

# Forecasting under Long Memory and Nonstationarity\*

Uwe Hassler and Marc-Oliver Pohle<sup>†</sup>

January 23, 2019

## Abstract

Long memory in the sense of slowly decaying autocorrelations is a stylized fact in many time series from economics and finance. The fractionally integrated process is the workhorse model for the analysis of these time series. Nevertheless, there is mixed evidence in the literature concerning its usefulness for forecasting and how forecasting based on it should be implemented.

Employing pseudo-out-of-sample forecasting on inflation and realized volatility time series and simulations we show that methods based on fractional integration clearly are superior to alternative methods not accounting for long memory, including autoregressions and exponential smoothing. Our proposal of choosing a fixed fractional integration parameter of  $d = 0.5$  a priori yields the best results overall, capturing long memory behavior, but overcoming the deficiencies of methods using an estimated parameter.

Regarding the implementation of forecasting methods based on fractional integration, we use simulations to compare local and global

---

\***Acknowledgements:** We thank Federico M. Bandi, A. Ronald Galant, Patrik Guggenberger, Joon Y. Park and Antonio Rubia for many helpful comments. Earlier versions of this paper were presented at Indiana University, Bloomington, at Penn State University, the 2nd Conference on New Trends and Developments in Econometrics (Bank of Portugal), the Statistische Woche Linz 2018 and the Hannover Workshop on Long Memory in October 2018.

<sup>†</sup>E-mail: hassler@wiwi.uni-frankfurt.de and pohle@econ.uni-frankfurt.de, Address: Campus Westend, RuW Building, 60323 Frankfurt, Germany

semiparametric and parametric estimators of the long memory parameter from the Whittle family and provide asymptotic theory backed up by simulations to compare different mean estimators. Both of these analyses lead to new results, which are also of interest outside the realm of forecasting.

**Keywords:** prediction; fractional integration; inflation; realized volatility

**JEL classification:** C22 (time series models), C53 (forecasting), C58 (financial econometrics)

## 1 Introduction

Long memory or strong persistence are considered as stylized facts in many time series from a variety of fields, from economics and finance over political sciences to hydrology. The strong persistence characterized by slowly decaying autocorrelations on the one hand makes modeling and subsequent estimation hard for these series. Fractional integration [FI] of order  $d$  is the most widely used model to capture strong persistence and long memory. Estimated values of  $d$  vary between 0 and 1, with  $d = 1/2$  separating the reign of stationarity from nonstationarity. FI thus offers a lot of flexibility in modeling persistence. This is a virtue and a burden at the same time since the estimation of  $d$  is notoriously difficult and troubled by large variances of slowly converging semiparametric estimators. On the other hand, the strong persistence in these series makes promises with regard to their forecastability. The past of long memory time series contains more information about their future compared to less persistent series, which can be exploited to produce high quality forecasts even for quite long forecasting horizons if the above-mentioned problems can be overcome.

Driven by the empirical relevance and the theoretical challenges in modeling and estimation a large literature on long memory time series and FI has emerged since the seminal work by Granger and Joyeux (1980) and Geweke

and Porter-Hudak (1983), with one major motivation being forecasting. Nevertheless, there is a lot of mixed evidence in the literature focusing explicitly on forecasting under long memory and many open questions remain, which we try to shed light on in this paper. Before taking a look at these unresolved issues, we shortly review the related literature.

Early simulation-based evidence on forecasting under long memory is provided by Ray (1993) and Crato and Ray (1996). They examine the usefulness of stationary autoregressive fractionally integrated moving average [ARFIMA] models for forecasting compared to autoregressions of large order, so-called long autoregressions [LARs]. In their simulations, they find inferior or at most equal predictive performance of the ARFIMA models, which they attribute to the difficulties in model selection and estimation of the long memory parameter  $d$ , and thus question the usefulness of these models for prediction. Smith and Yadav (1994) analyze the effects of using ARIMA( $p,1,q$ ) models for forecasting nonstationary ARFIMA models with  $d$  ranging from 0.7 to 1.3. They conclude that the losses from overdifferencing are minor, even though they do not account for the estimation uncertainty from the estimation of  $d$ , while underdifferencing can lead to more substantial losses. Brodsky and Hurvich (1999) demonstrate that FI models are superior in forecasting long memory time series compared to ARMA(1,1) models. Bhardwaj and Swanson (2006) report gains in predictive performance when using FI models compared to a number of short memory competitors for medium sample sizes. For larger sample sizes, these gains are reported to grow substantially.

Concerning empirical evidence, again Bhardwaj and Swanson (2006) report predictive superiority of FI models for financial data, but less clear-cut evidence for macroeconomic data. They attribute this finding mainly to the much larger sample sizes in finance, which lead to a more precise estimation of the long memory parameter. Furthermore, since Andersen et al. (2003) FI models are used in realized volatility forecasting and several papers analyze their performance compared to short memory competitors, see Hansen and Lunde (2011) for an overview. The evidence is mixed as some authors claim

equal predictive ability of short memory alternatives like the Heterogeneous Autoregressive model [HAR] by Corsi (2009).

Given the state of the literature on forecasting under long memory, several open issues remain: As the short review shows, a central question around which the literature has been circling in the past 20 years by providing evidence in both directions is the following: Does using long memory models when forecasting time series that exhibit strong persistence improve predictive performance or do simple short memory alternatives perform at least as well? Our first contribution lies in providing new evidence to answer this question: In systematic forecasting experiments using classical examples of strongly persistent time series from macroeconomics and finance, namely inflation and realized volatility, as well as simulated series, we find that FI models outperform a wide range of competing methods, including autoregressive [AR] models, exponential smoothing [ES] methods and for the realized volatility series also the HAR model. We come up with the idea of using a fixed- $d$  FI model with  $d = 0.5$ , which almost uniformly outperforms all non-FI competitors and, even for large sample sizes, shows a comparable performance to FI models with estimated  $d$ . Under certain practically relevant circumstances like small estimation samples and more persistent short memory components, we find substantial gains of the fixed- $d$  method, which increase with the forecasting horizon. Picking  $d = 0.5$  is a natural choice at the border of nonstationarity and somehow in the middle between the classical  $I(0)$  and  $I(1)$  paradigms, robustifying the forecasting process against strong persistence and not being troubled with the estimation of the exact value of  $d$ . These surprising results do not only equip us with a simple and successful method for forecasting under long memory, but also provide an explanation for the mixed evidence formerly reported in the literature: Our analysis suggests that when the task is to forecast a strongly persistent time series, the inclusion of a long memory component into the used model is beneficial under all examined circumstances, e.g. over all sample sizes and forecasting horizons. However, FI models with estimated  $d$  may sometimes be outperformed by AR models in small and medium-sized samples as parts

of the literature report. This may be caused by a bad specification of the forecasting method with e.g. an unfavorable estimator of the long memory parameter  $d$  or an inappropriate bandwidth choice for an otherwise well-suited semiparametric estimator. But even if a sensible specification is used this will happen due the fundamentally high uncertainty plaguing the estimation of  $d$ .

As our second contribution we analyze how to exactly construct good forecasts using FI models. Naturally, this becomes relevant after having provided evidence that FI models are useful in forecasting strongly persistent time series, but of course has to be answered in the first place to even establish this evidence. We focus on three crucial aspects here. The first one is the estimation of  $d$ . Because of the inherent difficulty of this task, there is a large literature on it and an ongoing debate, by which applied researchers or forecasting professionals can easily be overwhelmed. We carry out a simulation study to compare parametric, local and global semiparametric estimation methods, focusing on the representatives from the Whittle family of estimators since they have favorable theoretical properties and show a good practical performance in general. For the parametric Whittle estimator we use Akaike's information criterion [AIC] to select a model and make the estimator feasible. We find quite dramatic losses of the parametric and the global semiparametric estimators compared to the local Whittle, which we can trace back to the bad performance of the Whittle estimator even under small amounts of misspecification. Consequently, we use the local Whittle estimator with different ad hoc bandwidth choices for our forecasting experiments. The second aspect is the estimation of the mean, which is usually not even mentioned in the long memory forecasting literature, even though it is well known that this is a hard task under strong persistence. We compare the arithmetic mean to the estimators proposed by Robinson (1994) and Shimotsu (2010) through a simulation study and find that the estimator of Robinson performs best. To the best of our knowledge, the theoretical properties of Robinson's estimator have not been analyzed before. We provide asymptotic theory, which shows that substantial efficiency gains can be

made, explaining our simulation results. The third aspect concerns the following question: How much past information should be optimally used for estimating the model used for forecasting? While earlier work as e.g. Bhardwaj and Swanson (2006) implicitly or explicitly assumes that the complete available past of the series of interest should be used because of the decreased estimation uncertainty, we find that this is not necessarily the case: Our empirical results suggest that a medium-sized rolling forecasting window seems to yield the best forecasting performance, balancing out decreased estimation uncertainty and ongoing gradual structural change optimally.

As outlined in the previous paragraph, the third contribution of this paper is to provide new results of interest also outside the realm of forecasting concerning the estimation of the long memory parameter and the mean under long memory.

The rest of the paper is structured as follows: In the next section we review the classical approach to modeling strong persistence, namely fractional integration. In the third section, the estimation of the long memory parameter is treated, while the fourth section deals with the estimation of the mean. In the fifth section, we describe the forecasting methods, i.e. on the one hand how we specifically implement forecasting with FI models based on the evidence from the previous two sections, and on the other hand what the competing methods are and why we use them. The sixth and the seventh section contain the results from the empirical and theoretical forecasting experiments and finally, the eighth section concludes.

## 2 Modeling long memory

The term strong persistence is used to describe processes, where the more distant past of the process still strongly influences the present, i.e. where the autocorrelation function dies out very slowly. Formally, this can be characterized by the autocovariance function  $\gamma$  (which is defined as  $\gamma(h) = E[(y_t - \mu)(y_{t+h} - \mu)]$  for a process  $\{y_t\}$  with mean  $\mu$ ) not being summable,

$$\sum_{h=0}^H \gamma(h) \rightarrow \infty \quad \text{as } H \rightarrow \infty,$$

or equivalently by the long run variance  $\omega^2 = \sum_{h=-\infty}^{\infty} \gamma(h)$  not being finite.

We will use the terms strong persistence and long memory interchangeably here since a distinction between the two only gets relevant for antipersistent processes, see Hassler (2019).

The most widely used model for strong persistence or long memory is fractional integration of order  $d$ . We briefly recap the fractionally integrated process  $\{y_t\}$  of order  $d$ , for short  $y_t \sim I(d)$ , which relies on the fractional integration operator  $\Delta^{-d}$  with the usual binomial expansion in terms of the lag operator  $L$  and a short memory input process  $\{x_t\}$  defined properly below:

$$y_t = \mu + \Delta^{-d}x_t, \quad 0 < d < \frac{1}{2}, \quad \Delta^{-d} = (1 - L)^{-d} = \sum_{j=0}^{\infty} \pi_j(-d)L^j. \quad (1)$$

A recent discussion of the sequence  $\{\pi_j(-d)\}$  is contained in Hassler (2019, Sect. 5.3). In particular, the coefficients die out as  $j^{d-1}/\Gamma(d)$  and are hence not summable for  $d > 0$ . This causes long memory in that the autocovariances of  $\{y_t\}$  at lag  $h$  converge to zero at rate  $h^{2d-1}$  and hence are not summable, too. This strength of persistence is sometimes also measured by the rate at which the variance of the cumulated process grows:  $\text{Var}(\sum_{t=1}^T y_t)$  diverges with  $T^{2d+1}$ . For these properties to hold true, the input sequence  $\{x_t\}$  has to meet certain requirements. We collect them in the following assumption.

**Assumption 1** *Let  $\{x_t\}$  be a covariance-stationary process with*

$$x_t = c(L)\varepsilon_t = \sum_{j=0}^{\infty} c_j\varepsilon_{t-j}, \quad \text{E}(\varepsilon_t) = 0, \quad \text{E}(\varepsilon_t \varepsilon_{t+h}) = \begin{cases} \sigma^2, & h = 0 \\ 0, & h \neq 0 \end{cases},$$

and with ( $c_0 = 1$ )

$$\sum_{j=0}^{\infty} |c_j| < \infty \quad \text{and} \quad c(1) = \sum_{j=0}^{\infty} c_j \neq 0.$$

The process  $\{x_t\}$  is called integrated of order zero, for short  $x_t \sim I(0)$ . The restriction  $c(1) \neq 0$  implies that the data is not overdifferenced. Consequently, the long-run variance is finite and positive:

$$0 < \omega_x^2 := \sigma^2 \left( \sum_{j=0}^{\infty} c_j \right)^2 < \infty.$$

All stationary and invertible ARMA processes meet Assumption 1 and it further holds that  $\text{Var}(\sum_{t=1}^T x_t)$  grows at rate  $T$ .

