

A Service of

ZBW

Leibniz-Informationszentrum Wirtschaft Leibniz Information Centre for Economics

Manganelli, Simone

Working Paper A new theory of forecasting

ECB Working Paper, No. 584

Provided in Cooperation with: European Central Bank (ECB)

Suggested Citation: Manganelli, Simone (2006) : A new theory of forecasting, ECB Working Paper, No. 584, European Central Bank (ECB), Frankfurt a. M.

This Version is available at: https://hdl.handle.net/10419/153018

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



WWW.ECONSTOR.EU



WORKING PAPER SERIES NO. 584 / JANUARY 2006

A NEW THEORY OF FORECASTING

by Simone Manganelli





NO. 584 / JANUARY 2006

A NEW THEORY OF FORECASTING'

by Simone Manganelli²



In 2006 all ECB publications will feature a motif taken from the €5 banknote.





1 I would like to thank, without implicating them, Lorenzo Cappiello, Rob Engle, Lutz Kilian, Benoit Mojon, Cyril Monnet, Andrew Patton and Timo Teräsvirta for their encouragement and useful suggestions. I also would like to thank the ECB Working Paper Series editorial board and seminar participants at the ECB. The views expressed in this paper are those of the author and do not necessarily reflect those of the European Central Bank or the Eurosystem.

2 DG-Research, European Central Bank, Kaiserstrasse 29, D-60311 Frankfurt am Main, Germany; e-mail: simone.manganelli@ecb.int

This paper can be downloaded without charge from http://www.ecb.int or from the Social Science Research Network electronic library at http://ssrn.com/abstract_id=876949.

© European Central Bank, 2006

Address Kaiserstrasse 29 60311 Frankfurt am Main, Germany

Postfach 16 03 19 60066 Frankfurt am Main, Germany

Telephone +49 69 1344 0

Internet http://www.ecb.int

Fax +49 69 1344 6000

Telex 411 144 ecb d

All rights reserved.

Any reproduction, publication and reprint in the form of a different publication, whether printed or produced electronically, in whole or in part, is permitted only with the explicit written authorisation of the ECB or the author(s).

The views expressed in this paper do not necessarily reflect those of the European Central Bank.

The statement of purpose for the ECB Working Paper Series is available from the ECB website, http://www.ecb.int.

ISSN 1561-0810 (print) ISSN 1725-2806 (online)

CONTENTS

Abstract				
Non-technical summary				
1.	Introduction			
2.	The problem			
3.	An alternative theory of forecasting			
4.	Generalisation			
5.	Examples			
	5.1	Simulation	20	
		5.1.1 Discussion	20	
	5.2	Forecasting U.S. GDP	22	
	5.3	Mean-variance asset allocation	27	
		5.3.1 Accounting for heteroscedasticity	32	
6.	Conclusion			
7.	7. Appendix			
References				
European Central Bank Working Paper Series 3				

2

Abstract

This paper argues that forecast estimators should minimise the loss function in a statistical, rather than deterministic, way. We introduce two new elements into the classical econometric analysis: a subjective guess on the variable to be forecasted and a probability reflecting the confidence associated to it. We then propose a new forecast estimator based on a test of whether the first derivatives of the loss function evaluated at the subjective guess are statistically different from zero. We show that the classical estimator is a special case of this new estimator, and that in general the two estimators are asymptotically equivalent. We illustrate the implications of this new theory with a simple simulation, an application to GDP forecast and an example of mean-variance portfolio selection.

Keywords: Decision under uncertainty, estimation, overfitting, asset allocation.

JEL classification: C13, C53, G11.

Non-technical summary

Classical forecast estimators typically ignore non-sample information and estimation errors due to finite sample approximations. This paper shows how these two problems are closely connected. We argue that forecast estimators should minimise the loss function in a statistical sense, rather than in the usual deterministic way. We formally introduce into the classical econometric analysis two new elements: a subjective guess on the variable to be forecasted and a probability reflecting the confidence associated to it. These elements serve to summarise the non-sample information available to the decision-maker, and allow us to formalise the interaction between judgement and data in the forecasting process. We then propose a new forecast estimator based on a test of whether the first derivatives of the loss function evaluated at the subjective guess are statistically different from zero. If this is the case, the subjective guess becomes the forecast, since the null hypothesis that it minimises the loss function cannot be rejected at the given confidence level. Otherwise, the loss function is decreased as long as its first derivatives are significantly different from zero, and a new, model-based forecast is obtained.

Classical estimators deterministically set the first order conditions equal to zero. They can therefore be obtained as a special case of our estimator by choosing a confidence level equal to zero, which corresponds to ignoring any subjective guess. Moreover, under standard regularity conditions the classical estimator is shown to be asymptotically equivalent to our estimator. We argue that the forecasting process should be characterised by a clear separation between "experts" - who should provide the subjective guess and the confidence associated to it - and econometric modellers, who should test whether the given subjective guess is supported by the available data and models used for the analysis. We also argue that under the forecasting framework developed in this paper the responsibility for good or bad forecasts is shared between "experts" and econometricians: formulating a good initial guess may be as important as the quality of the econometric model.

We illustrate the implications of this new theory with three examples. In the first one, we perform a simple Monte Carlo simulation with different sample sizes and show how under certain circumstances our new estimator outperforms the classical sample mean. In the second example, we provide an application to GDP growth rate forecasts, explaining how an initial guess on the dependent variable may be mapped into an initial guess on the parameters of the econometrician's favourite model. In the third example, we show how this forecasting theory can be used to tackle some of the well-known implementation problems of mean-variance portfolio selection models. The practical effect of our estimator is to provide a shrinkage device to be applied directly to portfolio weights. We also discuss how GARCH-type heteroscedasticity can be accounted for.

1 Introduction

That some macroeconomic and financial econometric models forecast poorly is a well-known fact. The ongoing debate was effectively summarised by the different contributions to the 100^{th} volume of the *Journal of Econometrics*, an open forum on the current state and future challenges of econometrics. In this paper, we argue that standard estimation methods may be partly responsible for the poor forecasting performance of econometric models and we propose an alternative theory of forecasting.

Forecasting is intrinsically intertwined with decision-making. Forecasts serve their purpose by helping agents to make decisions about an uncertain future. Forecast errors generate costs to the decision-maker, to the extent that different forecasts command different decisions. Since forecast errors are unavoidable in a random world, the classical theory of forecasting builds on the assumption that agents wish to minimise the expected cost associated to these errors (Granger 1969, Granger and Newbold 1986, Granger and Machina 2005). Classical forecast estimators are then obtained as the minimisers of the sample equivalent of the unobservable expected cost. We emphasise two closely related problems with this theory. First, classical forecasts do not explicitly account for non-sample information available to the decision-maker, even though subjective judgement often plays an important role in many real world forecasting processes. Second, classical estimators minimise a loss function, which depends on unknown parameters and is only a finite sample approximation of the true loss function to be minimised. Since any finite sample approximation is subject to estimation error, minimisation of this function does not necessarily coincide with the minimisation of the true loss function.

