Research Statement

Drew D. Creal

January 2018

Introduction

My research focuses on estimation, inference, and applications of state space and time-varying parameter models in economics. State space models are widely used throughout macroeconomics and finance, as many dynamic economic models are special cases of this framework. This includes both structural models with optimizing agents and reduced form models. Popular models within this category are, for example, dynamic stochastic generalized equilibrium models, generalized autoregressive conditional heteroskedasticity (GARCH) models, stochastic volatility models, and stochastic intensity models. My research has three main themes: (i) developing new observation-driven models and their properties for applications in macroeconomics and finance with an emphasis on credit risk; (ii) improving estimation methods for specific classes of state space models that are particularly important to economists; and (iii) solving the high-dimensional integrals required to estimate parameter-driven state space models.

Time-varying parameter models are often separated into two broad classes called observation-driven models and parameter-driven (state space) models, which I define more formally below. Econometrically, the two classes of models are different because estimation of parameter-driven models typically requires solving a high-dimensional integral while estimation of observation-driven models does not. My research aims to make contributions to both of these areas.

In my work on observation-driven models, I have developed with my co-authors a new class of time series models called generalized autoregressive score models; see Creal, Koopman, and Lucas (2013). Generalized autoregressive score models offer a unified framework for understanding how many existing observation-driven models are related to one another. More importantly, researchers can use the framework to develop new models for new applications. For example, I have used this framework to create new models for dynamic copulas, time-varying covariance matrices, ordered logit models with time-varying probabilities for each category, and marked point processes with stochastic intensity functions. In subsequent work, my co-authors and I have extended the generalized autoregressive score family of models by introducing the idea of observation-driven mixed measurement models with common factors. The new collection of models have many applications in economics but a primary motivation of this research has been to provide new, improved models for the credit risk literature.

Parameter-driven (state space) models are the other sub-class of time-varying parameter models. The goal of my research in this area is to improve methods of estimation for specific classes of models that are important to economists. Recently, I have worked on the class of affine models, which are the workhorse model in the literature on the term structure of interest rates and which are becoming increasingly more important in credit risk. In Creal (2017), I provide a solution for the high-dimensional integrals needed to estimate several of the models within the affine family that have the dynamics for their latent variables specified as a Cox, Ingersoll, and Ross (1985) process. I solve most of the integrals analytically and demonstrate that it is possible to calculate the likelihood function exactly as well as estimates of the latent (state) variable. The paper illustrates the methods on an affine stochastic volatility model. However, the results in the paper can be extended to a larger class of models whose state variables follow the Cox, Ingersoll, and Ross (1985) process including stochastic intensity, stochastic duration, and stochastic transition models.

Finally, I will also discuss other work related to the term structure of interest rates. In collaboration with Cynthia Wu, I developed improved estimation methods for discrete-time, non-Gaussian affine term structure models. Our approach generalizes some of the recent results in the literature on Gaussian affine term structure models, which are used in both the macroeconomics and finance literatures. Our work is motivated by both the fact that interest rates are near their zero-lower bound and also the desire to model time-varying volatility in interest rates.

Observation-driven models: generalized autoregressive score models

Observation-driven time series models are popular in economics because they offer a simple method for building models that can better fit the observed data yet are reasonably easy to estimate. Leading examples of observation driven models include the generalized autoregressive conditional heteroskedasticity (GARCH) model of Engle (1982)/Bollerslev (1986) as well as the autoregressive conditional duration (ACD) model of Engle and Russell (1998). Almost all observation-driven time-varying parameter models found in economics have the following structure

$$y_t \sim p\left(y_t|x_t;\theta\right), \tag{1}$$

$$x_{t+1} = \mu + \Phi x_t + \Sigma s_t, \tag{2}$$

where y_t is an observed time series variable and x_t is the time-varying parameter. The conditional distribution of y_t given x_t is $p\left(y_t|x_t;\theta\right)$ and $s_t = s\left(y_1,\ldots,y_t,x_t;\theta\right)$ is a deterministic function of past observations; both of which may depend on a vector of unknown parameters θ. The recursion for the time-varying parameter x_t has the familiar structure of an autoregressive moving average.

In an observation-driven model, the time-varying parameter gets updated deterministically through time based upon past values of the data as in (2).[1] Conditional on the model's parameters θ and an initial value for x_0, the time-varying parameter is observable by construction. Consequently, it is possible to calculate the likelihood function as the product of the conditional distributions in (1). This makes estimation of observation-driven models relatively easier than the sub-class of parameter-driven models and this is why observation-driven models have historically been widely used.