The inverse filter, the fractional differencing filter,

$$\Delta^d = (1 - L)^d = \sum_{j=0}^{\infty} \pi_j(d) L^j,$$

removes the persistence or long memory from the FI process:  $\Delta^d y_t = x_t$ . Note that the mean is removed as well by fractional differencing,  $\Delta^d \mu = 0$ , because  $\sum_{j=0}^{\infty} \pi_j(d) = 0$  holds.

The stationary FI process from (1) is often called of type I since the work by Marinucci and Robinson (1999); a truncated version that becomes stationary only asymptotically is called type II. The truncation assumes zero values before  $t = 1$ . We assume such a truncated model instead of (1). The truncated filter becomes

$$\Delta_+^{-d} x_t := \Delta^{-d} x_t \mathbf{1}_{(t>0)}(t) = \sum_{j=0}^{t-1} \pi_j(-d) x_{t-j}, \quad t = 1, \dots, T, \quad (2)$$

where  $\mathbf{1}_{(t>0)}$  is the usual indicator function:

$$\mathbf{1}_{(t>0)}(t) = \begin{cases} 1, & t > 0 \\ 0, & \text{else} \end{cases}.$$

Analogously, one defines  $\Delta_+^d x_t := \Delta^d x_t \mathbf{1}_{(t>0)}(t)$ , such that by construction  $\Delta_+^d \Delta_+^{-d} = \mathbf{1}_{(t>0)}$  or  $\Delta_+^d \Delta_+^{-d} x_t = x_t \mathbf{1}_{(t>0)}(t)$ . Truncated fractional differencing and integration has been discussed more formally e.g. by Johansen (2008). Note that a truncated filter also allows to define truly nonstationary FI processes, i.e. that for the process

$$y_t = \mu + \Delta_+^{-d} x_t, \quad (3)$$



$d$  is unrestricted, in particular  $d \geq \frac{1}{2}$  is not ruled out. This is important as we do not want to rule out nonstationary processes when it comes to forecasting.

Being equipped with a model that can be employed when forecasting under strong persistence, the issues of how to estimate the long memory parameter  $d$  and the mean  $\mu$  come up naturally. These will be discussed in the next two sections.


### 3 Estimating the memory parameter

Due to the inherent difficulties in estimating the long memory parameter  $d$  there is a large literature on this topic. A natural first choice for an estimator is of course the exact maximum likelihood [ML] estimator (see e.g. Dahlhaus (1989) and Dahlhaus (2006)), which unfortunately has several deficiencies in this context, including its restriction to the stationary region, the required gaussianity assumption and the need for estimating the mean. The first two problems can be overcome by the conditional sum of squares estimator (see e.g. Beran (1995), Hualde and Robinson (2011) or Nielsen (2015)), which is a time domain approximation to exact ML, while the Whittle estimator (see e.g. Fox and Taqqu (1986) or Velasco and Robinson (2000)) additionally solves the third problem, being a frequency domain approximation, which makes estimation of the mean unnecessary. Nevertheless, all of these estimators require the full specification of the whole model.

As assuming knowledge of the full model might be quite heroic, semiparametric estimators have been proposed as an alternative to these fully parametric estimators. The most popular class of semiparametric estimators are local or narrowband estimators, which only use frequency domain information, i.e. the periodogram, in a vicinity of the origin, where the long memory component of the process dominates. They consequently only estimate the long memory parameter  $d$  and not all the parameters of a fully specified model. The Geweke Porter Hudak [GPH] estimator (see e.g. Geweke and Porter-Hudak (1983) and Hurvich et al. (1998)) and the local Whittle [LW] estimator (see e.g. Künsch (1987) and Robinson (1995)) and some of

its variants as the most prominent representatives of this class are settled in the frequency domain and thus do not require estimation of the mean, too. While the GPH estimator follows a simple idea and is very easy to implement, the LW estimator does not require a normality assumption and is more efficient. For these estimators, a bandwidth  $m$  has to be chosen, balancing a high variance caused by staying too close to the origin and using too little information and a bias induced by the contamination of the estimation through the short memory component of the process. As data-driven bandwidth selection does not work well in practice (see e.g. Hurvich and Deo (1999), Henry (2001) or Andrews and Guggenberger (2003)), ad hoc bandwidth choices as a function of the sample size  $T$  are often used. We will also follow this road. For example, the bandwidth for the LW estimator is usually chosen as  $m = O(T^\alpha)$  with  $0 < \alpha < 1$ . Besides the choice of the bandwidth, the price to pay for the local semiparametric estimators is a slower rate of decay of the asymptotic variance of  $O(\frac{1}{m})$  compared to  $O(\frac{1}{T})$  for the parametric estimators. For example, the asymptotic variance of the local Whittle estimator is  $\frac{1}{4m}$ . Regarding finite sample performance of the estimators, Nielsen and Frederiksen (2005) performed a systematic comparison within the classes of parametric and local semiparametric estimators through simulations. From the parametric estimators they recommend the use of the Whittle estimator as it clearly dominates its time domain competitors. For the semiparametric estimators, they find that the local Whittle estimator and some of its variants as for example the local polynomial Whittle (see Andrews and Sun (2004)) outperform the GPH estimator and its variants.

So-called global or broadband semiparametric estimators represent an alternative to the other two classes. As their name suggests, they use information over the whole frequency range, but their semiparametric nature comes in through the approximation of the short memory component via a certain parametric model of order  $k$ , which is growing with the sample size. The global semiparametric Whittle [GSW] estimator proposed by Bhansali et al. (2006) uses an AR( $k$ ) model here, while an alternative global estimator works with a log-periodogram regression of a fractional exponential model



**Figure 1.** asymptotic efficiency of  $\hat{d}_{LW}$  with bandwidth  $m = \lfloor T^{0.65} \rfloor$  relative to that of  $\hat{d}_{GSW}$  with lag length  $k = \lfloor \log(T) \rfloor$

(see Moulines and Soulier (1999), Moulines and Soulier (2000) and Hurvich and Brodsky (2001)). These estimators have the nice property that their asymptotic variance decays with a faster rate than the one of the semiparametric estimators. For example, the GSW has an asymptotic variance of  $\frac{k}{T}$ , where  $k = O(\log(T))$  is usually picked. Nevertheless, a closer comparison of the asymptotic variances of the two classes of estimators reveals a different picture due to the constants involved. Figure 1 shows the relative asymptotic efficiency of the LW relative to the GSW estimator, where the widely used choices  $m = \lfloor T^{0.65} \rfloor$  and  $k = \lfloor \log(T) \rfloor$  are picked. Thus, for all reasonable sample sizes, the uncertainty associated with the local Whittle estimator is smaller. To the best of our knowledge, the finite sample performance of global semiparametric estimators has not been systematically analyzed yet. For more details on the estimators and their properties see e.g. Beran et al. (2013) or Hassler (2019) and further references therein.

When it comes to forecasting specifically, a wide array of different estimators have been used in literature. There seems to be no consensus on which estimator to use, on the contrary, there is an ongoing debate. Baillie et al.

(2012) advocate the use of the maximum likelihood estimator, which they compare to the local Whittle. In a comment to this Arteche (2012) stresses that their results are unsurprising as they compared the performance of a parametric and a semiparametric estimator under a known model and did not really touch upon the issues of model misspecification and model selection.

Arising from the above considerations are three basic types of feasible estimators: Locally semiparametric estimators, globally semiparametric estimators and parametric estimators plus a model selection step. Unfortunately, we are not aware of any work that compares these three types of estimation procedures or two of them. To do this, the model selection step for the parametric estimators or at least an analysis of the consequences of model misspecification would have to be included. Furthermore, we are not aware of any analysis of the finite sample performance of global semiparametric estimators. To clarify these open issues before choosing estimators for forecasting, we perform a small simulation study. As touched upon above, the Whittle and the local Whittle estimator are favorable in terms of theoretical properties and finite sample performance within the classes of parametric and local semiparametric estimators respectively. Consequently, we use these two as representatives of the respective classes.<sup>1</sup> As the global semiparametric Whittle estimator is theoretically more appealing than the log-periodogram regression of the fractional exponential model, i.e. being more efficient and not needing a normality assumption, and as there are no studies concerning their finite sample performance, we use the GSW as a representative of the

---

<sup>1</sup>For the class of semiparametric estimators, we could have used variants of the local Whittle as well, e.g. the local polynomial Whittle, the fully extended local Whittle (see Abadir et al. (2007)) or the exact local Whittle estimator (see Shimotsu and Phillips (2005) and Shimotsu (2010)), which can even lead to improvements in some circumstances. In a recent study, Cheung and Hassler (2018) present experimental evidence from a systematic comparison of six variants of Whittle type estimators. Depending on the range of  $d$  they all have their merits and drawbacks, but overall there is no estimator dominating local Whittle uniformly. As it would not have been clear which from these estimators to choose for general forecasting purposes and as a large scale comparison of estimators is not within the scope of this paper, we use the basic variant of the local Whittle.

global semiparametric estimators in our study.

Even though consistency and limiting normality for the GSW estimator have only been established for the stationary region and under a restricted optimization window using a pre-estimator, we do not use this restriction and a pre-estimator since this would exclude the nonstationary region, which is not practical. The remaining procedure is very simple: An ARFI( $p, d$ ) model is used, where  $p = \lfloor \log(T) \rfloor$  and then the Whittle objective function is minimized to estimate  $d$ . The procedure for the Whittle estimator is very similar: We again assume an ARFI( $p, d$ ) model, pick  $p$  by Akaike's information criterion<sup>2</sup> (see Beran et al. (1998)) and then estimate  $d$  by minimizing the respective Whittle objective function. We use the AIC instead of alternative selection criteria as it is the usual choice in a forecasting context. For the local Whittle estimator, we pick the bandwidth as  $m = \lfloor T^\alpha \rfloor$  with  $\alpha \in \{0.5, 0.65, 0.8\}$ , which is standard in the literature.

As data generating processes we use FI( $d$ ) processes of type II, i.e. generated from model (3), with two different values of  $d$ ,  $d \in \{0.4, 0.7\}$ , and three different input processes for  $\{x_t\}$  from Assumption 1 with increasing persistence. We measure their degree of persistence by the variance ratio of long-run variance to variance,  $VR = \omega_x^2 / \gamma_x(0)$ . For standard normal iid innovations  $\{\varepsilon_t\}$ , the models are

$$iid : x_t = \varepsilon_t \quad \text{with } VR = 1,$$

$$AR(1) : x_t = 0.5 x_{t-1} + \varepsilon_t \quad \text{with } VR = 3,$$

$$MA(9) : x_t = \varepsilon_t + \frac{9}{10} \varepsilon_{t-1} + \dots + \frac{1}{10} \varepsilon_{t-9} \quad \text{with } VR = 7.857.$$

In Table 1 we present the mean squared errors (MSEs) for  $d = 0.4$  using 1000 repetitions from estimating  $d$  with the local Whittle estimator with bandwidth  $m = \lfloor T^\alpha \rfloor$  [ $LW(T^\alpha)$ ], the global semiparametric Whittle estimator and the Whittle estimator plus AIC [ $W(AIC)$ ]. For comparison we also present the Whittle estimator for ARFIMA( $p, d$ ) models for different values of  $p$  fixed a priori [ $W(p = p^*)$ ]. We report results for four different sample

---

<sup>2</sup>We set the maximum lag order to  $p_{max} = \lfloor \log(T) \rfloor$ , too.

sizes  $T \in \{60, 300, 1500, 7500\}$ . The smallest MSE in each row is marked in bold face.