To address the first problem, we formally introduce two new elements into the classical econometric analysis: a subjective guess on the variable to be forecasted and a probability reflecting the confidence of the decision-maker in this guess. These elements serve to summarise the non-sample information available to the decision-maker, and allow us to formalise the interaction between judgement and data in the forecasting process.

As for the second problem, estimation error is explicitly taken into consideration by testing whether the subjective guess statistically minimises the loss function. This is equivalent to testing whether the first derivatives evaluated at the parameter values implied by the subjective guess are statistically equal to zero at the given confidence level. If this is the case, the subjective guess is retained as the best estimate. Otherwise, the loss function is decreased as long as its first derivatives are significantly different from zero, and a new, model-based forecast is obtained.

Classical estimators deterministically set the first order conditions equal to zero. They can therefore be obtained as a special case of our estimator by choosing a confidence level equal to zero, which corresponds to ignoring any subjective guess. Moreover, under standard regularity conditions the classical estimator is shown to be asymptotically equivalent to the estimator proposed in this paper. As the sample size grows, the true loss function is approximated with greater precision, and the subjective guess becomes less and less relevant.

We illustrate the implications of this new theory with three examples. In the first one, we perform a simple Monte Carlo simulation with different sample sizes. We show how our new estimator outperforms the classical estimator, whenever the subjective guess is close enough to the true parameter value. We argue that the dichotomy between judgement and estimation implies that the forecasting process should be characterised by a clear separation between decision-makers - who should provide the judgement and their confidence in it - and econometricians, who should test whether such judgement is supported by the available data and models used for the analysis. It also implies that the responsibility for good or bad forecasts is shared between decision-makers and econometricians: having good judgement may be as important as the quality of the econometric model.

In the second example, we provide an application to GDP growth rate forecasts. Econometric models are complicated functions of parameters which are often devoid of economic meaning. It may therefore be difficult to express a subjective guess directly on these parameters. We suggest a simple and intuitive strategy to map the subjective guess on the variable of interest to the decision-maker into values for the parameters of the econometrician's favourite model. Specifically, these "judgemental parameter values" are obtained by minimising the loss function subject to the constraint that the forecast implied by the model is equal to the subjective guess. We illustrate how this works in the context of a simple autoregressive model.

In the third example, we show how the forecasting framework proposed in this paper can be used to tackle some of the well-known implementation problems of mean-variance portfolio selection models. For a given benchmark portfolio (the subjective guess, in the terminology used before), we derive the associated optimal portfolio which increases the empirical expected utility as long as the first derivatives are statistically different from zero. From this perspective, the confidence level associated to the subjective guess may be thought of as a device against overfitting, since it crucially determines when the increase in the objective function stops to be statistically significant. Alternatively, it may be interpreted as the cost of underperforming relative to the benchmark. The practical effect of our estimator is to provide a shrinkage device to be applied directly to portfolio weights. The two ends of the shrinkage are the benchmark portfolio and the classical mean-variance optimal portfolio. The amount of shrinkage is determined by the risk preferences of the decision-maker. We also discuss how GARCH-type heteroscedasticity can be accounted for.

The paper is structured as follows. In the next section, we use a stylised statistical model to highlight the problems associated to classical estimators. In section 3, we build on this stylised model to develop the heuristics behind our new forecasting theory. Section 4 contains a formal development of the new forecasting theory. The empirical applications are in section 5. Section 6 concludes.

2 The Problem

In this section we illustrate with a simple example the problem associated with classical estimators. The intuition is the following. Classical forecast estimators approximate the expected loss function with its sample equivalent. While asymptotically this approximation is perfect, in finite samples it is not. The quality of the finite sample approximation - which is out of the econometrician's control - will crucially determine the quality of the forecasts. Assume that $\{y_t\}_{t=1}^T$ is a series of i.i.d. normally distributed observations. We are interested in the forecast θ of y_{T+1} , using the information available up to time T. Let's denote the forecast error by $e \equiv y_{T+1} - \theta$. Suppose that the agent quantifies the cost of the error with a quadratic cost function, $C(e) \equiv e^2$. The optimal forecast minimises the expected cost:

$$\min_{\theta} E[(y_{T+1} - \theta)^2] \tag{1}$$

Setting the first derivative equal to zero, the optimal point forecast is given by the expected value of y_{T+1} , leading to the optimal forecast estimator $\hat{\theta}_T \equiv T^{-1} \sum_{t=1}^{T} y_t$. But $\hat{\theta}_T$ is the minimiser of $T^{-1} \sum_{t=1}^{T} [(y_t - \theta)^2]$ and the problem can be rewritten as:

$$\min_{\theta} \{ E[(y_{T+1} - \theta)^2] + \varepsilon_T(\theta) \}$$
(2)

where $\varepsilon_T(\theta) \equiv T^{-1} \sum_{t=1}^T [(y_t - \theta)^2] - E[(y_{T+1} - \theta)^2]$. $\varepsilon_T(\theta)$ is the error induced by the finite sample approximation of the expected cost function, which by the Law of Large Numbers converges to zero only as T goes to infinity. Therefore in finite samples, classical estimators don't minimise the expected cost, but also an unbounded error term $\varepsilon_T(\theta)$ which vanishes only asymptotically.

3 An Alternative Theory of Forecasting

The question is now whether we can find an alternative estimator which may have better properties than the classical one. To answer this question we introduce extra elements into the analysis: a subjective guess on the model's parameter and a probability summarising the confidence of the forecaster in this guess.

Let's go back to the first order condition of the optimal forecast problem (1):

$$E[y_{T+1} - \theta] = 0 \tag{3}$$

The sample equivalent of this expectation evaluated at $\tilde{\theta}$ is:

$$f_T(\tilde{\theta}) \equiv T^{-1} \sum_{t=1}^T [y_t - \tilde{\theta}]$$
(4)

where $\tilde{\theta}$ is some subjective guess. $f_T(\tilde{\theta})$ is the sample realisation of the first derivatives of the expected cost function. It is a random variable which may be different from zero just because of statistical error. Under the null hypothesis that $\tilde{\theta}$ is the optimal estimator, $f_T(\tilde{\theta}) \sim N(0, \sigma_T^2(\tilde{\theta}))$, where $\sigma_T^2(\tilde{\theta}) \equiv T^{-1}E[(y_t - \tilde{\theta})^2]$.