The contribution of Creal, Koopman, and Lucas (2013) is the choice of a specific function $s_t = s\left(y_1,\ldots,y_t,x_t;\theta\right)$ that determines how x_t evolves through time. We define s_t to be the scaled score function

$$s_t = S_t \nabla_t \qquad \nabla_t = \frac{\partial \log p\left(y_t|x_t;\theta\right)}{\partial x_t}$$

where ∇_t is the score of the observation density (1) and S_t is a scaling matrix, for example Fisher's information matrix $S_t = \mathrm{E}\left[\nabla_t \nabla_t'\right]^{-1}$. Due to the reliance of the transition equation (2) on the score function, we call this class of models the generalized autoregressive score model.

[1] This contrasts with the class of parameter-driven (state space) models where x_t is random variable drawn from its own probability distribution.

Interestingly, for different choices of the observation density $p(y_t|x_t;\theta)$ and scaling matrix S_t, the generalized autoregressive score framework nests well-known models in the literature. The class of generalized autoregressive score models includes as special cases the generalized autoregressive conditional heteroskedasticity (GARCH) model of Engle (1982)/Bollerslev (1986), the autoregressive conditional duration (ACD) model of Engle and Russell (1998), the autoregressive conditional multinomial (ACM) model of Russell and Engle (2005), the time-varying Bernoulli model of Cox (1958)/Shephard (1995), the multiplicative error model (MEM) of Engle and Gallo (2006), the time-varying Poisson regression model of Davis, Dunsmuir, and Streett (2003), as well as others.

More importantly, the generalized autoregressive score framework leads to new models. For example, our paper introduced a new methodology for creating dynamic copulas. Copulas are a statistical method for modeling the comovements (e.g. correlations or tail behavior) among a group of random variables separately from their marginal distributions. Copula models are becoming increasingly important in modern risk management. They are used by financial institutions to model the assets in a portfolio such as a loan book or a collateralized debt obligation. In the aftermath of the financial crisis, the copula models used by practitioners over the decade from 2000 to 2010 have been heavily criticized for good reasons. When the economy enters a crisis, the dependence structure among assets can change dramatically. Unfortunately, the existing copula models used by many financial institutions were overly simplistic and did not allow the dependence structure among assets to change over time. Making copulas dynamic however is not a simple task as it is difficult to understand what functions of the data should determine time-variation in a copulas' parameters. The generalized autoregressive score framework makes creating dynamic copulas easier. Our work contributes to the literature on time-varying copulas which is an active area of research in financial econometrics; see, e.g. Patton (2009). In Creal, Koopman, and Lucas (2013), we created several new time-varying copulas models from static models where there are no existing alternatives in the literature. Also, for the observation-driven copula models that did have an existing alternative (see e.g. Patton (2006)), we compared our new models to the existing models, and demonstrated a marked improvement in terms of fit and performance.

In Creal, Koopman, and Lucas (2011), we furthered the ideas of time-varying dependence by considering multivariate generalized autoregressive score models for time-varying covariance matrices. Multivariate extensions of the GARCH family of models have been developed by, among others, Bollerslev (1990) and Engle and Kroner (1995), with the most popular version being the dynamic conditional correlation (DCC) model of Engle (2002). As the generalized autoregressive score model for the univariate normal distribution with time-varying variance is equivalent to the GARCH model, one would correctly expect that estimates produced by the DCC model and the generalized autoregressive score model from the multivariate normal distribution are nearly identical. This is due to the fact that the recursion for the time-varying parameters in the DCC model is approximately the score of the multivariate normal distribution. However, tails of the conditional distribution of returns are well known to be heavier than the normal. Therefore, in Creal, Koopman, and Lucas (2011) and Zhang, Creal, Koopman, and Lucas (2012), we consider a wider class of distributions that offer more flexibility in capturing the data yet are essentially as simple to estimate as the DCC model. The estimates of the conditional volatilities and correlations produced by the new models can be substantially different from the existing models in the literature. This is because the score function from heavier tailed conditional distributions is less sensitive to outliers and large idiosyncratic movements in asset prices. The paper illustrated the improved performance of the new models over the DCC model on a dataset of equity returns and also in a simulation study. I note that Professor Rob Engle of New York University has recently implemented the univariate version of the Student's t generalized autoregressive score volatility model in the Volatility Laboratory (V-LAB), which provides real-time estimates of the volatility of a large cross-section of U.S. equities.