**Table 1.** MSEs of different estimators for  $d$  for simulated FI(0.4) processes with short memory component  $x_t$  and sample size  $T$ , 1000 repetitions

$x_t$	$T$	$LW(T^{0.5})$	$LW(T^{0.65})$	$LW(T^{0.8})$	$GSW$	$W(AIC)$	$W(p=0)$	$W(p=1)$	$W(p=3)$	$W(p=10)$
iid	60	0.1155	0.0405	<b>0.0171</b>	10.9716	11.4711	0.0212	0.4328	4.3762	14.0576
	300	0.0291	0.0092	0.0036	0.9612	0.7556	<b>0.0026</b>	0.02	0.2101	2.1613
	1500	0.0094	0.0027	0.0009	0.0331	0.0226	<b>0.0005</b>	0.0013	0.0051	0.0815
	7500	0.0037	0.0009	0.0002	0.0015	0.0005	<b>0.0001</b>	0.0002	0.0005	0.002
AR(1)	60	0.1322	<b>0.1135</b>	0.151	9.863	10.0318	0.1752	0.1357	4.0548	11.2652
	300	0.0316	<b>0.0265</b>	0.0966	1.6752	1.6751	0.1746	0.0532	0.7963	2.2098
	1500	0.0099	<b>0.005</b>	0.05	0.0674	0.0534	0.1725	0.0074	0.0167	0.1243
	7500	0.0035	0.0011	0.0229	0.0018	0.0011	0.1721	<b>0.0007</b>	0.0009	0.0022
MA(9)	60	0.5093	0.6676	0.5636	8.4441	7.4402	0.6814	<b>0.0722</b>	3.0257	10.1518
	300	0.0438	0.3271	0.62	1.8444	1.4754	0.7735	<b>0.0176</b>	0.8141	1.5582
	1500	0.0101	0.0393	0.5601	0.061	0.0665	0.8047	<b>0.0097</b>	0.1352	0.1845
	7500	<b>0.0038</b>	0.0041	0.3875	0.0058	0.005	0.8144	0.0101	0.0107	0.0186

As expected, for the iid short memory component the correctly specified Whittle estimator with  $p = 0$  performs best. The local Whittle is clearly the best of the feasible estimators, irrespective of the bandwidth choice. Nevertheless, a larger bandwidth is of course a better choice in this case as there are no short memory dynamics, which could cause a bias. Very surprising at the first glance are the extremely bad results of the Whittle estimator plus AIC and the global semiparametric Whittle estimator. Only for  $T = 7500$  their MSEs have the same order of magnitude as the ones of the local Whittle estimators. For the more persistent short memory components the same picture emerges regarding the comparison of the three feasible estimators. For the local Whittle estimator the smaller bandwidth choices now lead to improvements due to the increased short memory dynamics. A look at the last four columns of the table helps to understand why the Whittle estimator plus AIC and the local semiparametric Whittle estimator perform so badly. This seems to be caused by the drastic drop in performance of the Whittle estimator under even slight misspecification. For example if the model order is picked as 1 instead of 0 or 0 instead of 1, the accuracy most often drops below the one of the local Whittle estimator. For larger magnitudes of

misspecification, the accuracy continues to decrease strongly.<sup>3</sup> Thus, as the AIC or the deterministic order choice by the global semiparametric Whittle estimator often do not pick exactly the correct model, this slight misspecification in combination with the bad performance of the Whittle estimator under slight misspecification lead to these results.

**Table 2.** MSEs of different estimators for  $d$  for simulated FI(0.7) processes with short memory component  $x_t$  and sample size  $T$ , 1000 repetitions

$x_t$	$T$	$LW(T^{0.5})$	$LW(T^{0.65})$	$LW(T^{0.8})$	$GSW$	$W(AIC)$	$W(p=0)$	$W(p=1)$	$W(p=3)$	$W(p=10)$
iid	60	0.11185	0.03803	<b>0.01858</b>	11.39186	11.56437	0.01929	0.41093	4.37725	14.32783
	300	0.02869	0.00969	0.00389	0.94002	0.63425	<b>0.0027</b>	0.0176	0.20995	2.11467
	1500	0.01002	0.00317	0.00106	0.02898	0.01844	<b>0.0006</b>	0.00143	0.00431	0.07296
	7500	0.00421	0.0012	0.00028	0.00168	0.0008	<b>0.00014</b>	0.00031	0.00071	0.00211
AR(1)	60	0.12669	<b>0.09768</b>	0.10569	10.63713	9.72827	0.14552	0.16326	4.16247	12.31798
	300	0.03075	<b>0.02805</b>	0.08284	1.51804	1.23276	0.15889	0.05856	0.71517	2.14612
	1500	0.01097	<b>0.00658</b>	0.04769	0.06314	0.03568	0.16459	0.0077	0.01453	0.11269
	7500	0.0043	0.00171	0.02288	0.00201	0.00148	0.1689	<b>0.00103</b>	0.00107	0.00247
MA(9)	60	0.3995	0.42793	0.30085	11.03323	8.03117	0.40488	<b>0.13162</b>	3.8611	13.02868
	300	<b>0.04445</b>	0.25644	0.40295	2.19134	1.18821	0.5054	0.08068	0.68322	1.93014
	1500	<b>0.01161</b>	0.04029	0.45088	0.07743	0.08606	0.61862	0.0383	0.10376	0.16573
	7500	<b>0.00438</b>	0.00494	0.35516	0.00528	0.00568	0.70029	0.01046	0.00698	0.01254

For different strengths of long memory the conclusions do not change as can be seen from Table 2, which contains the results for nonstationary FI processes with  $d = 0.7$ .

The results of our simulation study suggest that global semiparametric estimators and parametric estimators with the prior use of a model selection criterion should be used with care. Only for very high sample sizes a reasonable performance can be expected. It is a safer choice to stick with a local semiparametric estimator like the local Whittle, which we will do throughout the rest of the paper.

<sup>3</sup>An exception here is that the Whittle estimator with  $p = 1$  performs quite well for the process with MA(9) innovations. This is probably due to the similarity of this MA(9) process to an AR(1) process with a coefficient of about 0.9.

## 4 Estimating the mean

Estimating the mean is naturally quite difficult under strong persistence since long trending spells of the time series away from its mean are common. This is reflected in the behavior of the sample mean  $\bar{y} = T^{-1} \sum_{t=1}^T y_t$ . From the behavior of  $\text{Var}(\sum_t y_t)$  given above, it is clear that  $\bar{y}$  converges all the more slowly to  $\mu$  the longer the memory is, see Samarov and Taqqu (1988). To see if improvements over  $\bar{y}$  are possible, in this section we analyze the behavior of the alternative mean estimators proposed by Robinson (1994) and Shimotsu (2010) theoretically as well as through simulations.

Note that an application of the truncated fractional difference filter to (3) results in a time-varying mean:

$$\Delta_+^d y_t = r_t \mu + x_t, \quad r_t = \sum_{j=0}^{t-1} \pi_j(d). \quad (4)$$

This suggests to follow the proposal by Robinson (1994) and to regress the filtered data  $\Delta_+^d y_t$  on the variable  $r_t$  to determine  $\hat{\mu}$  by ordinary least squares [OLS] from (4). Such an estimation is justified by the maximum likelihood principle if  $\{x_t\}$  is normal white noise. Robinson (1994) establishes that removing  $r_t \hat{\mu}$  does not affect the limiting distribution of his subsequent test. It seems that a comparison of this estimator with the arithmetic mean has never been carried out.

First, we determine the limiting variance under stationarity (more precisely: under  $d < 1/2$ ) in order to evaluate the efficiency.


**Proposition 1** *Assume Assumption 1 in model (4) for  $0 < d < 1/2$ . It then holds that*

$$T^{1-2d} \text{Var}(\hat{\mu}) \rightarrow \frac{\omega_x^2 (1-2d)(\pi d)^2}{\Gamma^2(d+1) \sin^2(\pi d)}$$

as  $T \rightarrow \infty$ , where  $\Gamma$  is the gamma function.

PROOF See Appendix.





**Figure 2.**  $REff(d)$  according to Proposition 1

The same rate applies for the arithmetic mean. Under mildly stronger assumptions Tanaka (1999, Coro. 2.3) found that

$$T^{1-2d}\text{Var}(\bar{y}) \rightarrow \frac{\omega_x^2}{\Gamma^2(1+d)(2d+1)},$$

see also Marinucci and Robinson (2000, Thm. 1). Let us define as relative efficiency

$$REff(d) := \lim_{T \rightarrow \infty} \frac{\text{Var}(\hat{\mu})}{\text{Var}(\bar{y})} = (1 - 4d^2) \frac{(\pi d)^2}{\sin^2(\pi d)}.$$

The graph of this function in Figure 2 shows that there is plenty of room for efficiency gains over the arithmetic mean for larger values of  $d$ , at least asymptotically. Note, however, that  $\hat{\mu}$  is a function of  $d$ ,  $\hat{\mu} = \hat{\mu}(d)$ , since the filter  $\Delta_+^d$  depends on  $d$ . A preliminary estimation  $\hat{d}$  is required to construct a feasible estimator  $\hat{\mu}(\hat{d})$ .

Next, we turn to the nonstationary case. It is clear that the sample mean diverges for  $d \geq 1/2$ : In fact, we know from Marinucci and Robinson (2000, Thm. 1) that  $\text{Var}(\bar{y})$  diverges with  $T^{2d-1}$  for  $d > 1/2$ . Note that the variance of  $\hat{\mu}$  from Proposition 1 also vanishes only under  $d < 1/2$ . For nonstationary

fractionally integrated processes the estimator is, however, still defined, and its variance is

$$\text{Var}(\widehat{\mu}(d)) = \frac{\gamma_x(0) \sum_{t=1}^T r_t^2 + 2 \sum_{h=1}^{T-1} \gamma_x(h) \sum_{t=1}^{T-h} r_t r_{t+h}}{\left(\sum_{t=1}^T r_t^2\right)^2}.$$

For  $d > 1/2$  the regressors  $r_t$  converge to zero so quickly that  $\sum_{t=1}^T r_t^2$  remains finite for growing  $T$  and  $\mu$  is not estimated consistently, but, contrary to  $\bar{y}$ , the estimator at least has a finite variance.

Shimotsu (2010) considered the starting value as an estimator under non-stationarity,  $\tilde{\mu} = y_1$  with  $\tilde{\mu} - \mu = x_1$ . Note that  $\widehat{\mu}(1) = \tilde{\mu}$ . For general  $d$ , the ratio of the variances of these two estimators is

$$\frac{\text{Var}(\widehat{\mu}(d))}{\text{Var}(\tilde{\mu})} = \frac{1}{\sum_{t=1}^T r_t^2} + 2 \sum_{h=1}^{T-1} \frac{\gamma_x(h) \sum_{t=1}^{T-h} r_t r_{t+h}}{\gamma_x(0) \left(\sum_{t=1}^T r_t^2\right)^2}.$$

If  $\{x_t\}$  was white noise, then  $\text{Var}(\widehat{\mu}(d))/\text{Var}(\tilde{\mu}) = 1/\sum_{t=1}^T r_t^2 \leq 1$ , but under serial correlation of  $\{x_t\}$  it is not clear which estimator is more precise. The full proposal of Shimotsu is to use either the arithmetic mean, the starting value or a linear combination of both, depending on the value of  $d$ :

$$\tilde{\mu}(d) = v(d) \bar{y} + (1 - v(d)) y_1,$$

where

$$v(d) = \begin{cases} 1, & d \leq 1/2 \\ \frac{1+\cos(4\pi d)}{2}, & 1/2 < d < 3/4 \\ 0, & d \geq 3/4. \end{cases}$$

A feasible version of this estimator  $\tilde{\mu}(\hat{d})$  also requires a preliminary estimation of the long memory parameter.