For a given confidence level α , let $\pm \kappa_{\alpha/2}$ denote the corresponding standard normal critical values and $\pm \hat{\eta}_{\alpha/2}(\theta) \equiv \pm \sqrt{\hat{\sigma}_T^2(\theta)} \kappa_{\alpha/2}$, where $\hat{\sigma}_T^2(\theta)$ is a suitable estimator of $\sigma_T^2(\theta)$. An intuitive estimator $\hat{\theta}_T^*$ given the confidence level α is:

$$\hat{\theta}_{T}^{*} = \begin{cases} \tilde{\theta} & \text{if} - \hat{\eta}_{\alpha/2}(\tilde{\theta}) < f_{T}(\tilde{\theta}) < \hat{\eta}_{\alpha/2}(\tilde{\theta}) \\\\ \arg\min_{\theta} [f_{T}(\theta) + \hat{\eta}_{\alpha/2}(\theta)]^{2} & \text{if} \ f_{T}(\tilde{\theta}) < -\hat{\eta}_{\alpha/2}(\tilde{\theta}) \\\\ \arg\min_{\theta} [f_{T}(\theta) - \hat{\eta}_{\alpha/2}(\theta)]^{2} & \text{if} \ f_{T}(\tilde{\theta}) > \hat{\eta}_{\alpha/2}(\tilde{\theta}) \end{cases}$$

That is, given the subjective guess $\tilde{\theta}$ and the confidence level α , if the null hypothesis $H_0: f_T(\tilde{\theta}) = 0$ cannot be rejected, the subjective guess $\tilde{\theta}$ becomes the forecast. Rejection of the null signals that the subjective guess can be statistically improved, until $f_T(\hat{\theta}_T^*)$ is exactly equal to its critical value. Note that as long as $\hat{\sigma}_T^2(\theta)$ is a consistent estimator of $\sigma_T^2(\theta)$, $\hat{\eta}_{\alpha/2}(\theta)$ converges to zero as T goes to infinity. Therefore $\hat{\theta}_T^*$ converges asymptotically to the classical estimator and is consistent.

This estimator has a natural economic interpretation in terms of the expected cost/utility function used in the forecasting problem. For a given subjective guess $\tilde{\theta}$, it answers the following question: Can the forecaster increase his/her expected utility in a statistically significant way? If the answer is no, i.e. if one cannot reject the null that the first derivatives evaluated at $\tilde{\theta}$ are equal to zero, $\tilde{\theta}$ should be taken as the forecast. If, on the contrary, the answer is yes, the econometrician will move the parameter θ as long as the first derivatives are statistically different from zero. S/he will stop only when $\hat{\theta}_T^*$ is such that the empirical expected utility cannot be increased any more in a statistically significant way.

The confidence level α may have different interpretations. It may be interpreted as the confidence of the forecaster in the subjective guess and in this case it should reflect the knowledge of the environment in which the forecast takes place. Alternatively, it may be thought of as the probability of committing type I errors, i.e. of rejecting the null when $\tilde{\theta}$ is indeed the optimal forecast. Finally, since it determines when the increase in expected utility stops to be statistically significant, it may be seen as a device against overfitting. The forecaster will choose a low α whenever s/he is reasonably confident in the subjective guess $\tilde{\theta}$, and/or if the cost of committing type I errors is high, and/or if she is concerned about overfitting.

Note that in the classical paradigm there is no place for subjective guesses and therefore $\alpha = 100\%$: in this case $\kappa_{\alpha/2} = 0$ and $\hat{\theta}_T^*$ is simply the solution obtained by setting the first derivatives (4) equal to zero.

4 Generalisation

In this section we generalise the analysis of the previous section. We formally define a new estimator which depends on a subjective guess and the confidence associated to it, and establish its relationships with classical estimators. This new estimator is obtained by adding a constraint on the first derivatives to the classical optimisation problem. We show that the classical estimator is a special case of our new estimator and that the two estimators are asymptotically equivalent.

Classical forecasts typically maximise some objective function that depends on parameters, data and sample size. Following the framework of Newey and McFadden (1994), denote with $\hat{Q}_T(\theta)$ such objective function, where θ is the vector of parameters belonging to the k-dimensional parameter space Θ . We assume the following: Condition 1 (Uniform Convergence) $\hat{Q}_T(\theta)$ converges uniformly in probability to $Q_0(\theta)$.

Condition 2 (Identification) $Q_0(\theta)$ is uniquely maximised at θ^0 .

Condition 3 (Compactness) Θ is compact.

Condition 4 (Continuity) $Q_0(\theta)$ is continuous.

We will refer to θ^0 as to the true (or pseudo-true) parameter (see White 1994 for a treatment of quasi-maximum likelihood estimation). These are the standard conditions needed for consistency results of extremum estimators (see theorem 2.1 of Newey and McFadden 1994). For the present context we need to impose also the following conditions:

Condition 5 $\theta^0 \in interior(\Theta)$.

Condition 6 (Differentiability) $\hat{Q}_T(\theta)$ is continuously differentiable.

Condition 7 (Asymptotic Normality) $\sqrt{T}\nabla_{\theta}\hat{Q}_{T}(\theta^{0}) \xrightarrow{d} N(0,\Sigma).$

The first derivative of $\hat{Q}_T(\theta)$ evaluated at a subjective guess $\tilde{\theta}$ is a kdimensional random variable. Denote with $\hat{\Sigma}_T$ a \sqrt{T} -consistent estimate of Σ . Under the null hypothesis $H_0: \tilde{\theta} = \theta^0$, the above conditions imply:

$$\hat{z}_T(\theta) \equiv T \nabla'_{\theta} \hat{Q}_T(\tilde{\theta}) \hat{\Sigma}_T^{-1} \nabla_{\theta} \hat{Q}_T(\tilde{\theta}) \xrightarrow{d} \chi_k^2 \tag{5}$$

The classical and new estimators are given in the following definitions.

Definition 1 (Classical Estimator) The classical estimator is $\hat{\theta}_T = \arg \max_{\theta} \hat{Q}_T(\theta)$.

Definition 2 (New Estimator) Let $\tilde{\theta}$ denote the subjective guess and, for a given confidence level α , let $\eta_{\alpha,k}$ denote the critical value of χ_k^2 . The **new** estimator $\hat{\theta}_T^*$ is defined as follows:

1. if
$$\nexists \hat{\theta} \in \Theta$$
 s.t. $\hat{z}_T(\hat{\theta}) > \eta_{\alpha,k}$ and $\hat{Q}_T(\hat{\theta}) > \hat{Q}_T(\hat{\theta})$:

$$\hat{\theta}_T^* = \tilde{\theta}; \tag{6}$$

2. if $\exists \ \ddot{\theta} \in \Theta \ s.t. \ \hat{z}_T(\ddot{\theta}) > \eta_{\alpha,k}$ and $\hat{Q}_T(\ddot{\theta}) > \hat{Q}_T(\tilde{\theta}), \ \hat{\theta}_T^*$ is the solution of the following constrained maximisation problem:

$$\max_{\theta} \hat{Q}_T(\theta) \tag{7}$$

s.t. $\hat{z}_T(\theta) = \eta_{\alpha,k}$

According to the above definition, the new estimator can be obtained from the classical maximisation problem, by adding a constraint. The role of the constraint is to make sure that the classical maximisation problem takes explicitly into account estimation errors. If at a given subjective guess $\tilde{\theta}$ and confidence level α , the likelihood cannot be increased in a statistically significant way, the subjective guess $\tilde{\theta}$ should be retained as the forecast estimator. If instead the null hypothesis that $\tilde{\theta}$ maximises the likelihood is rejected, the subjective guess can be statistically improved upon, until the first derivatives are not significantly different from zero. The slightly cumbersome conditions of points 1 and

2 in Definition 2 are needed to account for the fact that the first order conditions can generally have multiple roots.

Note that the test statistic in (5) relies on an asymptotic approximation to its distribution. If it is suspected that with the available sample size such approximation may be poor, one could resort to bootstrap methods to improve the accuracy of the estimator.

The following theorem shows that the new estimator is consistent and establishes its relationship with the classical estimator.

