Journal of Financial Economics 000 (2017) 1–19

Contents lists available at ScienceDirect

Journal of Financial Economics

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

The advantages of using excess returns to model the term structure☆

Adam Goliński*, Peter Spencer

Department of Economics and Related Studies, Heslington, University of York, York Y010 5DD, United Kingdom

ARTICLE INFO

Article history:
Received 15 October 2012
Revised 7 August 2015
Accepted 7 April 2016
Available online xxx

JEL Classification:
G12
C58

Keywords:
Term structure
Excess return framework
Long memory

ABSTRACT

We advocate the use of excess returns rather than yields or log prices in analysing the risk neutral dynamics of the term structure. We show that under standard assumptions, excess returns are affine in the risk neutral innovations in the factors. This framework has several important advantages. First, it allows for an easy estimation of models that are more flexible than the *AR*(1). Indeed, we estimate models with more general dynamics, like *ARFIMA*(*p*, *d*, *q*), almost as easily as *AR*(1). Second, within our framework the dimension of the unrestricted model is the same for the *AR*(1) as it is for the richer models, and does not expand in line with the state vector as it does in a yield or log price framework. This makes it appropriate to test all of these risk neutral dynamic specifications against the same *OLS* unrestricted alternative. Our results for the US Treasury bond market show that the unrestricted model is preferred to the *AR*(1) by the Bayesian Information Criterion, but the opposite conclusion is reached for more flexible models. A final advantage of the excess returns framework is that the pricing errors are much lower than for the equivalent log price system.

1. Introduction

This paper proposes a novel approach to the term structure of interest rates, which allows the cross section of returns to be modelled easily without necessarily adopting any specific model of the risk-neutral dynamics, other than to assume them to be linear. We exploit the fact that absent arbitrage and measurement error, the forward price of any security for any future period is its risk-neutral expectation for that period. Thus, it can be used to replace the risk-neutral expectation generated by an autoregressive model in the standard specification of the cross section, allowing less restrictive dynamic specifications to be used. It follows that the excess logarithmic return on any security can be viewed as a risk-neutral surprise or innovation plus a Jensen term determined by its volatility. Assuming that these excess returns have a factor structure, this approach yields an affine model that relates M excess bond returns that are measured with error to $K < M$ factor *innovations*, which can be represented by excess returns on bond portfolios that are assumed to be measured without error. The loadings of the excess returns on the factor innovations depend upon the risk-neutral factor dynamics, allowing these to be identified and estimated.

In the absence of no-arbitrage restrictions, this model can be estimated simply by regressing the M excess bond returns on the K factor innovations using Ordinary Least

☆ We are grateful to Karim Abadir, Richard Baillie, John Campbell, Laura Coroneo, Antonio Diaz Perez, Fabrizio Iacone, Francisco Jareño, Scott Joslin, Peter Phillips, Yongcheol Shin, Rob Taylor, Mike Wickens and to participants at the 2013 EFMA and WFC conferences for helpful comments. We are also indebted to a referee for exceptionally perceptive and helpful suggestions.

* Corresponding author.
E-mail addresses: adam.golinski@york.ac.uk (A. Goliński), peter.spencer@york.ac.uk (P. Spencer).

http://dx.doi.org/10.1016/j.jfineco.2017.05.001

Squares (*OLS*).[1] This $M \times K$ model can be restricted by making appropriate assumptions about arbitrage opportunities and the risk-neutral dynamics. We show that the excess return framework easily handles *ARFIMA* processes, which include the various autoregressive (*AR*), moving average (*MA*), and long memory (*LM*) models that have been used in the literature. *MA* and *LM* models are notoriously difficult to estimate, but are easily handled in the excess returns framework.[2] The unrestricted *OLS* excess returns model provides a convenient encompassing model for testing such restrictions and in particular, the first-order autoregressive *AR*(1) arbitrage-free specification used extensively in the term structure literature.

As (Bams and Schotman, 2003) and others have noted, a variant of the excess returns model can be obtained simply by forward-differencing the standard *AR* (1) affine term structure model (*ATSM*), which relates the *M* log prices or yields to the values of the *K* factor portfolios.[3] However, we show that this forward-difference relation only holds for the *AR*(1) specification; richer dynamic specifications increase the dimensions of the yield *ATSM* but not the return *ATSM*. For example, if the risk-neutral dynamics are second-order autoregressive (*AR*(2)), this doubles the number of lagged state variables in the risk-neutral dynamics and means that the *M* yields are not properly explained by the *K* factor portfolios; another *K* factors are in principle needed to do this.[4] However, working with innovations removes the second- and higher-order lags (and, in the case of *MA* models, lagged error terms) from the risk-neutral dynamics and preserves the $M \times K$ structure of the excess returns model.