To assess how the theoretical findings concerning these three estimators carry over to practical performance and how the feasible estimators work, we conducted a simulation study. As data generating processes, we used again FI( $d$ ) processes, i.e. model (3), with the same three short memory

**Table 3.** MSEs of different estimators for  $\mu$  for simulated FI( $d$ ) processes with short memory component  $x_t$  and sample size  $T$ ,  $d$  is estimated by LW with  $m = \lfloor T^{0.65} \rfloor$ , 10000 repetitions

d	$x_t$	T = 60			T = 300			T = 1500			T = 7500		
		$\bar{y}$	Shimotsu	Robinson	$\bar{y}$	Shimotsu	Robinson	$\bar{y}$	Shimotsu	Robinson	$\bar{y}$	Shimotsu	Robinson
0.2	iid	<b>0.0714</b>	0.0792	0.0843	<b>0.0281</b>	<b>0.0281</b>	0.0285	0.0107	0.0107	<b>0.0104</b>	0.0041	0.0041	<b>0.0039</b>
	AR(1)	<b>0.291</b>	0.4799	0.4733	<b>0.1073</b>	0.108	0.1275	0.0417	0.0417	<b>0.0408</b>	0.0158	0.0158	<b>0.0153</b>
	MA(9)	<b>2.1529</b>	3.765	3.6557	<b>0.8372</b>	3.3225	3.1316	<b>0.3194</b>	0.3195	0.4502	0.1217	0.1217	<b>0.1196</b>
0.4	iid	0.307	0.3383	<b>0.3043</b>	0.229	0.2283	<b>0.1904</b>	0.1655	0.1654	<b>0.1232</b>	0.1195	0.1195	<b>0.0827</b>
	AR(1)	1.248	1.1351	<b>1.0047</b>	0.8761	0.7844	<b>0.6641</b>	0.6445	0.6364	<b>0.442</b>	0.4669	0.4669	<b>0.3155</b>
	MA(9)	9.2067	3.8866	<b>3.6284</b>	6.8461	3.9072	<b>3.9044</b>	4.9629	<b>3.2068</b>	3.2374	3.602	3.597	<b>2.2469</b>
0.6	iid	1.2811	0.9747	<b>0.7306</b>	1.8127	1.054	<b>0.6361</b>	2.4894	1.1748	<b>0.5631</b>	3.4138	1.4086	<b>0.5388</b>
	AR(1)	5.1844	1.5449	<b>1.4219</b>	6.925	1.5802	<b>1.374</b>	9.6651	2.2303	<b>1.4011</b>	13.3503	3.9409	<b>1.459</b>
	MA(9)	38.0374	3.9581	<b>3.3317</b>	54.1704	3.9209	<b>3.514</b>	74.725	<b>3.9517</b>	5.1493	103.2586	12.5138	<b>7.1468</b>

input processes  $\{x_t\}$  with increasing persistence. We used three different values of  $d$ , namely  $d \in \{0.2, 0.4, 0.6\}$ , and as sample sizes we used again  $T \in \{60, 300, 1500, 7500\}$ . Table 3 presents the MSEs over 10000 repetitions, where at each instance the estimator with the smallest MSE is marked in bold face again. Overall, the estimator proposed by Robinson dominates the other two, while Shimotsu's estimator usually shows the second best performance and the arithmetic mean the worst. The gains over the arithmetic mean can be quite large and are rising with  $d$  and with the persistence of the short memory component. Only for  $d = 0.2$  and the smaller sample sizes the arithmetic mean outperforms the other two estimators. But in general, it seems to be worth to accept the increased uncertainty coming with the estimation of  $d$  in the first step of the feasible estimators to overall get a more precise estimate of the mean.

Hence, we will work with  $\hat{\mu}(d)$  for a given  $d$  and with  $\hat{\mu}(\hat{d})$  for estimated  $\hat{d}$  in the next sections when it comes to forecasting real and simulated series using FI models since this estimator is consistent for  $d < 1/2$ , has a finite variance for larger orders of integration and shows a good performance in our simulations.

An alternative strategy would be to ignore the mean and simply work with  $\Delta_+^d y_t$  because  $\Delta_+^d 1 \rightarrow \Delta^d 1 = 0$  for  $t \rightarrow \infty$ :  $\Delta_+^d y_t \approx x_t$ . Clearly, this cannot be recommended since  $\Delta_+^d 1 = \sum_{j=0}^{t-1} \pi_j(d)$  for small  $d$  converges very slowly with  $t$  getting large.

## 5 Forecasting methods

In this section we present in detail how we use the FI model to construct forecasting methods for strongly dependent time series and we describe the relevant competing methods, which do not employ long memory models. We assume a stretch of data of size  $T$ ,  $y_1, \dots, y_T$ , used to forecast  $h$  steps ahead:

$\widehat{y}_{T+h}$ .

We use methods that explicitly account for long memory in that they (try to) remove the strong persistence from the series by filtering with a fractional difference filter, forecast the filtered data, and then recolor these forecasts by applying the fractional integration filter. Because of the findings from the previous section, the mean is accounted for from the filtered data by Robinson's method. With model (4) in mind we hence proceed as follows for a given  $d$ . Compute  $\Delta_+^d y_t$  in order to estimate  $\widehat{\mu}$ , and save the residuals:

$$\xi_t := \Delta_+^d y_t - r_t \widehat{\mu}, \quad t = 1, \dots, T.$$

They are fed into an autoregression on  $k$  lags to account for short memory in  $\{x_t\}$ , where  $k$  is determined by Akaike's information criterion, AIC:<sup>4</sup>

$$\xi_t = \widehat{a}_1 \xi_{t-1} + \dots + \widehat{a}_k \xi_{t-k} + \widehat{\varepsilon}_t, \quad t = k+1, \dots, T.$$

The autoregressive scheme is then used to recursively forecast  $h$  steps ahead:<sup>5</sup>

$$\widehat{\xi}_{T+h} := \widehat{a}_1 \widehat{\xi}_{T+h-1} + \dots + \widehat{a}_k \widehat{\xi}_{T+h-k}. \quad (5)$$

Finally, forecasts of the original sequence are given by

$$\widehat{y}_{T+h} = \widehat{\mu} + \Delta_+^{-d} \widehat{\xi}_{T+h}. \quad (6)$$

---

<sup>4</sup>Note that no intercept is included as the original series has already been demeaned. Furthermore, the inclusion of an intercept clearly deteriorated the forecasting performance in our experiments. The maximum lag length considered is  $12 \lfloor (T/100)^{0.25} \rfloor$ .

<sup>5</sup>Consider some forecasted sequence  $\widehat{z}_{T+h-j}$ . Whenever  $T+h-j \leq T$ , then  $\widehat{z}_{T+h-j} = z_{T+h-j}$ .

Because of  $\Delta_+^{-d}\Delta_+^d = \mathbf{1}_{(t>0)}$ , this two-step procedure is equivalent to computing the recursion

$$\widehat{y}_{T+h} = - \sum_{j=1}^{T+h-1} \pi_j(d)\widehat{y}_{T+h-j} + r_t\widehat{\mu} + \widehat{\xi}_{T+h},$$

which is also often applied when forecasting with FI models.

Before the fractional differencing filter  $\Delta_+^d$  can be applied, a value of  $d$  has to be chosen. The estimation of the long memory parameter has been discussed in the third section. Resulting from our analysis of the different estimators there, we use the local Whittle estimator with bandwidth  $m$  that we pick as  $m = \lfloor T^\alpha \rfloor$  with  $\alpha \in \{0.5, 0.65, 0.8\}$  (In the tables throughout the next two chapters this method is abbreviated by FI( $T^\alpha$ )). We use three different bandwidth choices to analyze if the forecasting performance strongly depends on the bandwidth.

Even though the local Whittle estimator yields good estimation results compared to other estimators, it still has a high variance due to the inherent difficulty of estimating the long memory parameter. Motivated by this, we propose the unorthodox method of using a fixed long memory component in the forecasting method, i.e. a FI model with an a priori fixed parameter  $d$ , to circumvent the estimation problem. A natural choice here is to set  $d = 0.5$  as this is the border to nonstationarity and somehow the middle between the classical  $I(0)$  and  $I(1)$  paradigms. Furthermore, the dynamics of many economic time series are believed to lie somewhere close to the borderline between stationarity and nonstationarity. The exact choice of  $d$  does probably not matter too much anyway - as long as it is not as far away from the truth as the outcomes of classical long memory estimators may often be - since the flexible short memory component in form of the AR model may adapt to the fixed long memory component to a certain degree. Thus, using a FI(0.5) model within the forecasting method (which will also be abbreviated by FI(0.5)) may be a good middle ground, robustifying the forecasting procedure against strong persistence by being able to account for slowly decaying autocorrelations, but not running into problems with the

very volatile long memory estimators.

One classical competitor of the true fractionally integrated model is its special case from the  $I(1)$  world, at which we arrive by setting  $d = 1$ . The FI(1) model, which is an ARI( $p,1$ ) model when using the method described above, plays an important role in modeling and forecasting economic time series since the seminal work by Nelson and Plosser (1982) and is thus a relevant competitor, appearing in the long memory forecasting literature since Smith and Yadav (1994). Another FI(1) model that is a classical benchmark and a sanity check is the no change forecast,  $\hat{y}_{T+h} = y_T$ , arising from a random walk, which we also included in our experiments. As it was dominated by the other procedures, we do not report its results in the tables throughout the rest of the paper.

The fiercest competitor from the literature is certainly the autoregressive model. Since Ray (1993), it has been repeatedly argued that modeling fractional integration does not provide better forecasts than simply fitting an autoregression, in this context often called long autoregression. The autoregression (now with intercept) of the unfiltered series (as we are in the special case of a FI(0) model) builds again on a lag length  $k$  determined by AIC:

$$y_t = \hat{\nu} + \hat{a}_1 y_{t-1} + \dots + \hat{a}_k y_{t-k} + \hat{\varepsilon}_t, \quad t = k + 1, \dots, T.$$

Autoregressions are the classical benchmark in forecasting economic time series, being simple and at the same time very hard to beat, see e.g. Stock and Watson (1998). For forecasting under long memory this competitor is even more relevant due to the repeated empirical and simulation-based claims that forecasts based on LAR models would be superior to forecasts based on FI models and even theoretical claims that they would be able to capture true long memory. For a clarification and a recent theoretical underpinning of the latter see Demetrescu and Hassler (2016). As a simplification of the LAR model without the need of selecting a lag length, an often used benchmark is the AR(1) model, which we also included in our experiments. As it is dominated by the other procedures, we refrain from reporting its results, too.

Besides the autoregressive forecasts, the second dominant benchmark method used in the general forecasting literature is exponential smoothing. It is known to perform well as a highly flexible general-purpose forecasting method for a wide variety of situations, see e.g. De Gooijer and Hyndman (2006). Some theoretical work has provided a solid foundations for the method in form of state space models (see Hyndman et al. (2002)) as well as demonstrated that exponential smoothing arises as the optimal method in certain highly relevant situations, for example under certain forms of structural change, see e.g. Clements and Hendry (2006) and Pesaran et al. (2013). Thus, we also include exponential smoothing in our forecasting experiments even though it has traditionally not been used in the long memory forecasting literature.<sup>6</sup>

Finally, we also used the arithmetic mean,  $\hat{y}_{T+h} = \bar{y}$ , as a standard benchmark and sanity check and to get an idea up to which forecasting horizon the forecasts from the other methods contain valuable information.

## 6 Empirical evidence

In this section we present the results of systematic pseudo-out-of-sample [POOS] forecasting experiments on real time series with the methods described in the previous section to assess their predictive performance. We use inflation and realized volatility series as these are classical examples of real-world long memory time series from economics and finance. We present robustness checks in several directions, changing amongst others the time period and the country. Finally, we also assess the effect of the estimation window size and compare rolling to expanding window forecasts.

---

<sup>6</sup>We use the `ets` function from the R `forecast` package, where a trend component is chosen or not (additive or multiplicative, with or without dampening) according to (corrected) AIC. We do not use a seasonal component. The reference for the forecast package is: Hyndman R., G. Athanasopoulos, C. Bergmeir, G. Caceres, L. Chhay, M. O'Hara-Wild, F. Petropoulos, S. Razbash, E. Wang, F. Yasmeeen (2019). `forecast`: Forecasting functions for time series and linear models. R package version 8.5, <http://pkg.robjhyndman.com/forecast>.

## 6.1 Inflation

Inflation is the classical example of a macroeconomic variable exhibiting long memory. Since Hassler and Wolters (1995) and Baillie et al. (1996) many papers have analyzed inflation time series using long memory models. Surprisingly, in the long memory forecasting literature, there is not much work on inflation forecasting. Bhardwaj and Swanson (2006) for example use the complete Stock and Watson dataset to analyze the forecasting performance of FI models for macroeconomic data, which of course also contains many variables that do not show long memory characteristics and are thus better modeled and forecasted by other approaches. As our goal is to analyze how to best forecast time series which show long memory behavior, we focus on variables which do so.

Firstly, we use monthly year-on-year US consumer price inflation.<sup>7</sup> The series spans from January 1948 to November 2017 and is shown in part (a) of Figure 3. The autocorrelogram of the series shown in part (b) of this figure decays very slowly as is characteristic for long memory time series. We ran a pseudo-out-of-sample forecasting experiment with a rolling estimation window of size 180 for six different forecasting horizons  $h \in \{1, 3, 6, 12, 24, 48\}$ . The predictive performance of the competing methods is measured by the mean squared forecast error. In Table 4, we report relative mean squared forecast errors in relation to the method based on the FI(0.5) model. The smallest relative mean squared error is marked in bold face, the second smallest in italics.

Regarding the results, a first striking observation is that the FI(0.5) method uniformly outperforms all methods based on short memory models over all horizons. From these short memory methods only LAR shows a reasonable performance with losses between 2 and 17 percent compared

---

<sup>7</sup>We calculated the inflation from the US consumer price index without seasonal adjustment. The exact reference for the series is: U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items [CPIAUCNS], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CPIAUCNS>, January 3, 2018.




to  $\text{FI}(0.5)$ , which are increasing with the forecasting horizon. This suggests that modeling long memory explicitly when forecasting under long memory seems to be beneficial even in this very simple form of a fixed long memory component. Furthermore, when forecasting further into the future, this seems to become more and more important.