Theorem 1 (Properties of the New Estimator) Under Conditions 1-7 the estimator $\hat{\theta}_T^*$ of Definition 2 satisfies the following properties:

- 1. If $\alpha = 100\%$, $\hat{\theta}_T^*$ is the classical estimator;
- 2. If $\alpha > 0$, $\hat{\theta}_T^*$ converges in probability to the classical estimator.

Proof. See Appendix.

The intuition behind this result is that as the sample size grows the distribution of the first derivatives will be more and more concentrated around its true mean. If the subjective guess $\tilde{\theta}$ coincides with the true parameter, the chi-square statistic (5) will be lower than its critical value for large T and according to Definition 2 the estimator will be $\hat{\theta}_T^* = \tilde{\theta}$. If, on the other hand, the true parameter is different from the subjective guess, the constraint of the maximisation problem in Definition 2, coupled with Conditions 1-7, will guar-

Working Paper Series No. 584

antee that $\hat{\theta}_T^*$ will get closer and closer to the true parameter, as the sample size grows.

5 Examples

We illustrate some of the implications of our theory with three examples. The first one is based on a Monte Carlo simulation. We show how the performance of the new estimator crucially depends on the quality of the subjective guess and the confidence associated to it. We argue that the choice of the subjective guess should be independent of the econometric model used for estimation.

The second example is an application to U.S. GDP forecast. We show how one can map a subjective guess on future GDP growth rates into subjective guesses on the parameters of the econometrician's favourite model. We provide an illustration using an autoregressive model to forecast quarterly GDPs.

In the third example, we estimate the optimal portfolio weights maximising a mean-variance utility function. We highlight how the theory proposed in this paper naturally takes into account the impact of estimation errors - which typically plague standard mean-variance optimisers - by shrinking portfolio weights estimates from a given benchmark towards the classical estimates.

5.1 Simulation

We generated random draws from a standard normal distribution, with different sample sizes, T = 5, 20, 60, 120, 240, 1000. For each series, we estimated the classical forecast estimator $(\hat{\theta}_T = T^{-1} \sum_{t=1}^T y_t)$ and the one proposed in this paper $(\hat{\theta}_T^*)$, using a quadratic loss function, i.e. $\hat{Q}_T(\theta) \equiv -C(\theta) =$ $-T^{-1} \sum_{t=1}^T (y_t - \theta)^2$. In the estimation of $\hat{\theta}_T^*$, we set $\alpha = 10\%$. Next, we computed the expected costs associated to these estimators and to the subjective guess \tilde{a} $(E^i[C(\hat{\theta}_T)], E^i[C(\hat{\theta}_T^*)]$ and $E^i[C(\tilde{\theta})]$ with a Monte Carlo simulation (with 10,000 random draws from the normal distribution). We repeated this procedure 5000 times and then averaged the expected utilities, i.e. $E[C(\hat{\theta}_T)] = \sum_{i=1}^{5000} E^i[C(\hat{\theta}_T)]/5000$ and the same for the other estimators. The results are reported in table 1.

The major points to be highlighted are the following. First, the new estimator $\hat{\theta}_T^*$ may be biased but is consistent. Second, in small samples the classical estimator $\hat{\theta}_T$ performs worse than $\hat{\theta}_T^*$ when the subjective guess $\tilde{\theta}$ is reasonably close to the true value. Third, in large samples, the performance of $\hat{\theta}_T$ and $\hat{\theta}_T^*$ becomes roughly equivalent, independently of the subjective guess $\tilde{\theta}$.

5.1.1 Discussion

These results have implications for the organisation of the forecasting process of any institution interested in forecasting. There should be a clear separation

		Т	5	20	60	120	240	1000
$\tilde{ heta}$	$E[C(\tilde{\theta})]$	$E[C(\hat{\theta}_T)]$	1.2068	1.0497	1.0167	1.0083	1.0039	1.0009
0	1		1.0045	1.0012	1.0000	1.0003	0.9999	0.9999
0.05	1.0025		1.0071	1.0035	1.0022	1.0024	1.0020	1.0016
0.1	1.01	$E[C(\hat{\theta}_T^*)]$	1.0142	1.0102	1.0085	1.0082	1.0070	1.0034
0.5	1.2500		1.2151	1.1452	1.0621	1.0317	1.0150	1.0036
1	2.0000		1.6572	1.2019	1.0629	1.0317	1.0150	1.0036

Table 1: Monte Carlo comparison of expected cost functions associated to different estimators. We formatted in bold the cases where the new estimator outperforms the classical estimator.

between the decision-maker providing the subjective guess and the confidence associated to it, and the econometrician whose task is to check whether such a subjective guess is supported by the available data or whether it can be improved. In particular, the formulation of the subjective guess should be independent of the econometric model used to evaluate it. In the previous example a subjective guess based on the OLS estimator would never be rejected by the data, but it would also have very little value added.

Within this new framework, the responsibility of good or bad forecasts is shared between the decision-maker and the econometrician. High confidence in a bad subjective guess would inevitably result in poor forecasts (in small samples). Therefore, under the new forecasting framework developed in this paper, formulating a good subjective guess may become as important as having a good econometric model.

5.2 Forecasting U.S. GDP

We illustrate how the theory of section 4 can be applied to forecast the U.S. real GDP. A possible difficulty in implementing the theory is related to the formulation of a subjective guess on parameters of an econometric model about which the decision-maker may know nothing or very little. We propose a simple strategy to map a subjective guess on the variable of interest to the decision-maker (GDP in this case) into subjective guesses on the parameters of the econometrician's favourite model.

In principle, it is possible to express a subjective guess directly on the parameter vector θ or indirectly on the dependent variable y_{T+1} to be forecasted. If the decision-maker can formulate a guess on θ , the theory of section 4 can be applied directly. In most circumstances, however, it may be more natural to have a judgement about the future behaviour of y_{T+1} , rather than about abstract model parameters. Let's denote this subjective guess as \tilde{y}_{T+1} . Using the notation of section 4, this can be translated into a subjective guess on θ as follows:

$$\tilde{\theta} = \arg \max_{\theta} \hat{Q}_T(\theta)$$
(8)

s.t. $\hat{y}_{T+1}(\theta) = \tilde{y}_{T+1}$



where $\hat{y}_{T+1}(\theta)$ is the model's forecast conditional on the parameter vector θ . The subjective guess \tilde{y}_{T+1} is mapped into a subjective guess on the parameter vector by choosing the $\tilde{\theta}$ that maximises the likelihood subject to the constraint that the forecast at time T is equal to \tilde{y}_{T+1} .

