf{Y}^s_{t+2}$, $\mathbf{\Lambda}^s_{t+2}, \ldots$ recursively, where s is an index of a Monte Carlo experiment. These trajectories have the associated importance weights w^s_t based on Monte Carlo filtering step. For a sample of future trajectories $\mathbf{Y}^s_{\ge t}$, $s = 1, \ldots, S$ with associated importance weights one can estimate various forecast statistics like point forecasts, interval forecasts and so on. For example, to get an interval forecast for $q = \sum_{h=1}^H Y_{i,t+h}$, one simulates a sample of q^s and calculates the relevant sample quantiles.

To reduce the size of Monte Carlo samples where some trajectories can have negligible weights one can obtain a smaller set of trajectories with equal weights by using resampling. (This can provide a dramatic increase in the speed of computations when one needs Monte Carlo sample to be sorted). For resampling a simple method called systematic sampling is used. To obtain a sample of size M one selects initial trajectories with indices s_j such that

$$\sum_{k=1}^{s_j-1} w^k_t < \frac{j-1+u}{M} \le \sum_{k=1}^{s_j} w^k_t$$

for $j = 1, \ldots, M$ where $u \sim U[0, 1]$ is a random offset.

B. DATA SOURCES

Series A U.S. Bureau of Labor Statistics, CPIAUCNS — consumer price index for all urban consumers: all items, monthly, not seasonally adjusted. Retrieved on 2009-11-09 from `http://research.stlouisfed.org/fred2/categories/9`.

Series B International Financial Statistics database, International Monetary Fund. Retrieved on 2009-04-21 from `http://www.imfstatistics.org`.

Series C Livingston survey. Retrieved on 2009-11-09 from `http://www.philadelphiafed.org/research-and-data/real-time-center/livingston-survey/`.

Series D U.S. Bureau of Labor Statistics, CPILFESL — consumer price index for all urban consumers: all items less food & energy, monthly, seasonally adjusted. Retrieved on 2009-11-09 from `http://research.stlouisfed.org/fred2/categories/9`.

Series E Japan Statistics Bureau, Labour Force Survey, unemployment rate, seasonally adjusted, both sexes, total. Retrieved on 2010-05-03 from `http://www.stat.go.jp/english/data/roudou/lngindex.htm`.

Series F International Labour Organization, LABORSTA — Labour Statistics Database, Table: B4 Unemployment, Korea, general level, total. Retrieved on 2010-05-13 form `http://laborsta.ilo.org/`.