This is the key message of this paper: the return or forward-difference format provides a very convenient framework for testing risk-neutral models of the factor dynamics because it is not necessary to restrict attention to the *AR*(1) model analysed by Bams and Schotman (2003), Bauer (2017), and Adrian et al. (2013), and others who use returns data. Moreover, we find that empirically, the measurement errors in the return framework are much smaller than in the corresponding yield model (see Fig. 5). Also, as found by Dai and Singleton (2000); Duffee (2011a), and (Adrian et al., 2013), the yield pricing errors exhibit a large degree of autocorrelation, which disappears when working with returns (see Table 3). These results suggest that the latter provides a more powerful framework for asset pricing and model testing. We use this framework to re-assess the performance of *AR*(1) and other *ATSMs*.

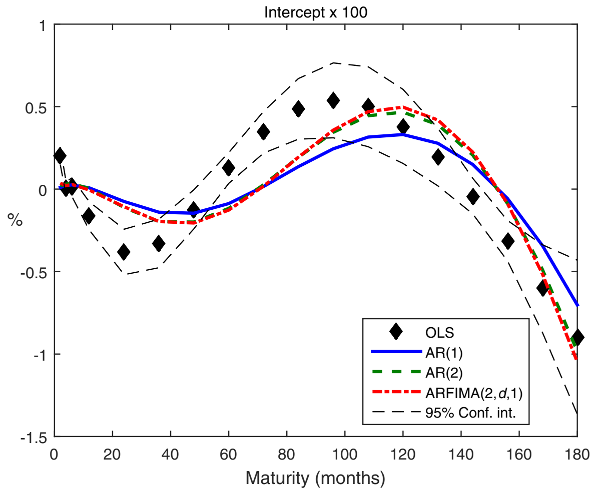

Fig. 1. Intercept coefficients by maturity of the *AR*(1), *AR*(2), and *ARFIMA*(2, *d*, 1) excess return *DTSM* models and the unrestricted *OLS* benchmark. The dashed lines represent the 95% *OLS* confidence interval. The sample period is January 1983 to December 2011.

Several researchers have noted that the factor loadings estimated for an *AR*(1) yield model are typically very close to those of an unrestricted *OLS* regression and that the measurement errors are only a little larger. Hamilton and Wu (2014) and (Duffee, 2011a) compare these specifications and show that the *AR*(1) restrictions are nevertheless strongly rejected statistically: the differences between the *AR*(1) and *OLS* benchmark are numerically small but statistically significant. Balduzzi and Chiang (2012) also find strong evidence against the *AR*(1) *ATSM* when testing it against the *OLS*. These findings reflect the fact that the pricing errors in the Treasury bond market are very small, which means, as Cochrane and Piazzesi (2008) observe, that the risk-neutral model parameters are very precisely determined, making data for this market very good at comparing the performance of rival models.

We confirm the rejection of the *AR*(1) restrictions in our data set but, in a more constructive way, we show that richer dynamic models with parameters that are constrained by asset pricing theory can beat the unrestricted benchmark model. Fig. 2 shows the return loadings for the *AR*(1) model alongside the unconstrained *OLS* return loadings and their 95 % confidence interval. As in previous models, the fit looks reasonable, but many of the *AR*(1) loadings on the first two factors lie outside the confidence interval. More worrying, Fig. 1 shows that the intercept coefficients of the *AR*(1) model are visibly and significantly flatter than those of the *OLS* benchmark.[5] However, the more flexible models do a better job, generating loadings that largely coincide with the *OLS* estimates. Reflecting these observations, on the basis of the Bayesian

[1] We show that if the factor dynamics have unrestricted linear representation (allowing them to be represented as an infinite series of current and lagged risk neutral innovations) and the no-arbitrage restriction that follows from the factor structure is neglected, this generates the $M \times K$ unrestricted benchmark *OLS* model of the cross section of returns. In contrast, it seems hard to interpret the unrestricted *OLS* yield level benchmark as a model of the factor dynamics.

[2] Goliński and Zaffaroni (2016) model yields with long memory dynamics using infinite-dimensional state vectors and estimate the system using a state-space truncation.

[3] Thus, the factor loadings matrices of these systems should in principle be identical.

[4] Duffee (2011b) and (Adrian et al., 2013) have suggested that the performance of the *AR*(1) model might be improved by introducing more factors, but this would also increase the stochastic dimension of the cross section (see Section 3.2.1).

[5] Hamilton and Wu (2014) also find that the badly fitting intercept is the main reason that the *AR*(1) model is rejected against the *OLS* model.