Let's consider, for concreteness, an application to quarterly GDP forecasting, using an AR(4) model:

$$y_t = \theta_0 + \sum_{i=1}^4 \theta_i y_{t-i} + \varepsilon_t \tag{9}$$

If the model is estimated via OLS, we have:

$$\hat{Q}_T(\theta) \equiv -T^{-1} \sum_{t=1}^T [y_t - \hat{y}_t(\theta)]^2$$
(10)

where $\hat{y}_t(\theta) \equiv \theta_0 + \sum_{i=1}^4 \theta_i y_{t-i}$. The score evaluated at $\tilde{\theta}$ is

$$\nabla_{\theta} \hat{Q}_T(\tilde{\theta}) = 2T^{-1} \sum_{t=1}^T \hat{\varepsilon}_t(\tilde{\theta}) \nabla_{\theta} \hat{y}_t(\tilde{\theta})$$
(11)

where $\hat{\varepsilon}_t(\tilde{\theta}) \equiv y_t - \hat{y}_t(\tilde{\theta})$ and $\nabla_{\theta}\hat{y}_t(\tilde{\theta}) \equiv [1, y_{t-1}, y_{t-2}, y_{t-3}, y_{t-4}]'$. We estimate the asymptotic variance-covariance matrix of the score using standard heteroscedasticity-consistent estimators (White 1980):

$$\hat{\Sigma}_T \equiv 4T^{-1} \sum_{t=1}^T \hat{\varepsilon}_t(\tilde{\theta})^2 \nabla_{\theta} \hat{y}_t(\tilde{\theta}) \nabla'_{\theta} \hat{y}_t(\tilde{\theta})$$
(12)

We estimate this model using quarterly data for the U.S. real GDP growth rates. The data are taken from the FRED[®] database¹. The data has been seasonally adjusted and our sample runs from Q1 1983 to Q3 2005, with 90 observations. The growth rates are computed as log differences.

	$\tilde{y}_{T+1} = 3\%$		$\tilde{y}_{T+1} = 5\%$		
$\hat{ heta}_T$	$ ilde{ heta}$	${\hat heta}_T^*$	$ ilde{ heta}$	$\hat{ heta}_T^*$	
1.65	1.54	1.54	2.73	2.69	
0.23	0.21	0.21	0.42	0.30	
0.36	0.37	0.37	0.30	0.20	
-0.16	-0.18	-0.18	-0.03	-0.14	
0.04	0.05	0.05	-0.04	-0.13	

Table 2: Subjective guesses and estimated parameters associated to different subjective guesses on Q4 2005 GDP growth rates (3% and 5%). A subjective guess of 3% is not rejected by the data and maps into parameter values very close to the OLS $\hat{\theta}_T$. A subjective guess of 5%, instead, is rejected by the data, resulting in parameter estimates different from the parameter guess.

One quarter ahead GDP forecasts (guess = 3%)



Figure 1: Plot of the forecasts associated to the subjective guess of 3% GDP growth rate for Q4 2005 (thick dashed line) and to the OLS estimate (thin solid line). Note that the forecast at the end of the sample of the dashed line is exactly 3% by construction.





Figure 2: Plot of the forecasts associated to the subjective guess of 5% GDP growth rate for Q4 2005 (thick dashed line) and to the new estimator (thin solid line). The subjective guess of 5% is rejected by the data and results in very different forecasts associated to the estimated parameters.



For illustratory purposes, we consider two different subjective guesses for GDP growth in the next quarter (Q4 2005), $\tilde{y}_{T+1} = 3\%$ and $\tilde{y}_{T+1} = 5\%$, both with a confidence level $\alpha = 10\%$. The results are reported in table 2 and figures 1-2. As we can see from the table, $\tilde{y}_{T+1} = 3\%$ maps into a parameter guess $\tilde{\theta}$ which cannot be rejected by the data ($\hat{\theta}_T^* = \tilde{\theta}$). These parameter values are also very close to the OLS estimates $\hat{\theta}_T$, resulting in very similar forecasts (see figure 1). Note that in figure 1 the forecast at Q4 2005 associated to $\hat{\theta}_T^*$ is equal to 3%, the original subjective guess ($\tilde{y}_{T+1} = 3\%$).

The other subjective guess, $\tilde{y}_{T+1} = 5\%$, is instead rejected by the data at the chosen confidence level, resulting in parameter estimates $\hat{\theta}_T^*$ which are different from the parameter guess $\tilde{\theta}$. The implications can be seen in figure 2, where we plot the in-sample forecasts associated to these two parameter values. There are very remarkable differences between the two plotted time series, the one based on $\hat{\theta}_T^*$ (labelled "New Estimator") being a couple of percentage points lower than the one based on $\tilde{\theta}$ (labelled "Guess"). The out of sample forecast at Q4 2005 associated to $\hat{\theta}_T^*$ is 3.49% and the OLS forecast is 3.19%, both definitely lower than the subjective guess of 5%.

5.3 Mean-Variance Asset Allocation

In this section we illustrate how our theory can be applied to the mean-variance portfolio selection problem. Markowitz's (1952) mean-variance model provides

¹See http://research.stlouisfed.org/fred2.

the standard benchmark for portfolio allocation. It formalises the intuition that investors optimise the trade off between returns and risks, resulting in optimal portfolio allocations which are a function of expected return, variance (the proxy used for risk) and the degree of risk aversion of the decision-maker. Despite its theoretical appeal, it is well known that standard implementations of this model produce portfolio allocations with no economic intuition and little (if not negative) investment value. These problems were initially pointed out, among others, by Jobson and Korkie (1981), who used a Monte Carlo experiment to show that estimated mean-variance frontiers can be quite far away from the true ones. The crux of the problem is colourfully, but effectively highlighted by the following quotation of Michaud (1998, p. 3):

"[Mean-variance optimizers] overuse statistically estimated information and magnify the impact of estimation errors. It is not simply a matter of garbage in, garbage out, but, rather, a molehill of garbage in, a mountain of garbage out."

The problem can be restated in terms of the theory developed in section 4. Classical estimators maximise the empirical expected utility, without taking into consideration whether this maximisation is statistically significant or not. Our theory provides a natural alternative. For a given benchmark portfolio (the subjective guess $\tilde{\theta}$ in the notation of section 4) and a confidence level α , the resulting optimal portfolio is the one which increases the empirical expected utility as long as the first derivatives are statistically different from zero.

Working Paper Series No. 584

To formalise this discussion, consider a portfolio with N + 1 assets. Denote with θ the N-vector of weights associated to the first N assets entering a given portfolio, and denote with $y_t(\theta)$ the portfolio return at time t, where the dependence on the individual asset weights has been made explicit. Since all the weights must sum to one, note that $\theta_{N+1} = 1 - \sum_{i=1}^{N} \theta_i$, where θ_{N+1} is the weight associated to the $(N + 1)^{th}$ asset of the portfolio. Let's assume an investor wants to maximise a trade-off between mean and variance of portfolio returns, resulting in the following expected utility function:

$$Q_{0}(\theta) \equiv U[\theta; y_{T+1}] = E[y_{T+1}(\theta)] - \lambda V[y_{T+1}(\theta)]$$
(13)
= $E[y_{T+1}(\theta)] - \lambda \{ E[y_{T+1}^{2}(\theta)] - E[y_{T+1}(\theta)]^{2} \}$

where λ describes the investor's attitude towards risk. The empirical analogue is:

$$\hat{Q}_{T}(\theta) \equiv \hat{U}_{T}[\theta; \{y_{t}\}_{t=1}^{T}] = T^{-1} \sum_{t=1}^{T} y_{t}(\theta) - \lambda \{T^{-1} \sum_{t=1}^{T} y_{t}^{2}(\theta) - [T^{-1} \sum_{t=1}^{T} y_{t}(\theta)]^{2}\}$$
(14)