A second striking and at first rather surprising observation concerns the performance of the methods based on a FI model with estimated  $d$ : Even though we have taken great care in choosing a well-performing estimator compared to the existing alternatives with the local Whittle, the  $\text{FI}(0.5)$  method again almost uniformly outperforms  $\text{FI}(T^{0.5})$ ,  $\text{FI}(T^{0.65})$  and  $\text{FI}(T^{0.8})$ . Thus, irrespective of the choice of the bandwidth, here the fixed- $d$  method is clearly superior to the methods with estimated long memory component and the gap widens with increasing forecasting horizon. The latter observation will be analyzed in detail in the next section with the help of simulations. As discussed in the previous section, including a fixed long memory part in a forecasting method could thus provide a good way to robustify against long memory, while not being plagued by the problems usually arising when one tries to estimate the exact strength of the memory. Although we have chosen  $d = 0.5$ , picking any value between 0 and 1 for  $d$  is of course possible and can lead to improvements upon the limiting cases LAR and  $\text{FI}(1)$ , which both perform already quite well.

Thirdly, looking further at the performance of the estimated- $d$  long memory methods, they perform better than the best short memory method, i.e. LAR, irrespective of the bandwidth chosen. The only exception is the very large forecast horizon  $h = 48$ , where their performance strongly drops off. Comparing the three different bandwidth choices no pattern emerges which would indicate an optimal bandwidth choice here, even though the performance clearly varies with the bandwidth. This indicates that choosing a bandwidth here is indeed a difficult task.

To find out whether these observations continue to hold more generally, we first take a look at several robustness checks, then move on to another type of time series, namely financial time series in the form of realized volatility, and



(a) timeplot



(b) autocorrelogram

**Figure 3.** plots for monthly year-on-year US CPI inflation from January 1948 to November 2017

**Table 4.** MSEs relative to FI(0.5) from POOS forecasting experiments for monthly year-on-year US CPI inflation, January 1948 to November 2017, rolling window of size 180

horizon	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean
$h = 1$	<i>1.01</i>	1.015	<b>1.007</b>	1.026	1.021	1.426	74.922
$h = 3$	<b>1.007</b>	<i>1.019</i>	1.021	1.033	1.037	1.252	13.686
$h = 6$	<b>0.998</b>	<i>1.008</i>	1.022	1.022	1.065	1.167	5.602
$h = 12$	1.038	<i>1.022</i>	1.036	<b>1.015</b>	1.078	1.152	2.268
$h = 24$	1.157	<i>1.063</i>	<i>1.063</i>	<b>1.012</b>	1.169	1.34	1.3
$h = 48$	1.802	1.306	1.206	<b>1.11</b>	<i>1.16</i>	1.819	1.292

finally use simulated series, too. We executed robustness checks in several directions, starting with a shorter time span of the original inflation series, then switching from year-on-year to month-on-month inflation and finally changing the country, looking at German inflation. Furthermore, at the end of this section we look at an analysis of the effects of the size of the rolling window and of the use of an expanding window. Table 5 includes the results for the shortened series starting from 1970 to analyze the effects of the specific time period used. Qualitatively nothing changes and all observations from above continue to hold.

**Table 5.** MSEs relative to FI(0.5) from POOS forecasting experiments for monthly year-on-year US CPI inflation, January 1970 to November 2017, rolling window of size 180

horizon	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean
$h = 1$	<b>1.01</b>	1.024	<i>1.011</i>	1.035	1.024	1.518	41.128
$h = 3$	<b>1.004</b>	<i>1.035</i>	1.037	1.054	1.038	1.314	7.008
$h = 6$	<b>0.986</b>	<i>1.021</i>	1.044	1.048	1.074	1.195	3.037
$h = 12$	<b>1.065</b>	1.1	1.127	<i>1.097</i>	1.115	1.252	1.635
$h = 24$	1.282	<i>1.203</i>	1.248	<b>1.115</b>	1.379	1.337	1.596
$h = 48$	1.513	1.299	1.231	<b>0.962</b>	1.085	<i>1.034</i>	1.194

In the next step we change the definition of inflation and look at an-

annualized month-on-month inflation from January 1948 to November 2017.<sup>8</sup> Naturally, month-on-month inflation is much more volatile and less persistent than year-on-year inflation, which is visible in Figure 4. Despite the quite different characteristics of the series, the pseudo-out-of-sample forecasting results presented in Table 6 are very similar: FI(0.5) is still clearly the best method, even though the methods with estimated  $d$  now perform better than before. Furthermore, exponential smoothing now performs much better than before even though still not better than the long autoregression or the long memory methods.

**Table 6.** MSEs relative to FI(0.5) from POOS forecasting experiments for annualized month-on-month US CPI inflation, February 1947 to November 2017, rolling window of size 180


horizon	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean
$h = 1$	1.018	<b>1.003</b>	<b>1.003</b>	1.033	1.013	1.044	1.652
$h = 3$	1.027	<i>0.999</i>	1.013	1.084	1.022	<b>0.988</b>	1.348
$h = 6$	1.024	<b>1</b>	<i>1.008</i>	1.074	1.048	1.064	1.399
$h = 12$	<i>1.02</i>	<b>1.008</b>	1.023	1.022	1.05	1.073	1.216
$h = 24$	<i>1.008</i>	<b>0.993</b>	1.028	1.109	1.147	1.175	1.094
$h = 48$	<b>1.001</b>	<i>1.003</i>	1.02	1.042	1.256	1.088	1.14

To check robustness with respect to the country, we use German year-on-year inflation,<sup>9</sup> which is depicted in Figure 5. The results reported in Table 7 are similar again: FI(0.5) outperforms all other methods by a margin increasing in horizon and often being very large.<sup>10</sup> What is different here is


<sup>8</sup>We calculated the inflation from the seasonally adjusted US consumer price index. The exact reference for the series is: U.S. Bureau of Labor Statistics, Consumer Price Index for All Urban Consumers: All Items [CPIAUCSL], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CPIAUCSL>, January 3, 2018.

<sup>9</sup>Here, the exact reference is: Organization for Economic Cooperation and Development, Consumer Price Index: OECD Groups: All Items for Germany [CPALTT01DEM659N], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CPALTT01DEM659N>, February 8, 2018.

<sup>10</sup>This does of course not hold true for the arithmetic mean, which in relative terms gets better as the informational content of the conditioning information used by the other



(a) timeplot



(b) autocorrelogram

**Figure 4.** plots for annualized month-on-month US CPI inflation from January 1948 to November 2017

that the LAR method seems to perform better than the  $\text{FI}(T^\alpha)$  methods. Such a behavior could explain the mixed evidence reported in the literature concerning the usefulness of the FI model for forecasting: Due to the high uncertainty in estimating  $d$ , here the LAR method seems to be superior to the long memory methods if we only look at the ones with estimated  $d$ . But by using the fixed- $d$  method, we can resolve this issue as we clearly see that modeling the long memory explicitly without having to estimate its strength leads to improved forecasting results.

**Table 7.** MSEs relative to  $\text{FI}(0.5)$  from POOS forecasting experiments for monthly year-on-year German CPI inflation, January 1960 to November 2017, rolling window of size 180


horizon	$\text{FI}(T^{0.5})$	$\text{FI}(T^{0.65})$	$\text{FI}(T^{0.8})$	$\text{FI}(1)$	LAR	ES	Mean
$h = 1$	1.01	<b>1.003</b>	1.011	1.014	<i>1.005</i>	1.161	28.975
$h = 3$	1.041	1.032	<i>1.028</i>	1.039	<b>1.012</b>	1.103	8.808
$h = 6$	1.076	1.061	<i>1.059</i>	1.069	<b>1.033</b>	1.093	4.158
$h = 12$	1.168	1.148	1.126	1.137	<i>1.07</i>	<b>1.065</b>	1.796
$h = 24$	1.364	1.283	<i>1.256</i>	1.262	<b>1.214</b>	1.297	1.272
$h = 48$	1.849	1.588	1.527	1.495	<i>1.343</i>	1.522	<b>0.979</b>

## 6.2 Realized Volatility


Another area where many time series show long memory characteristics and the FI model is widely used is finance. A prime example of long memory time series in this field are return volatility series. Since the rise of high frequency data in finance and the introduction of realized volatility measures (see e.g. Andersen and Bollerslev (1998) or Andersen et al. (2001)), the latter are widely used in financial volatility forecasting, being more accurate measures of volatility and better proxies of unobserved conditional variances than squared returns (see e.g. Hansen and Lunde (2011)). FI models are natural candidates for forecasting realized volatility and since Andersen et al.

---

methods gets smaller and smaller.



(a) timeplot



(b) autocorrelogram

**Figure 5.** plots for monthly year-on-year German CPI inflation from January 1960 to November 2017

(2003) have often been applied for that purpose and compared to the performance of competing methods (see again Hansen and Lunde (2011) for a literature overview). A model that is also widely employed in forecasting realized volatility is the HAR model by Corsi (2009). The idea behind the model is that different volatility components arise over different time horizons (trading days, weeks and months) due to different types of market participants and their trading behavior. Being just a restricted AR model, it is very simple and shows a good forecasting performance and has thus become a workhorse model for volatility forecasting. Consequently, we add it as a very important competitor to the forecasting experiments for realized volatility.

A standard way to measure realized volatility is to use the five-minute sub-sampled realized variance or its square root respectively (for details see e.g. Shephard and Sheppard (2010)). To reduce skewness and kurtosis, usually logs of this measure are taken. We follow these conventions and use the logs of the square root of the five-minute sub-sampled realized variance of the S&P 500 from January 2000 to September 2018 and execute a pseudo-out-of-sample forecasting experiment with a rolling window of size 500 on it.<sup>11</sup> The series and its autocorrelogram are depicted in Figure 6. The extremely slowly decaying autocorrelation function clearly indicates long memory. Although, apart from that, its characteristics are quite different from the inflation series, the results shown in Table 8 are very similar: The short memory methods are uniformly outperformed by FI(0.5) and the gap tends to get bigger with the forecasting horizon. Now, the forecasts based on the HAR model are the best forecasts based on a short memory model, showing a decent performance, especially for short horizons. With regards to the long memory methods, here the procedures with estimated  $d$  show slight improvements of up to two percent over FI(0.5). This difference compared to the inflation experiments may be due to the now less complicated dynamics of the underlying short memory process, which makes the estimation of  $d$  easier. More light will be

---

<sup>11</sup>This time series as well as the other realized variance series stem from the "Oxford-Man Institute's realized library, version 0.3", Oxford-Man Institute, University of Oxford by Heber G., A. Lunde, N. Shephard and K. Sheppard (2009).



shedded on this question through the simulations in the next section.

**Table 8.** MSEs relative to FI(0.5) from POOS forecasting experiments for daily log realized volatility of the S&P 500, from 3.01.2000 to 20.09.2018, rolling window of size 500


horizon	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean	HAR
$h = 1$	1.001	<i>0.998</i>	<b>0.995</b>	1.026	1.017	1.018	3.043	1.004
$h = 3$	1	<b>0.997</b>	<b>0.997</b>	1.046	1.038	1.056	2.32	1.011
$h = 5$	0.998	<b>0.995</b>	<i>0.997</i>	1.063	1.054	1.096	2.028	1.009
$h = 10$	1.002	<b>0.994</b>	<i>0.998</i>	1.078	1.087	1.13	1.756	1.019
$h = 20$	0.992	<b>0.984</b>	<i>0.991</i>	1.106	1.118	1.147	1.505	1.019
$h = 40$	<i>0.982</i>	<b>0.98</b>	0.993	1.136	1.124	1.182	1.334	1.032
$h = 80$	<b>0.98</b>	<i>0.981</i>	0.987	1.181	1.114	1.227	1.209	1.055

To check robustness with respect to the time period and especially to the financial crisis, we also executed the analysis on the latter part of the series starting in 2012. As the results in Table 9 show, no large changes occur. Now, there are virtually no differences between the FI ( $T^\alpha$ ) methods and FI(0.5) and HAR performs a little worse than before.


