ARTICLE IN PRESS

Journal of Econometrics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Theory-coherent forecasting[☆]

Raffaella Giacomini[a], Giuseppe Ragusa[b,*]

[a] Department of Economics, University College London, Gower Street, London WC1E6BT, UK

[b] Department of Economics and Finance, Luiss University, Viale Romania 32, 00197 Rome, Italy

ARTICLE INFO

Article history:
Available online xxxx

JEL classification:
C14
C22
C53

Keywords:
Density forecast
Forecast evaluation
Semiparametric projections
Exponential titling

ABSTRACT

We consider a method for producing multivariate density forecasts that satisfy moment restrictions implied by economic theory, such as Euler conditions. The method starts from a base forecast that might not satisfy the theoretical restrictions and forces it to satisfy the moment conditions using exponential tilting. Although exponential tilting has been considered before in a Bayesian context (Robertson et al. 2005), our main contributions are: (1) to adapt the method to a classical inferential context with out-of-sample evaluation objectives and parameter estimation uncertainty; and (2) to formally discuss the conditions under which the method delivers improvements in forecast accuracy. An empirical illustration which incorporates Euler conditions into forecasts produced by Bayesian vector autoregressions shows that the improvements in accuracy can be sizable and significant.

1. Introduction

Economic theory often provides moment conditions which restrict the dynamic behavior of key macroeconomic variables, but which cannot be used directly to produce forecasts that are coherent with theory. An expectational Euler condition, for example, imposes a nonlinear restriction on the joint density of future consumption and real interest rates, conditional on current observables, and thus potentially provides valuable information for forecasting the future path of both variables. This paper considers a method for constructing multivariate density forecasts that are coherent with theory. The idea is to start from a base multivariate density forecast which does not necessarily satisfy the theoretical restrictions and then use exponential tilting to obtain a new density forecast that by construction satisfies the moment conditions. In practice, the method consists of an importance sampling procedure where draws from the base density forecasts are reweighted using weights obtained by numerical optimization. The computational cost of the procedure is low, only depending on the number of theoretical restrictions one wishes to incorporate.

Even though exponential tilting has been considered before in microeconometrics and Bayesian econometrics (in particular by Robertson et al., 2005), the contribution of this paper is to formalize the method in a classical inferential context with out-of-sample evaluation objectives, address the issue of parameter estimation uncertainty and present formal conditions under which the incorporation of theoretical restrictions leads to forecast accuracy gains. Our main result shows that a tilted density forecast which incorporates moment conditions that are true in population but may depend on consistently estimated parameters is more accurate than the base density forecast, provided that accuracy is measured by the logarithmic scoring rule of Amisano and Giacomini (2007). One practical implication of this result is the recommendation to separately estimate the parameters in the base density and the parameters in the moment condition.

The paper offers a way to incorporate economic theory into forecasting without resorting to estimation of full-fledged Dynamic Stochastic General Equilibrium (DSGE) models, and is part of a small literature which proposes hybrid approaches that combine elements of economic theory and reduced-form modeling. There are several reasons why a hybrid approach might be an appealing alternative to forecasting with fully specified theoretical models. First, a DSGE model in general does not directly provide a conditional density that can be used for forecasting, which means that one typically would end up forecasting with an approximated density (either from the linearized model or, less frequently, from higher order approximations or numerical solutions of the model). A few articles have investigated the forecast performance of linearized DSGE models compared to reduced-form models (e.g., Smets and Wouters, 2003; Edge et al., 2010; Christoffel et al., 2010), but it remains to be seen whether the results are robust to different choices of DSGE priors, different time periods

☆ We thank the Domenico Giannone, Norman Swanson, three anonymous referees and seminar participants at Princeton, Columbia, Institute for Advanced Studies at Vienna, University of Bath, Université libre de Bruxelles for useful comments.

* Corresponding author. Tel.: +39 06 8522 5640; fax: +39 06 8522 5949.
E-mail addresses: r.giacomini@ucl.ac.uk (R. Giacomini), gragusa@luiss.it (G. Ragusa).

http://dx.doi.org/10.1016/j.jeconom.2014.04.014
0304-4076/© 2014 Published by Elsevier B.V.