The first order conditions are:

$$\nabla_{\theta} \hat{U}_{T}[\theta; \{y_{t}\}_{t=1}^{T}] = T^{-1} \sum_{t=1}^{T} \nabla_{\theta} y_{t}(\theta) - \lambda \{T^{-1} 2 \sum_{t=1}^{T} y_{t}(\theta) \nabla_{\theta} y_{t}(\theta) - 2[T^{-1} \sum_{t=1}^{T} y_{t}(\theta)]T^{-1} \sum_{t=1}^{T} \nabla_{\theta} y_{t}(\theta)\}$$

where $\nabla_{\theta} y_t(\theta) \equiv y_t^N - y_t^{N+1}\iota$, y_t^N is an N-vector containing the returns at time t of the first N assets, y_t^{N+1} is the return at time t of the $(N+1)^{th}$ asset, and

α	100%	0.01%	$10^{-8}\%$	$10^{-14}\%$
$\operatorname{Var}(\hat{\boldsymbol{\theta}}_T^*)$	0.0199	0.0057	0.0018	0.0001
\min	-0.16	-0.07	-0.01	0.02
\max	0.45	0.19	0.13	0.11

Table 3: This table reports variance, minimum and maximum of the optimal vector of portfolio weights associated to different confidence levels, starting from an equally weighted portfolio. The lower the confidence level, the closer the optimal allocation to the benchmark portfolio.

 ι is an N-vector of ones.

We apply the methodology developed in section 4 to monthly log returns of 15 stocks of the Dow Jones Industrial Average (DJIA) index, as of July 15, 2005. The sample runs from January 1, 1987 to July 1, 2005, for a total of 223 observations. We set $\lambda = 4$ and use as subjective guess the equally weighted portfolio. We tried different confidence levels α , which in this case can be interpreted as the cost of underperforming relative to the benchmark portfolio. In table 3 we report variance, minimum and maximum of the optimised portfolio weights. As we can see from this table, a low α results in a lower dispersion of the portfolio weights, implying a portfolio closer to the equally weighted benchmark. This can be visually seen in figure 3.

The case with $\alpha = 100\%$ corresponds to the standard implementation of the mean-variance model (i.e., it corresponds to the case where the sample



Figure 3: Plots of the optimal weights associated to the different confidence levels α .

estimates of expected returns and variance-covariances are substituted into the analytical solution of the optimal portfolio weights). To see why, rewrite (13) as

$$U[\theta; y_{T+1}] = \bar{\theta}' E[y_{T+1}] - \lambda \bar{\theta}' V[y_{T+1}] \bar{\theta}$$
(15)

where $\bar{\theta} \equiv [\theta', \theta_{N+1}]'$. Standard implementations maximise this expected utility analytically (subject to the constraint that the weights sum to 1) and then substitute into the solution the sample estimates for $E[y_{T+1}]$ and $V[y_{T+1}]$. This is equivalent to first substituting the sample estimates for $E[y_{T+1}]$ and $V[y_{T+1}]$ and then maximising with respect to θ , which in turn is equivalent to directly maximising (14).

Figure 3 offers an alternative interpretation of our estimator: it is a shrinkage estimator, shrinking the portfolio weights from the benchmark $\tilde{\theta}$ towards the classical estimator. The confidence level α determines the amount of shrinkage: the higher the α , the stronger the shrinkage effect towards the classical weights. Setting $\alpha = 100\%$ corresponds to a complete shrinkage, resulting in the usual optimal portfolio allocation.

5.3.1 Accounting for heteroscedasticity

The previous analysis has assumed constant means and variance-covariance matrices. In practice, it is well known that financial returns at higher frequencies are heteroscedatic. This feature of the data can be effectively captured by Engle's (1982) and Bollerslev's (1986) GARCH models. One concrete drawback of these models in applications to multivariate asset allocation problems is the curse of dimensionality: as the number of assets included in the portfolio increases, the number of parameters to be estimated in multivariate GARCH models grows exponentially. Manganelli (2004) proposes a solution to this problem, which can be directly extended to the new theory developed in this paper. The idea is to work with univariate portfolio GARCH models. The multivariate dimension of the portfolio allocation problem is recovered via the derivatives of the estimated GARCH variance with respect to the portfolio weights.

Assume that portfolio returns $y_t(\theta)$ are modelled as a zero-mean process with a GARCH(p,q) conditional variance:

$$y_t(\theta) = \sqrt{h_t(\beta, \theta)}\varepsilon_t \qquad \varepsilon_t \sim (0, 1)$$
(16)
$$h_t(\beta, \theta) = z'_t(\beta, \theta) \cdot \beta$$

where $z_t(\beta, \theta) \equiv [1, y_{t-1}^2(\theta), ..., y_{t-p}^2(\theta), h_{t-1}(\beta, \theta), ..., h_{t-q}(\beta, \theta)]'$ and we have made explicit the dependence on the GARCH parameter β and the vector of portfolio weights θ . The utility (13) becomes $U[\beta, \theta; y_{T+1}(\theta)] = -\lambda h_{T+1}(\beta, \theta)$ and its empirical analogue is:

$$U[\hat{\beta}_T, \theta; \{y_t(\theta)\}_{t=1}^T] = -\lambda h_{T+1}(\hat{\beta}_T, \theta)$$

where $h_{T+1}(\hat{\beta}_T, \theta)$ represents the GARCH variance evaluated at the maximum likelihood estimate $\hat{\beta}_T$. Notice that in this context $\hat{\beta}_T$ depends on the portfolio weights θ : changes in θ imply a different time series $\{y_t(\theta)\}_{t=1}^T$ and therefore different estimates $\hat{\beta}_T$. The first derivatives are:

$$\nabla_{\theta} U[\hat{\beta}_{T}, \theta; \{y_{t}(\theta)\}_{t=1}^{T}] = -\lambda \nabla_{\theta} h_{T+1}(\hat{\beta}_{T}, \theta)$$
$$= -\lambda [\nabla_{\theta} z'_{T+1}(\hat{\beta}_{T}, \theta) \cdot \hat{\beta}_{T} + \nabla_{\theta} \hat{\beta}'_{T} \cdot z_{T+1}(\hat{\beta}_{T}, \theta)]$$

The formula to compute $\nabla_{\theta} \hat{\beta}_T$ is given in theorem 1 of Manganelli (2004). To obtain the distribution of the first derivatives, we use a mean value expansion around β^0 . Under the null hypothesis that the allocation θ is optimal, the first derivatives evaluated at the true parameter should be zero and we get:

$$\nabla_{\theta} U[\hat{\beta}_{T}, \theta; \{y_{t}(\theta)\}_{t=1}^{T}] = \nabla_{\theta} U[\beta^{0}, \theta; \{y_{t}(\theta)\}_{t=1}^{T}] + \nabla_{\theta\beta} U[\beta^{*}_{T}, \theta; \{y_{t}(\theta)\}_{t=1}^{T}] \cdot (\hat{\beta}_{T} - \beta^{0})$$
$$= -\lambda \nabla_{\theta\beta} h_{T+1}(\beta^{*}_{T}, \theta) \cdot (\hat{\beta}_{T} - \beta^{0})$$

where β_T^* lies between $\hat{\beta}_T$ and β^0 .