**Table 9.** MSEs relative to FI(0.5) from POOS forecasting experiments for daily log realized volatility of the S&P 500, from 2.01.2012 to 20.09.2018, rolling window of size 500

horizon	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean	HAR
$h = 1$	<i>1.003</i>	1.005	<b>1</b>	1.026	1.019	1.053	2.884	1.015
$h = 3$	<i>1.001</i>	1.004	<b>0.999</b>	1.061	1.043	1.13	1.924	1.031
$h = 5$	<b>1</b>	1.004	<i>1.002</i>	1.087	1.064	1.215	1.618	1.022
$h = 10$	1.004	<b>0.999</b>	<b>0.999</b>	1.109	1.102	1.277	1.439	1.032
$h = 20$	<i>0.994</i>	0.995	<b>0.99</b>	1.177	1.118	1.297	1.286	1.032
$h = 40$	1.003	<i>1.002</i>	<b>0.999</b>	1.23	1.105	1.377	1.198	1.049
$h = 80$	1.037	<i>1.015</i>	<b>0.991</b>	1.391	1.107	1.545	1.118	1.133

We also changed the country and used the respective series for the DAX, again from January 2000 to September 2018. It is plotted in Figure 7 together with its autocorrelation function, which looks very similar to the previous



(a) timeplot



(b) autocorrelogram

**Figure 6.** plots for daily S&P 500 realized volatility from 3.01.2000 to 20.09.2018


one. The results as shown in Table 10 do not change substantially, too, with the only discrepancy that now for the smaller bandwidths and for the larger horizons some performance gains of the methods with estimated  $d$  over the fixed- $d$  method occur.

**Table 10.** MSEs relative to FI(0.5) from POOS forecasting experiments for daily log realized volatility of the DAX, from 3.01.2000 to 20.09.2018, rolling window of size 500

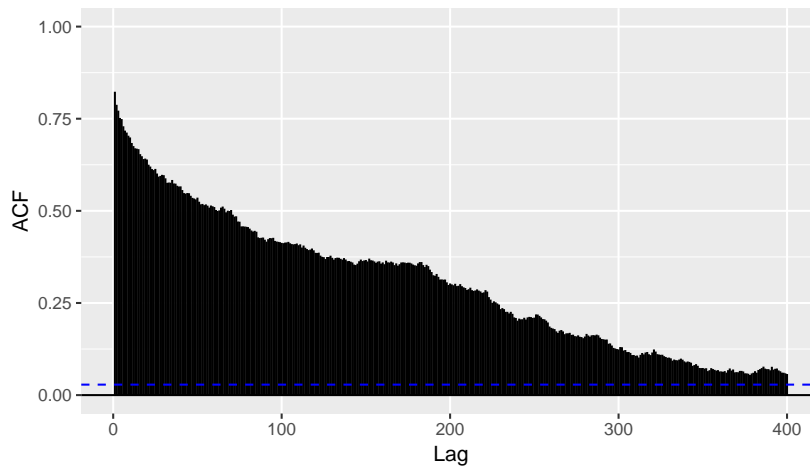
horizon	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean	HAR
$h = 1$	<i>0.997</i>	<b>0.996</b>	1.001	1.019	1.014	1.011	3.405	<i>0.997</i>
$h = 3$	<i>0.992</i>	<b>0.989</b>	0.998	1.021	1.022	1.022	2.693	0.996
$h = 5$	<i>0.991</i>	<b>0.987</b>	0.995	1.039	1.04	1.044	2.423	0.999
$h = 10$	<i>0.989</i>	<b>0.982</b>	0.996	1.047	1.063	1.065	2.058	1.014
$h = 20$	<i>0.98</i>	<b>0.965</b>	0.993	1.05	1.096	1.073	1.75	1.022
$h = 40$	<i>0.96</i>	<b>0.946</b>	0.993	1.043	1.135	1.06	1.519	1.045
$h = 80$	<i>0.939</i>	<b>0.937</b>	0.993	1.057	1.135	1.078	1.325	1.058

### 6.3 Experiments on the optimal estimation window size

An important and often overlooked specification issue when setting up a forecasting method is the size of the estimation window, see e.g. Pesaran and Timmermann (2007). Should all available information on the past of the series of interest be used to decrease the variance of the estimators or should only the more recent past be included to robustify the procedure against structural change? Certainly, the answer to this question depends on the characteristics of the individual variables. But as it is often implicitly assumed that a longer stretch of past data leads to better results, we performed experiments on the optimal window size on the inflation and realized volatility series used already in this paper to check this assumption. At the same time, the experiments in this section serve as a robustness check for the choice of the window size in the first two parts of this section. Moreover, the effects of using an expanding instead of a rolling estimation window are



(a) timeplot



(b) autocorrelogram

**Figure 7.** plots for daily DAX realized volatility from 3.01.2000 to 20.09.2018

analyzed.

We again employed pseudo-out-of-sample forecasting for different horizons and now also for different sizes of the rolling estimation window and for an expanding window. This enables us to compare the accuracy of forecasts obtained by using rolling windows of different sizes and an expanding window. To ensure a better comparability, we let the methods all produce forecasts for the same time period, which means that for the shorter window sizes the remaining first parts of the series, which are needed for the estimation with the longer sizes, are not used.

**Table 11.** MSEs from POOS experiments for different rolling window sizes and for an expanding window with initial size 360 for monthly year-on-year US CPI inflation, January 1948 to November 2017, the first forecast is for January 1988 for all sizes, forecast horizon  $h = 1$

window size	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI(0.5)	FI (1)	LAR	ES	Mean
60	<b>0.1322</b>	<i>0.1333</i>	0.1338	0.1337	0.1343	0.1536	0.1515	4.3899
120	0.1051	0.1032	<b>0.1019</b>	0.1039	<i>0.1031</i>	0.1072	0.1505	5.5664
180	0.1036	0.1044	<i>0.1033</i>	<b>0.1028</b>	0.1055	0.1049	0.1527	6.815
240	<b>0.1029</b>	0.1035	0.1047	0.1034	0.105	<i>0.1032</i>	0.1558	8.2456
360	0.106	0.1048	<i>0.1047</i>	<b>0.104</b>	<i>0.1047</i>	0.1053	0.1481	9.4418
<i>expanding</i>	<i>0.1055</i>	0.1067	0.1066	<b>0.105</b>	0.1064	0.1061	0.1445	8.0776

Tables 11 and 12 contain the results for the US inflation series from the beginning of this section for the forecasting horizons  $h = 1$  and  $h = 6$  respectively and for rolling windows of sizes varying from 60 to 360 and an expanding window of initial size 360. For the other inflation series and other horizons the results are very similar and so we do not report them here. We now report absolute MSFEs to enable a comparison over the rows and hence a column for the FI(0.5) method is added. There are basically two ways of reading the tables: Firstly, the focus of interest can be the best forecasting performance over all methods and window sizes. Then only the results of the best method for each window size needs to be considered, which amounts to a comparison of the numbers in bold face. Secondly, when interested in the dependence of the performance of a single method on the window size

one should stay in the respective column. Taking the first perspective, the forecasting performance increases strongly from a window size of 60 to a size of 120. From then on, the MSE does not change much anymore, but seems to rise slightly for the largest size of 360 and for the expanding window. Looking at the individual methods, the same pattern emerges for the FI-based methods and the autoregression. Hence, for the monthly inflation series a rolling window with a medium estimation window size between 120 and 240 seems to be a good choice, balancing out an increased estimation uncertainty for too small sizes and the vulnerability to gradual structural change for too large sizes optimally.

**Table 12.** MSEs from POOS experiments for different rolling window sizes and for an expanding window with initial size 360 for monthly year-on-year US CPI inflation, January 1948 to November 2017, the first forecast is for June 1988 for all sizes, forecast horizon  $h = 6$

window size	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	$FI(0.5)$	FI (1)	LAR	ES	Mean
60	<i>1.6622</i>	<b>1.6441</b>	1.7114	1.8034	1.707	3.1693	1.9453	5.1008
120	1.6596	<i>1.6114</i>	<b>1.599</b>	1.6434	1.6472	1.8125	1.8672	6.0981
180	<i>1.6103</i>	1.6421	1.6676	<b>1.6073</b>	1.6701	1.7181	1.8615	7.3252
240	1.653	<b>1.5387</b>	1.5603	<i>1.5417</i>	1.5441	1.5881	1.8457	8.6649
360	1.7027	1.6369	1.6364	<b>1.6018</b>	1.64	<i>1.6064</i>	1.8826	9.697
<i>expanding</i>	<i>1.6351</i>	1.7755	1.7525	<b>1.6172</b>	1.7143	1.665	1.8027	8.2271

For the realized volatility series from above we did the same analysis for rolling window sizes from 125 to 1000 and an expanding window of size 1000. We report the results for forecast horizons  $h = 1$  and  $h = 5$  for the initial series from above in Tables 13 and 14, the results for other horizons and window sizes are very similar as well. Again, picking a too small window size of 125 leads to a drop in performance. From the window size 250 on the MSEs stay almost constant for increasing window sizes of the rolling window up to 1000 and are almost identical to the results when using an expanding window. Thus, for the realized volatility series it does hardly matter which window size is chosen as long as it is not too small.

**Table 13.** MSEs from POOS experiments for different rolling window sizes and for an expanding window with initial size 1000 for daily log realized volatility of the S&P 500, from 3.01.2000 to 20.09.2018, the first forecast is for the 1001st observation for all sizes, forecast horizon  $h = 1$

window size	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI(0.5)	FI (1)	LAR	ES	Mean
125	0.0959	<i>0.0949</i>	<b>0.0945</b>	0.0951	0.1003	0.098	0.0951	0.1969
250	0.0929	<i>0.0921</i>	<b>0.0915</b>	0.0924	0.097	0.094	0.0937	0.2283
500	0.092	<i>0.0916</i>	<b>0.0912</b>	0.0917	0.0944	0.0934	0.0938	0.2807
750	0.0919	<i>0.0917</i>	<b>0.0915</b>	0.0918	0.0941	0.0934	0.0942	0.3063
1000	0.092	<b>0.0917</b>	<b>0.0917</b>	0.0919	0.0933	0.0935	0.0943	0.3266
<i>expanding</i>	0.0919	0.0919	<b>0.0917</b>	<i>0.0918</i>	0.0929	0.0923	0.0947	0.369

**Table 14.** MSEs from POOS experiments for different rolling window sizes and for an expanding window with initial size 1000 for daily log realized volatility of the S&P 500, from 3.01.2000 to 20.09.2018, the first forecast is for the 1005th observation for all sizes, forecast horizon  $h = 5$

window size	HW	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI(0.5)	FI (1)	LAR	ES	Mean
125	0.1504	0.1494	0.1458	<b>0.1453</b>	<b>0.1453</b>	0.1623	0.1562	0.1585	0.2064
250	0.1485	0.1432	0.1418	<i>0.1416</i>	<b>0.1413</b>	0.1599	0.151	0.1571	0.2337
500	0.1488	0.1415	<i>0.141</i>	<b>0.1408</b>	0.1416	0.1513	0.1497	0.1565	0.2847
750	0.1481	0.1425	<b>0.1408</b>	<i>0.1409</i>	0.1415	0.1504	0.1493	0.1564	0.309
1000	0.1472	0.1423	<b>0.1411</b>	<i>0.1412</i>	<i>0.1412</i>	0.1482	0.1485	0.1558	0.3289
<i>expanding</i>	0.1442	0.1414	<i>0.1411</i>	<b>0.1409</b>	0.1413	0.146	0.1438	0.1538	0.37

## 7 Simulation evidence

To analyze if the empirical findings from the previous section still hold when true FI processes are forecasted and if some of these findings can be better explained by systematic variations in the data generating process, we conducted a simulation study. We simulated from the FI process given in equation 3 with different values of the long memory parameter  $d$  and the three different short memory components  $\{x_t\}$  already used in sections 3 and 4. For each instance we generated 1000 series of size  $T + h$ , held back the last  $h$  observations and used the first  $T$  observations to generate the forecasts for period  $T + h$ . We chose  $T \in \{60, 300, 1500\}$  and  $h \in \{1, 3, 12\}$ . As in the previous section we present the MSEs relative to FI(0.5) and the smallest and second

smallest MSEs in each row of the tables are marked in bold face and italics respectively.

Table 15 contains the results for FI(0.4) processes. The FI(0.5) method again almost uniformly outperforms the long autoregression and the exponential smoothing method. This confirms the empirical results and further makes a case for the use of true long memory models when forecasting strongly persistent time series. Additionally, as observed with the empirical POOS experiments, the predictive performance of both LAR and ES gets worse relative to FI(0.5) as the forecasting horizon increases. Unsurprisingly, adequately accounting for long memory seems to get more critical with increasing forecasting horizon. While the dynamics over a short horizon may be approximated by only a short memory component, over longer horizons the long memory component increasingly dominates the forecastable part of the process as the influence of the short memory component dies out much faster. Furthermore, as in the empirical experiments, the long autoregression clearly shows the best performance from the methods that do not employ a long memory model.