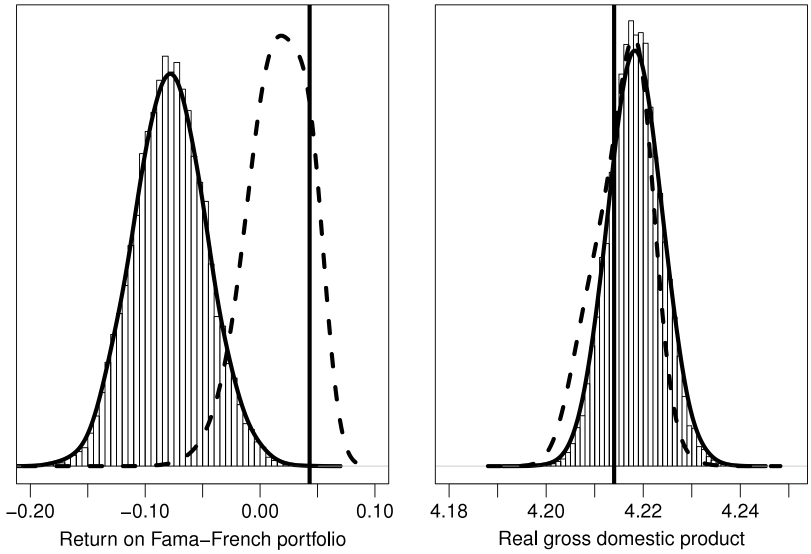

Fig. 1. The figure shows two out-of-sample density forecasts for the real GDP and for the return on a Fama–French portfolio at one particular point in time (1988:Q1): the histogram is the one-step-ahead density forecast implied by a BVAR with 22 variables which include real GDP, non-durables and services real consumption, the federal funds rate and the return on the Fama–French portfolio. In each graph, the dashed line is the projected density forecast that incorporates the Euler equation; the solid vertical line is the realization of the variable.

and different specifications for the reduced form model, to name a few. Second, the user has to take a stand on many aspects of the model for which theory provides no guidance, necessitating ad hoc modeling choices. Third, DSGE models are not capable of incorporating the rich datasets that have proved helpful in reduced-form forecasting of variables such as inflation and real output. Examples of "hybrid" approaches have been considered in the context of linearized DSGE models are Schorfheide (2000), who proposes priors based on the DSGE model to perform Bayesian inference in the VAR model, and Del Negro and Schorfheide (2004) and Carriero and Giacomini (2011), who consider combinations of the DSGE and the VAR model. Even though the focus of these methods is on estimation, they could in principle be used for forecasting. The method considered in this paper has several advantages over these existing hybrid approaches as it allows full flexibility in the choice of the base model, which can be driven by considerations about its empirical performance instead of the requirement that the base model contain the same variables as the DSGE model. Further, it does not require one to put equal faith in all of the restrictions embedded in a DSGE model, but to choose which restrictions to impose. Finally, the restrictions can be nonlinear, whereas the approaches mentioned above necessitate restricting attention to linearized models.

A noteworthy feature of the tilting procedure is that it yields a new multivariate forecast which has a known analytical form but in general is not a member of a known family of distributions (for example, even if the base density is a multivariate normal, the tilted density will not be normal in all but a few special cases). The fact that the method gives a density that is not in the same family as the original density forecast is a useful feature of the approach, as it allows one to understand the effects of imposing theoretical restrictions on the entire shape of the distribution, including the marginal densities and the dependence structure. To illustrate how the tilting method can modify the marginal density forecast and why this could result in accuracy improvements for the individual variables, we show in Fig. 1 an actual example from our empirical application. The application considers as a base model a Bayesian VAR with 22 variables including real consumption (C_t), the real return on the Fama and French (1993) portfolio R_t and real GDP, and incorporates the Euler condition $E_t\left[\beta\left(\frac{C_{t+1}}{C_t}\right)^{-\alpha}R_{t+1}-1\right]=0$, with $\beta=0.999$ and $\alpha=0.6$.

The left panel of Fig. 1 shows the effect of tilting at a particular point in time on the density forecast for R_t, a variable which is directly restricted by the Euler equation, whereas the right panel shows that the tilting also affects the density forecast for real GDP, which does not enter the Euler condition directly but is nonetheless indirectly modified by the procedure along with all other variables in the base model. The histograms represent the density forecasts implied by the base model and the dashed lines are the tilted density forecasts. The vertical lines are the realizations of the variables. In both cases the incorporation of the Euler equation restrictions modifies the shape of the density forecasts implied by the base model by making them left skewed and by shifting them towards the actual realizations of the variables, thus yielding more accurate point and density forecasts for both variables. Even though the figure only illustrates the benefits of tilting for asset returns at one point in time, we show in Section 4 that this tends to be true on average.

2. Motivating example

This section shows a simple example where an analytical expression for the tilted density can be easily obtained, which provides some intuition for the method. Suppose that the true conditional density $h_t(y_{t+1})$ of the variable of interest Y_{t+1} is unknown apart from its conditional mean μ_t, which implies the moment condition:

$$E_t\left[Y_{t+1}-\mu_t\right]=\int\left(y_{t+1}-\mu_t\right)h_t(y_{t+1})dy_{t+1}=0.$$

Suppose that one has available a one-step-ahead density forecast $f_t(y)\sim N(\widehat{\mu}_t,1)$, which does not necessarily satisfy the moment conditions in that $\widehat{\mu}_t$ may be different from μ_t. In order to obtain a new density forecast which by construction have the correct mean, the tilting procedure finds a new density forecast $\widetilde{f}_t(y)$ with mean μ_t and which is closest to $f_t(y)$ according to the Kullback–Leibler measure of divergence. The solution of this constrained optimization can be shown to be

$$\begin{aligned}\widetilde{f}_t(y)&=f_t(y)\exp\left\{\eta_t+\tau_t(y-\mu_t)\right\} &(1)\\ &=\frac{1}{\sqrt{2\pi}}\exp\left\{-\frac{1}{2}(y-\widehat{\mu}_t)^2+\eta_t+\tau_t(y-\mu_t)\right\}.\end{aligned}$$

Please cite this article in press as: Giacomini, R., Ragusa, G., Theory-coherent forecasting. Journal of Econometrics (2014), http://dx.doi.org/10.1016/j.jeconom.2014.04.014