By the standard GARCH results, $\sqrt{T}(\hat{\beta}_T - \beta^0) \stackrel{d}{\longrightarrow} N(0, \Phi)$. See Bollerslev, Engle and Nelson (1994) for a discussion of how to estimate Φ . Since $\hat{\beta}_T$ is a consistent estimator of β^0 , under the null hypothesis that a given allocation θ is optimal, the first derivatives of the utility function will have the following asymptotic distribution:

$$\sqrt{T}\nabla_{\theta} U[\theta, \hat{\beta}_T; \{y_t(\theta)\}_{t=1}^T] \xrightarrow{d} N(0, \lambda^2 \nabla_{\theta\beta} h_{T+1}(\beta^0, \theta) \cdot \Phi \cdot \nabla_{\theta\beta} h_{T+1}(\beta^0, \theta)')$$



The term $\nabla_{\theta\beta}h_{T+1}(\beta^0, \theta)$ can be consistently estimated by plugging $\hat{\beta}_T$ into its analytical expression:

$$\nabla_{\theta\beta}h_{T+1}(\hat{\beta}_T,\theta) = \nabla_{\theta}z'_{T+1}(\hat{\beta}_T,\theta) + (\hat{\beta}'_T \otimes I_N) \cdot \nabla_{\theta\beta}z'_{T+1}(\hat{\beta}_T,\theta) + \\ + \nabla_{\theta}\hat{\beta}'_T \cdot (\nabla_{\beta}z'_{T+1}(\hat{\beta}_T,\theta))' + (z'_{T+1}(\hat{\beta}_T,\theta) \otimes I_N) \cdot \nabla_{\theta\beta}\hat{\beta}'_T$$

where \otimes denotes the Kronecker product and I_N is an identity matrix of dimension N, the dimension of θ . It is now possible to derive the chi-square statistic and apply the methodology of section 4.

6 Conclusion

Classical forecast estimators typically ignore non-sample information and estimation errors due to finite sample approximations. In this paper we pointed out how these two problems are connected. We argued that forecast estimators should optimise the objective function in a statistical sense, rather than in the usual deterministic way. We formally introduced into the classical econometric analysis two new elements: a subjective guess on the variable to be forecasted and a confidence associated to it. Their role is to explicitly take into consideration the non-sample information available to the decision-maker. These elements served to define a new estimator, which statistically optimises the objective function, and to formalise the interaction between judgement and data in the forecasting process.

We provided three empirical applications, which give strong support to our theory. We argued that there should be a clear separation between the decisionmaker - who should provide the subjective guess and the confidence associated to it - and the econometrician - whose task is to check whether such subjective guess is supported by the available data or whether it can be improved. We showed how a subjective guess on the variable to be forecasted can be mapped into a subjective guess on the parameters of the econometrician's favourite model. Finally, we illustrated how our new estimator may provide a satisfactory solution to the well-known implementation problems of the mean-variance asset allocation model.

7 Appendix

Proof of Theorem 1 (Properties of the New Estimator) - 1. If $\alpha = 100\%$, $\eta_{\alpha,k} = 0$ and the constraint in (7) becomes $\hat{z}_T(\theta) = 0$. This implies $\nabla_{\theta}\hat{Q}_T(\theta) = 0$, which coupled with (7) implies $\hat{\theta}_T^* = \hat{\theta}_T$, where $\hat{\theta}_T$ is defined in Definition 1.

2. Let $\theta^0 \equiv p \lim_{T \to \infty} \hat{\theta}_T$. We need to show that $\hat{\theta}_T^* \xrightarrow{p} \theta^0$. By (7) and Condition 2, this is equivalent to show that $||\nabla_{\theta}\hat{Q}_T(\hat{\theta}_T^*)|| \xrightarrow{p} 0$. Suppose by contradiction that $||\nabla_{\theta}\hat{Q}_T(\hat{\theta}_T^*)|| \xrightarrow{p} c \neq 0$. Then $\Pr(\hat{z}_T(\hat{\theta}_T^*) > \eta_{\alpha,k}) =$ $\Pr(\nabla_{\theta}'\hat{Q}_T(\hat{\theta}_T^*)\hat{\Sigma}_T^{-1}\nabla_{\theta}\hat{Q}_T(\hat{\theta}_T^*) > \eta_{\alpha,k}/T)$. But since $\nabla_{\theta}'\hat{Q}_T(\hat{\theta}_T^*)\hat{\Sigma}_T^{-1}\nabla_{\theta}\hat{Q}_T(\hat{\theta}_T^*)$ is bounded in probability above zero and $\eta_{\alpha,k}/T$ converges to 0 as T goes to infinity, for any $q \in [0, 1)$ there must exist a T^* such that, for any $T > T^*$, $\Pr(\nabla_{\theta}'\hat{Q}_T(\hat{\theta}_T^*)\hat{\Sigma}_T^{-1}\nabla_{\theta}\hat{Q}_T(\hat{\theta}_T^*) > \eta_{\alpha,k}/T) > q$. This implies a violation of the constraint in (7) and therefore a contradiction.

Working Paper Series No. 584

References

- Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics*, 31: 307-327.
- Bollerslev, T., R.F. Engle and D.B. Nelson (1994), ARCH Models, in *Handbook of Econometrics*, Volume IV, edited by R.F. Engle and D.L. McFadden, Elsevier Science.
- [3] Engle, R.F. (1982), Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, *Econometrica*, 50(4): 987-1006.
- [4] Granger, C. W. J. (1969), Prediction with a Generalized Cost of Error Function, Operational Research Quarterly, 20: 451-468.
- [5] Granger, C. W. J. and M. Machina (2005), Forecasting and Decision Theory, forthcoming in the *Handbook of Economic Forecasting*, edited by G. Elliott, C. W. J. Granger and A. Timmermann.
- [6] Granger, C. W. J. and P. Newbold (1986), Forecasting Economic Time Series, Academic Press, London.



- [7] Jobson, J. D. and B. Korkie (1981), Estimation for Markowitz Efficient Portfolios, Journal of the American Statistical Association, 75: 544-554.
- [8] Manganelli, S. (2004), Asset Allocation by Variance Sensitivity Analysis, Journal of Financial Econometrics 2(3): 370-389.
- [9] Markowitz, H.M. (1952), Portfolio Selection, Journal of Finance, 39: 47-61.
- [10] Michaud, R. O. (1998), Efficient Asset Allocation. A practical Guide to Stock Portfolio Optimization and Asset Allocation, Harvard Business School Press, Boston, Massachusetts.
- [11] Newey, W. and D. McFadden (1994), Large Sample Estimation and Hypothesis Testing, in *Handbook of Econometrics*, Volume IV, edited by R.F. Engle and D.L. McFadden, Elsevier Science.
- White, H. (1980), A Heteroskedaticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity, *Econometrica*, 48: 817-838.
- [13] White, H. (1994), Estimation, Inference and Specification Analysis, Cambridge, UK: Cambridge University Press.