The time between the random arrival of an event (e.g. stock trades, news announcements, credit ratings transitions) is often of economic interest. A point process is a popular statistical model for the random arrival of an event, where the probability of the next event occurring over a short period of time is called the intensity. Conditional on an event occurring, the realized value of the event may take one of a finite number of outcomes, which is known as a "mark". In the paper introducing generalized autoregressive score models, another contribution was a new model for marked point processes with time-varying intensities. In our application to credit risk, the set of different random events (marks) are the possible credit ratings transitions of a group of firms. One of the possible transitions is default of the firm. The model captures the existence of a common business cycle or "frailty" factor that drives default probabilities higher than what is possible with only observable covariates such as accounting variables. We illustrated that the model performs comparably with other models in the credit risk literature such as Duffie, Eckner, Horel, and Saita (2009) but is substantially easier to estimate.

We extended the generalized autoregressive score framework further in Creal, Schwaab, Koopman, and Lucas (2014) by introducing the idea of mixed measurement panel data models where data of many different types (discrete, continuous, and ordered categorical) may depend on a set of common time-varying parameters due to their shared dependence on the business cycle. Our methodology was used to jointly model the credit ratings transitions for a panel of over 7000 firms, over 1000 observed values of the loss given default, and a panel of macroeconomic variables. Upon estimation of the model, it can be used to forecast credit risk conditions in the economy and to construct predictive loss distributions for portfolios of corporate bonds at different forecasting horizons. The paper also provides a methodology for conducting stress tests of credit portfolios from which one can determine capital requirements using the high percentiles of the simulated portfolio loss distributions. Our modeling framework provides a relatively simple observation-driven alternative to the parameter-driven frailty models of McNeil and Wendin (2007), Koopman, Lucas, and Monteiro (2008), and Duffie, Eckner, Horel, and Saita (2009). An important part of the contribution is that our modeling framework allows for the identification of three components of credit risk simultaneously (macroeconomic risk, rating migration/default risk, and loss given default risk), whereas other models only concentrate on defaults, defaults and ratings, or defaults and macroeconomic risk. Our results illustrate that the tails of the loss distributions grow substantially after accounting simultaneously for these sources of uncertainty.

In addition to introducing the autoregressive conditional heteroskedasticity (ARCH) model, Engle (1982) also developed a specification test for the null hypothesis that the errors in a linear model are conditionally homoskedastic, that is there are no ARCH effects. His Lagrange multiplier (LM) test is widely used and is standard output for linear regression models in many software programs. As the ARCH model is a special case of the generalized autoregressive score model, there naturally exists an LM test that generalizes Engle's ARCH test to other parametric families of distributions. In Calvori, Creal, Koopman, and Lucas (2017), we introduce the GAS-LM test and compare it to other structural breaks tests found in the literature across a range of models. Results from Monte Carlo experiments in the paper show that the test performs as one would expect when compared to a likelihood ratio (LR) or Wald test under the null hypothesis that the parameters of a model are constant. When the null hypothesis is false and the proposed alternative is close to the true DGP, the LR and Wald tests are more powerful. However, the generalized autoregressive score LM test is potentially more robust as the other tests can be either sensitive to the alternative chosen by the user or to the tuning parameters used to conduct the test.

Parameter-driven models: general state space models

In the second part of my research program, I am interested in applications and the methodological development of the class of parameter-driven (state space) models. A general state space model for an observed time series y_t has the form

$$\begin{aligned} y_t &\sim p\left(y_t|x_t;\theta\right), \\ x_t &\sim p\left(x_t|x_{t-1};\theta\right), \end{aligned}$$

where x_t is the state variable (time-varying parameter) and $p\left(y_t|x_t;\theta\right)$ remains the conditional distribution of y_t given x_t. The difference between an observation-driven model and a parameter-driven model is how the state variable evolves through time. Instead of defining the evolution of the state variable x_t as a deterministic function of past observations, the state variable is now a realization from the conditional probability distribution $p\left(x_t|x_{t-1};\theta\right)$. In statistical terms, the state variable is latent and the likelihood function is in the form of a high-dimensional integral

$$p(y_1 \ldots, y_T;\theta) = \int \ldots \int p(y_1, \ldots, y_T|x_0, x_1 \ldots, x_T;\theta) p(x_0, x_1, \ldots, x_T;\theta) \mathrm{d}x_0, \ldots \mathrm{d}x_T$$

This makes estimation of the model more challenging as the integrals have known solutions in a limited number of settings; e.g. linear, Gaussian state spaces models (Kalman (1960)) and finite-state Markov switching models (Baum and Petrie (1966)/Hamilton (1989)).