When comparing the performance of the different FI-based methods, the ones using the local Whittle estimator with different bandwidth choices to estimate the long memory parameter do not outperform the one using an a priori fixed  $d = 0.5$  (although this is not the true  $d$ ), not even for the largest sample size  $T = 1500$ . As expected, the performance of the FI( $T^\alpha$ ) methods varies with the bandwidth and for more persistent  $\{x_t\}$  larger bandwidths lead to better results. The smaller the sample size and the larger the persistence of  $\{x_t\}$  gets, the worse do the FI( $T^\alpha$ ) methods come off against FI(0.5). This can be explained by the lack of accuracy and the bias respectively that plague the local Whittle estimator under these circumstances. Regarding the influence of the forecasting horizon, exactly under these circumstances, where the estimation of  $d$  is difficult, a similar phenomenon occurs as for LAR and ES: The relative performance compared to FI(0.5) decreases with  $h$ , sometimes leading to very large losses for  $h = 12$ . This behavior is problematic as these conditions often arise in practice: Longer forecasting horizons are



naturally of interest under long memory, small sizes of the estimation sample will rather be the rule than the exception, especially given the analysis at the end of the last section concerning optimal estimation window sizes, and substantive short memory dynamics may well arise. As discussed above, the long memory component becomes more and more important the longer the forecasting horizon gets. Thus, as the difficulty of disentangling long and short memory dynamics and estimating the long memory component often leads to a very wrong specification of the long memory component, the predictive performance of the  $FI(T^\alpha)$  methods suffers. Our experiments suggest that it works better to accept a wrong value of  $d$  a priori and choose it ad hoc so that it lies at least not very far away from the dynamics usually observed in economic time series. Together with a short memory component which may somehow adapt to this, a good and robust forecasting performance can be achieved.

Table 16 presents the results for simulated nonstationary long memory processes, namely  $FI(0.7)$  processes. As expected, as the true value is further away from 0.5 now, the other methods gain a little ground relative to  $FI(0.5)$ . But still no large gains are possible by estimating  $d$ , the improvements under the best circumstances and the best bandwidth choice for  $FI(T^\alpha)$  never reach 5 percent. Otherwise, no substantive changes occur compared to Table 15. Exponential smoothing performs better than above, but is still not as good as the long autoregression.

The results from the simulation study are perfectly in line with the ones from the ones from the empirical section and additionally help to better understand them: For the inflation series we observed that the methods with estimated  $d$  showed a worse performance than the method with the fixed  $d = 0.5$  and the performance dropped over the forecasting horizon, usually with very large losses for large horizons. As we observed an analogous behavior for series with more persistent short memory components, the inflation series seem to fit in this category. For realized volatility, the two methods show a very similar performance, with perhaps slight advantages for the methods with estimated  $d$ . Thus, the realized volatility series series do not seem to

have very persistent short memory components. Furthermore, the estimation window sizes used for the realized volatility series are longer due to the higher frequency of the financial series.

Overall, we conclude from the forecasting experiments with empirical and simulated series that explicitly accounting for long memory by using FI models leads to improved results. When using a FI model, fixing the parameter to  $d = 0.5$  a priori offers a generally well-performing and simple method, which is hard to beat and could serve as a benchmark for forecasting under long memory. In some situations, slight improvements over this method are possible by using a semiparametric estimator for  $d$ , having to accept a higher downside risk.

## 8 Concluding remarks

In this paper we are concerned with the issue of forecasting strongly persistent time series. By pseudo-out-of-sample forecasting experiments on inflation and realized volatility time series and by simulations we find clear evidence that accounting explicitly for long memory by using a fractionally integrated model improves predictive performance in comparison to classical short memory competitors. We are able to explain the former mixed evidence on this issue and to present a simple and successful method for forecasting by proposing the use of a fixed-parameter FI model with  $d = 0.5$ , which overcomes the problem of the high uncertainty associated with the estimation of the long memory parameter and is nevertheless able to capture decisive long memory features.

While our results hence suggest that FI models are a useful tool in forecasting strongly persistent time series, which are widespread in economics, finance and many other fields, there seems not much too be gained by estimating  $d$ , i.e. by trying to determine the exact strength of the long memory. Further support for this is provided by our analysis of the optimal estimation window size, where we find that medium window sizes provide an optimal balance between high estimation variances and gradual structural change.

To construct proper forecasting methods from FI models, the estimation of the mean and of the long memory parameter (at least if one wants to use models with an estimated parameter or wants to study the effects of the estimation) are crucial. We investigate both and our results are of general interest outside the realm of forecasting as well. We compare the arithmetic mean and Shimotsu's and Robinson's mean estimators by a theoretical analysis, providing asymptotic theory for the latter for the first time, and by a simulation study to find that the Robinson estimator can yield efficiency gains and shows a good performance in practice. Regarding the estimation of  $d$ , we compare the Whittle estimator plus AIC, the global semiparametric Whittle and the local Whittle estimator to examine if parametric estimators plus model selection step, global or local semiparametric estimators yield more accurate results. Previous literature has already identified the members of the Whittle family as favourable. As the parametric and the global semiparametric estimator incur drastic losses compared to the local Whittle in our simulations, which we can trace back to the very bad performance of the Whittle estimator under even small amounts of misspecification, we conclude that they should be used with care and that local semiparametric estimators should usually be a better choice.

Not to go beyond the scope of this paper, we limit our analysis to a univariate framework, but we believe that many of our fundamental findings will continue to hold and be helpful in multivariate settings. Furthermore, for the fixed- $d$  FI processes we picked the value  $d = 0.5$  at the beginning of our analysis and do not report results for alternative choices, withstanding the statistician's immanent desire to optimize parameter values (after just having overcome the estimation problem). Nevertheless, it is certainly possible to go in this direction and to develop methods for the selection of the value of  $d$  based on predictive performance. Baillie et al. (2014) for example do a similar thing and pick the bandwidth for semiparametric estimation of  $d$  by cross-validation. In both of these fields there is certainly room for future research.

## Appendix

PROOF OF PROPOSITION 1 The variance is given by

$$\text{Var}(\widehat{\mu}(d)) = \text{Var} \left( \sum_{t=1}^T r_t x_t \right) \left( \sum_{t=1}^T r_t^2 \right)^{-2},$$

where  $r_t = \sum_{j=0}^{t-1} \pi_j(d)$ . In a first step, we deal with the denominator of this variance, in a second step we will turn to the numerator.

Step 1) From Hassler (2019, Lemma 5.1) we have  $r_t = \pi_{t-1}(d-1)$ , which implies  $r_t \sim t^{-d}/\Gamma(1-d)$ . From Giraitis et al. (2012, Prop. 2.3.1) we have  $\sum_{t=1}^T t^{-2d} \sim T^{1-2d}/(1-2d)$ . Now, define  $s_T = \sum_{t=1}^T r_t^2$  and  $\sigma_T := \sum_{t=1}^T t^{-2d}$ , where the latter constitutes a strictly monotone and divergent sequence. Note that  $(s_T - s_{T-1})/(\sigma_T - \sigma_{T-1}) \rightarrow 1/\Gamma^2(1-d)$ . Hence, we can use the so-called Stolz-Cesàro Theorem, see Knopp (1990, pp. 76, 77), and conclude that  $s_T \sim \sigma_T/\Gamma^2(1-d)$ . This provides

$$\sum_{t=1}^T r_t^2 \sim \frac{T^{1-2d}}{(1-2d)\Gamma^2(1-d)}.$$

Step 2) The numerator becomes

$$\text{Var} \left( \sum_{t=1}^T r_t x_t \right) = \gamma_x(0) \sum_{t=1}^T r_t^2 + 2 \sum_{h=1}^{T-1} \gamma_x(h) \sum_{t=1}^{T-h} r_t r_{t+h}.$$

With  $r_{t+h} = r_t + \sum_{j=t}^{t+h-1} \pi_j(d)$  we obtain

$$\frac{\sum_{t=1}^{T-h} r_t r_{t+h}}{\sum_{t=1}^T r_t^2} = 1 - \frac{\sum_{t=T-h+1}^T r_t^2 - \sum_{t=1}^{T-h} r_t \sum_{j=t}^{t+h-1} \pi_j(d)}{\sum_{t=1}^T r_t^2} = 1 - \frac{b_h}{b_T},$$

where  $b_h$  and  $b_T$  are defined implicitly. Note that  $\pi_0(d) = 1$  and  $\pi_j(d)$  is negative and monotonically increasing for  $j > 0$ . Consequently, the sequence  $\{b_h\}$  is positive and monotonically increasing, while  $b_T$  diverges according to the previous step. Hence,

$$\sum_{h=1}^{T-1} \frac{b_h \gamma_x(h)}{b_T} \rightarrow 0 \quad \text{as } T \rightarrow \infty,$$

by what is sometimes called Kronecker's Lemma. It follows that

$$\frac{\text{Var}\left(\sum_{t=1}^T r_t x_t\right)}{\sum_{t=1}^T r_t^2} \rightarrow \gamma_x(0) + 2 \sum_{h=1}^{\infty} \gamma_x(h) = \omega_x^2.$$

Hence, the proof is complete.

## References

- Abadir, K. M., W. Distaso, and L. Giraitis (2007). Nonstationarity-extended local Whittle estimation. *Journal of Econometrics* 141, 1353–1384.
- Andersen, T. G. and T. Bollerslev (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and H. Ebens (2001). The distribution of realized stock return volatility. *Journal of Financial Economics* 61, 43–76.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. *Econometrica* 71, 579–625.
- Andrews, D. W. K. and P. Guggenberger (2003). A bias-reduced log-periodogram regression estimator for the long-memory parameter. *Econometrica* 71, 675–712.
- Andrews, D. W. K. and Y. Sun (2004). Adaptive local polynomial Whittle estimation of long-range dependence. *Econometrica* 72, 569–614.
- Arteche, J. (2012). Parametric vs. semiparametric long memory: Comments on prediction from ARFIMA models: Comparison between MLE and semiparametric estimation. *International Journal of Forecasting* 1(28), 54–56.
- Baillie, R. T., C.-F. Chung, and M. A. Tieslau (1996). Analysing inflation by the fractionally integrated ARFIMA-GARCH model. *Journal of Applied Econometrics* 11, 23–40.

- Baillie, R. T., G. Kapetanios, and F. Papailias (2014). Bandwidth selection by cross-validation for forecasting long memory financial time series. *Journal of Empirical Finance* 29, 129–143.
- Baillie, R. T., C. Kongcharoen, and G. Kapetanios (2012). Prediction from ARFIMA models: Comparisons between MLE and semiparametric estimation procedures. *International Journal of Forecasting* 28, 46–53.
- Beran, J. (1995). Likelihood estimation of the differencing parameter for invertible short and long memory autoregressive integrated moving average models. *Journal of the Royal Statistical Society. Series B* 57, 659–672.
- Beran, J., R. J. Bhansali, and D. Ocker (1998). On unified model selection for stationary and nonstationary short- and long-memory autoregressive processes. *Biometrika* 85, 921–934.
- Beran, J., Y. Feng, S. Ghosh, and R. Kulik (2013). *Long-Memory Processes: Probabilistic Properties and Statistical Methods*. Springer.
- Bhansali, R. J., L. Giraitis, and P. S. Kokoszka (2006). Estimation of the memory parameter by fitting fractionally differenced autoregressive models. *Journal of Multivariate Analysis* 97, 2101–2130.
- Bhardwaj, G. and N. R. Swanson (2006). An empirical investigation of the usefulness of ARFIMA models for predicting macroeconomic and financial time series. *Journal of Econometrics* 131, 539–578.
- Brodsky, J. and C. M. Hurvich (1999). Multi-step forecasting for long-memory processes. *Journal of Forecasting* 18, 59–75.
- Clements, M. P. and D. F. Hendry (2006). Forecasting with breaks. In M. P. Clements and D. F. Hendry (Eds.), *Handbook of economic forecasting*, pp. 605–657. Elsevier.
- Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics* 7(2), 174–196.