European Central Bank Working Paper Series

For a complete list of Working Papers published by the ECB, please visit the ECB's website (http://www.ecb.int)

- 531 "Market power, innovative activity and exchange rate pass-through in the euro area" by S. N. Brissimis and T. S. Kosma, October 2005.
- 532 "Intra- and extra-euro area import demand for manufactures" by R. Anderton, B. H. Baltagi,F. Skudelny and N. Sousa, October 2005.
- 533 "Discretionary policy, multiple equilibria, and monetary instruments" by A. Schabert, October 2005.
- 534 "Time-dependent or state-dependent price setting? Micro-evidence from German metal-working industries" by H. Stahl, October 2005.
- 535 "The pricing behaviour of firms in the euro area: new survey evidence" by S. Fabiani, M. Druant, I. Hernando, C. Kwapil, B. Landau, C. Loupias, F. Martins, T. Y. Mathä, R. Sabbatini, H. Stahl and A. C. J. Stokman, October 2005.
- 536 "Heterogeneity in consumer price stickiness: a microeconometric investigation" by D. Fougère, H. Le Bihan and P. Sevestre, October 2005.
- 537 "Global inflation" by M. Ciccarelli and B. Mojon, October 2005.
- 538 "The price setting behaviour of Spanish firms: evidence from survey data" by L. J. Álvarez and I. Hernando, October 2005.
- 539 "Inflation persistence and monetary policy design: an overview" by A. T. Levin and R. Moessner, November 2005.
- 540 "Optimal discretionary policy and uncertainty about inflation persistence" by R. Moessner, November 2005.
- 541 "Consumer price behaviour in Luxembourg: evidence from micro CPI data" by P. Lünnemann and T. Y. Mathä, November 2005.
- 542 "Liquidity and real equilibrium interest rates: a framework of analysis" by L. Stracca, November 2005.
- 543 "Lending booms in the new EU Member States: will euro adoption matter?" by M. Brzoza-Brzezina, November 2005.
- 544 "Forecasting the yield curve in a data-rich environment: a no-arbitrage factor-augmented VAR approach" by E. Mönch, November 2005.
- 545 "Trade integration of Central and Eastern European countries: lessons from a gravity model" by M. Bussière, J. Fidrmuc and B. Schnatz, November 2005.
- 546 "The natural real interest rate and the output gap in the euro area: a joint estimation" by J. Garnier and B.-R. Wilhelmsen, November 2005.
- 547 "Bank finance versus bond finance: what explains the differences between US and Europe?" by F. de Fiore and H. Uhlig, November 2005.

- 548 "The link between interest rates and exchange rates: do contractionary depreciations make a difference?" by M. Sánchez, November 2005.
- 549 "Eigenvalue filtering in VAR models with application to the Czech business cycle" by J. Beneš and D. Vávra, November 2005.
- 550 "Underwriter competition and gross spreads in the eurobond market" by M. G. Kollo, November 2005.
- 551 "Technological diversification" by M. Koren and S. Tenreyro, November 2005.
- 552 "European Union enlargement and equity markets in accession countries" by T. Dvorak and R. Podpiera, November 2005.
- 553 "Global bond portfolios and EMU" by P. R. Lane, November 2005.
- 554 "Equilibrium and inefficiency in fixed rate tenders" by C. Ewerhart, N. Cassola and N. Valla, November 2005.
- 555 "Near-rational exuberance" by J. Bullard, G. W. Evans and S. Honkapohja, November 2005.
- 556 "The role of real wage rigidity and labor market frictions for unemployment and inflation dynamics" by K. Christoffel and T. Linzert, November 2005.
- 557 "How should central banks communicate?" by M. Ehrmann and M. Fratzscher, November 2005.
- 558 "Ricardian fiscal regimes in the European Union" by A. Afonso, November 2005.
- 559 "When did unsystematic monetary policy have an effect on inflation?" by B. Mojon, December 2005.
- 560 "The determinants of 'domestic' original sin in emerging market economies" by A. Mehl and J. Reynaud, December 2005.
- 561 "Price setting in German manufacturing: new evidence from new survey data" by H. Stahl, December 2005.
- 562 "The price setting behaviour of Portuguese firms: evidence from survey data" by F. Martins, December 2005.
- 563 "Sticky prices in the euro area: a summary of new micro evidence" by L. J. Álvarez, E. Dhyne, M. M. Hoeberichts, C. Kwapil, H. Le Bihan, P. Lünnemann, F. Martins, R. Sabbatini, H. Stahl, P. Vermeulen and J. Vilmunen, December 2005.
- 564 "Forecasting the central bank's inflation objective is a good rule of thumb" by M. Diron and B. Mojon, December 2005.
- 565 "The timing of central bank communication" by M. Ehrmann and M. Fratzscher, December 2005.
- 566 "Real versus financial frictions to capital investment" by N. Bayraktar, P. Sakellaris and P. Vermeulen, December 2005.
- 567 "Is time ripe for a currency union in emerging East Asia? The role of monetary stabilisation" by M. Sánchez, December 2005.



- 568 "Exploring the international linkages of the euro area: a global VAR analysis" by S. Dées,F. di Mauro, M. H. Pesaran and L. V. Smith, December 2005.
- 569 "Towards European monetary integration: the evolution of currency risk premium as a measure for monetary convergence prior to the implementation of currency unions" by F. González and S. Launonen, December 2005.
- 570 "Household debt sustainability: what explains household non-performing loans? An empirical analysis" by L. Rinaldi and A. Sanchis-Arellano, January 2006.
- 571 "Are emerging market currency crises predictable? A test" by T. A. Peltonen, January 2006.
- 572 "Information, habits, and consumption behavior: evidence from micro data" by M. Kuismanen and L. Pistaferri, January 2006.
- 573 "Credit chains and the propagation of financial distress" by F. Boissay, January 2006.
- 574 "Inflation convergence and divergence within the European Monetary Union" by F. Busetti, L. Forni, A. Harvey and F. Venditti, January 2006.
- 575 "Growth in euro area labour quality" by G. Schwerdt and J. Turunen, January 2006.
- 576 "Debt stabilizing fiscal rules" by P. Michel, L. von Thadden and J.-P. Vidal, January 2006.
- 577 "Distortionary taxation, debt, and the price level" by A. Schabert and L. von Thadden, January 2006.
- 578 "Forecasting ECB monetary policy: accuracy is (still) a matter of geography" by H. Berger, M. Ehrmann and M. Fratzscher, January 2006.
- 579 "A disaggregated framework for the analysis of structural developments in public finances" by J. Kremer, C. Rodrigues Braz, T. Brosens, G. Langenus, S. Momigliano and M. Spolander, January 2006.
- 580 "Bank interest rate pass-through in the euro area: a cross country comparison" by C. K. Sørensen and T. Werner, January 2006.
- 581 "Public sector efficiency for new EU Member States and emerging markets" by A. Afonso,L. Schuknecht and V. Tanzi, January 2006.
- 582 "What accounts for the changes in U.S. fiscal policy transmission?" by F. O. Bilbiie, A. Meier and G. J. Müller, January 2006.
- 583 "Back to square one: identification issues in DSGE models" by F. Canova and L. Sala, January 2006.
- 584 "A new theory of forecasting" by S. Manganelli, January 2006.