- Crato, N. and B. K. Ray (1996). Model selection and forecasting for long-range dependent processes. *Journal of Forecasting* 15, 107–125.
- Dahlhaus, R. (1989). Efficient parameter estimation for self-similar processes. *The Annals of Statistics* 17, 1749–1766.
- Dahlhaus, R. (2006). Correction: Efficient parameter estimation for self-similar processes. *The Annals of Statistics* 34, 1045–1047.
- De Gooijer, J. G. and R. J. Hyndman (2006). 25 years of time series forecasting. *International journal of forecasting* 22(3), 443–473.
- Demetrescu, M. and U. Hassler (2016). (When) Do long autoregressions account for neglected changes in parameters? *Econometric Theory* 32, 1317–1348.
- Fox, R. and M. S. Taqqu (1986). Large-sample properties of parameter estimates for strongly dependent stationary Gaussian time series. *The Annals of Statistics* 14, 517–532.
- Geweke, J. and S. Porter-Hudak (1983). The estimation and application of long memory time series models. *Journal of Time Series Analysis* 4, 221–238.
- Giraitis, L., H. L. Koul, and D. Surgailis (2012). *Large Sample Inference for Long Memory Processes*. Imperial College Press.
- Granger, C. W. J. and R. Joyeux (1980). An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1, 15–29.
- Hansen, P. R. and A. Lunde (2011). Forecasting volatility using high frequency data. In M. P. Clements and D. F. Hendry (Eds.), *The Oxford handbook of economic forecasting*, Chapter 19, pp. 525–558. Oxford University Press.
- Hassler, U. (2019). *Time Series Analysis with Long Memory in View*. Wiley.

- Hassler, U. and J. Wolters (1995). Long memory in inflation rates: International evidence. *Journal of Business & Economic Statistics* 13, 37–45.
- Henry, M. (2001). Robust automatic bandwidth for long memory. *Journal of Time Series Analysis* 22, 293–316.
- Hualde, J. and P. M. Robinson (2011). Gaussian pseudo-maximum likelihood estimation of fractional time series models. *The Annals of Statistics* 39, 3152–3181.
- Hurvich, C. M. and J. Brodsky (2001). Broadband semiparametric estimation of the memory parameter of a long-memory time series using fractional exponential models. *Journal of Time Series Analysis* 22, 221–249.
- Hurvich, C. M., R. Deo, and J. Brodsky (1998). The mean squared error of Geweke and Porter-Hudak’s estimator of the memory parameter of a long-memory time series. *Journal of Time Series Analysis* 19, 19–46.
- Hurvich, C. M. and R. S. Deo (1999). Plug-in selection of the number of frequencies in regression estimates of the memory parameter of a long-memory time series. *Journal of Time Series Analysis* 20, 331–341.
- Hyndman, R. J., A. B. Koehler, R. D. Snyder, and S. Grose (2002). A state space framework for automatic forecasting using exponential smoothing methods. *International Journal of forecasting* 18(3), 439–454.
- Johansen, S. (2008). A representation theory for a class of vector autoregressive models for fractional processes. *Econometric Theory* 24, 651–676.
- Knopp, K. (1990). *Theory and Application of Infinite Series* (Republication of 1951 Blackie & Son ed.). Dover Publications.
- Künsch, H. (1987). Statistical aspects of self-similar processes. In Y. Prohorov and V. V. Sazanov (Eds.), *Proceedings of the First World Congress of the Bernoulli Society, Vol. 1*, pp. 67–74. VNU Science Press.



- Marinucci, D. and P. M. Robinson (1999). Alternative forms of fractional Brownian motion. *Journal of Statistical Planning and Inference* 80, 111–122.
- Marinucci, D. and P. M. Robinson (2000). Weak convergence of multivariate fractional processes. *Stochastic Processes and their Applications* 86, 103–120.
- Moulines, E. and P. Soulier (1999). Broadband log-periodogram regression of time series with long-range dependence. *The Annals of Statistics* 27, 1415–1439.
- Moulines, E. and P. Soulier (2000). Data driven order selection for projection estimator of the spectral density of time series with long range dependence. *Journal of Time Series Analysis* 21, 193–218.
- Nelson, C. R. and C. I. Plosser (1982). Trends and random walks in marcoeconomic time series: Some evidence and implications. *Journal of Monetary Economics* 10, 139–162.
- Nielsen, M. Ø. (2015). Asymptotics for the conditional-sum-of-squares estimator in fractional time-series models. *Journal of Time Series Analysis* 36, 154–188.
- Nielsen, M. Ø. and P. H. Frederiksen (2005). Finite sample comparison of parametric, semiparametric, and wavelet estimators of fractional integration. *Econometric Reviews* 24, 405–443.
- Pesaran, M. H., A. Pick, and M. Pranovich (2013). Optimal forecasts in the presence of structural breaks. *Journal of Econometrics* 177(2), 134–152.
- Pesaran, M. H. and A. Timmermann (2007). Selection of estimation window in the presence of breaks. *Journal of Econometrics* 137(1), 134–161.
- Ray, B. K. (1993). Modeling long-memory processes for optimal long-range prediction. *Journal of Time Series Analysis* 14, 511–525.

- Robinson, P. M. (1994). Efficient tests of nonstationary hypotheses. *Journal of the American Statistical Association* 89, 1420–1437.
- Robinson, P. M. (1995). Gaussian semiparametric estimation of long range dependence. *Annals of Statistics* 23, 1630–1661.
- Samarov, A. and M. S. Taqqu (1988). On the efficiency of the sample mean in long-memory noise. *Journal of Time Series Analysis* 9, 191–200.
- Shephard, N. and K. Sheppard (2010). Realising the future: Forecasting with high-frequency-based volatility (HEAVY) models. *Journal of Applied Econometrics* 25(2), 197–231.
- Shimotsu, K. (2010). Exact local Whittle estimation of fractional integration with unknown mean and trend. *Econometric Theory* 26, 501–540.
- Shimotsu, K. and P. C. B. Phillips (2005). Exact local Whittle estimation of fractional integration. *The Annals of Statistics* 33, 1890–1933.
- Smith, J. and S. Yadav (1994). Forecasting costs incurred from unit differencing fractionally integrated processes. *International Journal of Forecasting* 10, 507–514.
- Stock, J. H. and M. W. Watson (1998). A comparison of linear and nonlinear univariate models for forecasting macroeconomic time series. Technical report, National Bureau of Economic Research.
- Tanaka, K. (1999). The nonstationary fractional unit root. *Econometric Theory* 15, 549–582.
- Velasco, C. and P. M. Robinson (2000). Whittle pseudo-maximum likelihood estimation for nonstationary time series. *Journal of the American Statistical Association* 95, 1229–1243.

**Table 15.** MSEs relative to FI(0.5) for different forecasting methods for simulated series of size  $T$  and a forecast horizon  $h$ , the series are generated by FI(0.4) processes with short memory component  $x_t$ , 1000 repetitions

$x_t$	$T$	$h$	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean
iid	60	1	1.011	<i>0.974</i>	<b>0.964</b>	1.117	1.049	<i>0.974</i>	1.243
		3	1.043	<i>1.019</i>	<b>0.995</b>	1.236	1.072	1.123	1.041
		12	1.058	1.014	<i>0.997</i>	1.281	1.093	1.221	<b>0.985</b>
	300	1	1.006	<i>0.999</i>	<b>0.993</b>	1.036	1.028	1.002	1.567
		3	1.042	<i>1.017</i>	<b>1.008</b>	1.118	1.065	1.099	1.329
		12	1.006	<i>0.991</i>	<b>0.987</b>	1.197	1.049	1.24	1.093
	1500	1	1.002	<i>0.999</i>	<b>0.995</b>	1.01	1.013	1.069	1.754
		3	1.004	<i>0.998</i>	<b>0.995</b>	1.061	1.014	1.079	1.376
		12	<i>0.988</i>	0.99	<b>0.986</b>	1.121	1.022	1.219	1.151
AR(1)	60	1	<i>1.001</i>	1.017	1.018	1.056	1.023	<b>0.98</b>	3.096
		3	<b>1.025</b>	1.049	1.053	1.186	<i>1.037</i>	1.147	1.285
		12	1.097	1.079	1.098	1.326	<i>1.069</i>	1.327	<b>1.027</b>
	300	1	<b>1.002</b>	<i>1.003</i>	1.012	1.023	1.015	1.037	3.897
		3	<i>1.011</i>	<b>1.007</b>	1.036	1.105	1.03	1.217	1.939
		12	<i>1.019</i>	<b>1.017</b>	1.075	1.278	1.051	1.434	1.151
	1500	1	<i>1.002</i>	<b>0.999</b>	<i>1.002</i>	1.028	1.013	1.071	5.345
		3	<b>0.99</b>	<b>0.99</b>	1.016	1.063	1.017	1.223	1.898
		12	<b>0.993</b>	<i>0.996</i>	1.035	1.173	1.035	1.466	1.298
MA(9)	60	1	1.052	1.053	1.054	<b>1.026</b>	<i>1.043</i>	1.064	16.356
		3	<i>1.099</i>	1.103	1.107	<b>1.051</b>	1.11	1.137	3.369
		12	1.511	1.579	1.383	<i>1.239</i>	1.359	1.638	<b>1.036</b>
	300	1	<b>1.002</b>	1.013	1.032	1.018	<i>1.011</i>	1.039	22.784
		3	<i>1.012</i>	1.051	1.089	1.045	<b>1.01</b>	1.121	4.666
		12	<b>1.027</b>	1.235	1.485	1.282	<i>1.033</i>	1.744	1.327
	1500	1	0.998	<i>0.995</i>	1.01	1.001	<b>0.989</b>	1.031	27.057
		3	<i>1.001</i>	1.002	1.042	1.025	<b>0.999</b>	1.171	5.704
		12	<b>1.002</b>	<i>1.018</i>	1.231	1.164	1.024	1.755	1.433

**Table 16.** MSEs relative to FI(0.5) for different forecasting methods for simulated series of size  $T$  and a forecast horizon  $h$ , the series are generated by FI(0.7) processes with short memory component  $x_t$ , 1000 repetitions

$x_t$	$T$	$h$	FI ( $T^{0.5}$ )	FI ( $T^{0.65}$ )	FI ( $T^{0.8}$ )	FI (1)	LAR	ES	Mean
iid	60	1	1.008	0.982	<i>0.974</i>	1.06	1.058	<b>0.943</b>	3.488
		3	1.008	<i>0.992</i>	<b>0.965</b>	1.06	1.04	1.024	1.787
		12	1.062	<i>0.981</i>	<b>0.959</b>	1.05	1.208	1.038	1.165
	300	1	0.989	0.989	<i>0.975</i>	1.001	1.005	<b>0.964</b>	6.663
		3	1.001	<i>0.976</i>	<b>0.968</b>	1.024	1.035	1.022	4.165
		12	0.994	<i>0.969</i>	<b>0.954</b>	1.08	1.095	1.124	2.121
	1500	1	0.992	<i>0.989</i>	<b>0.987</b>	1.002	0.999	1.031	14.521
		3	0.985	<i>0.98</i>	<b>0.979</b>	1.01	1.005	1.042	7.857
		12	0.999	<i>0.997</i>	<b>0.99</b>	1.073	1.053	1.133	3.866
AR(1)	60	1	<i>1.016</i>	1.043	1.03	1.038	1.096	<b>0.997</b>	11.316
		3	1.049	1.04	<i>1.036</i>	1.04	1.106	<b>0.986</b>	3.062
		12	1.131	<b>1.078</b>	<i>1.085</i>	1.089	1.193	1.162	1.331
	300	1	<i>0.988</i>	0.989	0.994	0.998	0.994	<b>0.979</b>	25.032
		3	<b>0.986</b>	<i>0.988</i>	1.015	1.021	1.012	1.063	7.634
		12	<i>0.979</i>	<b>0.973</b>	1.023	1.046	1.048	1.09	2.217
	1500	1	<i>0.993</i>	<b>0.99</b>	<i>0.993</i>	0.998	1	1.019	58.401
		3	<i>0.989</i>	<b>0.987</b>	1.012	1.022	1.009	1.082	13.297
		12	<b>0.977</b>	<b>0.977</b>	1.03	1.063	1.026	1.145	4.757
MA(9)	60	1	1.046	1.049	1.041	<b>1.022</b>	1.058	<i>1.034</i>	76.782
		3	1.051	1.038	<i>1.033</i>	<b>0.983</b>	1.102	1.047	10.086
		12	1.327	1.313	<i>1.116</i>	<b>1.019</b>	1.402	1.558	1.405
	300	1	<b>0.999</b>	1.011	1.017	1.014	1.027	<i>1.001</i>	186.109
		3	<b>1.008</b>	1.034	1.047	<i>1.01</i>	1.017	1.019	24.082
		12	<b>1.027</b>	1.179	1.274	1.071	<i>1.06</i>	1.244	3.309
	1500	1	1.006	1.003	1.01	<i>1.002</i>	<b>0.998</b>	1.005	391.329
		3	<b>0.998</b>	1	1.024	1.003	<i>0.999</i>	1.046	47.794
		12	<b>0.987</b>	1.002	1.17	1.022	<i>0.993</i>	1.226	5.836