My dissertation at the University of Washington focused on a set of statistical tools called sequential Monte Carlo methods or particle filters. Particle filters are Monte Carlo methods designed to solve the high-dimensional integrals that are encountered when estimating state space models. The early work from my dissertation led to several papers applying these methods to the estimation of the business cycle (Creal, Koopman, and Zivot (2010)), models for forecasting inflation (Creal (2012)), and non-Gaussian continuous-time models for conditional heteroskedasticity (Creal (2008)).

In my most recent papers, I investigated the class of affine models, which enjoy widespread use in macroeconomics and finance. These models lead to closed-form solutions for bond and options prices. In Creal (2017), I demonstrate how to calculate the likelihood function (solve the high-dimensional integral) exactly for a class of stochastic volatility models whose latent variance follows a Cox, Ingersoll, and Ross (1985) process. In this model, the state variable is actually two-dimensional with one variable that is continuous (the time-varying variance) and another variable that is discrete. The key insight is to integrate out the continuous-valued latent variance analytically leaving only the latent discrete variable whose support is over the non-negative integers. The remaining discrete variable can be integrated out by reformulating the original model as a Markov-switching model, which are well-known in econometrics (Hamilton (1989)). The methods developed from this paper can be extended to a larger class of affine models including stochastic intensity models, which are a key part of the literature on default risk.

Related to this work, I have written several papers with Cynthia Wu on affine term structure models. Affine term structure models are the benchmark models central banks use to conduct monetary policy. The current generation of models used at central banks are Gaussian models, which offer a reasonable means of capturing conditional first moments. The Great Recession and the ensuing zero lower bound unfortunately require non-Gaussian models. In Creal and Wu (2015), we developed a new approach for calculating the maximum likelihood estimator for the class of discrete-time, non-Gaussian affine models. The approach to estimation that we have developed generalizes the recent approaches for Gaussian affine term structure models of Joslin, Singleton, and Zhu (2011) and Hamilton and Wu (2012).

One of the main empirical lessons from Creal and Wu (2015) is that stochastic volatility models with spanned factors do not fit the term structure well nor do the class of unspanned stochastic volatility (USV) models that are popular in the finance literature. USV models impose cross-sectional restrictions that prevent the models from fitting the conditional mean of yields. In Creal and Wu (2017), we propose an alternative class of affine models that allows for unspanned volatility factors that do not impose the cross-sectional restrictions present in USV models. We demonstrate that our model can fit both the conditional mean and conditional volatility of yields.

Creal and Wu (2017) also made a second contribution to the growing literature on whether uncertainty shocks influence the real economy; see , e.g. Bloom (2009), Bloom (2014), and Jurado, Ludvigson, and Ng (2015). The early papers in this literature all use an observable variable as a proxy for uncertainty and then place this variable into a vector autoregression to study the impact of an uncertainty shock on real variables. In contrast, our model allows the conditional mean of yields and macroeconomic variables to be a function of the latent stochastic volatility. Consequently, uncertainty is internally consistent and is estimated with the observed data. In our paper, we studied the impact that shocks to the stochastic volatility of interest rates and the real economy have on real variables.

Finally, I would like to take the opportunity to discuss other papers that I co-wrote with colleagues at Chicago Booth. These papers apply tools from the state space modeling literature to solve practical economic problems. In Creal, Gramacy, and Tsay (2014), we develop a statistical methodology that creates credit ratings for a large database of both public and private companies directly from the firms' asset prices and their other observable characteristics. Even after the recent financial crisis when the ratings of independent agencies (Moody's and Standard & Poors) were heavily criticized, credit ratings from these agencies are still used in risk management departments of financial institutions. Our methodology is intended to be an alternative, objective measure. From a statistical perspective, we use simple methods that are reliable so that the ratings are transparent and are directly related to the data. The ratings methodology leads to categories that have an economic meaning; something which is sorely lacking in the ratings of the independent agencies. Our methodology also allows firms that are not actively traded to be rated by matching traded firms to non-traded firms based on their observable characteristics. We compare our ratings to the ratings of the independent agencies and find that they lead the agencies' ratings by a considerable margin of time.

In Creal and Tsay (2015), we develop a class of flexible stochastic copula models that are parameterized so that estimation and inference remains feasible even when the dimension of the data is high. High dimensional data introduces two problems for econometricians. First, the number of parameters to be estimated may exceed the number of data points and second the computational burden of estimation may explode as the amount of data increases. Our copula models solve both of these problems by introducing a sensible factor structure that imposes shrinkage to reduce the number of parameters while still allowing the dependence between observable random variables (e.g. correlations) to change flexibly through time. In our empirical application, we model the joint dependence between a panel of 100 firm's equity and debt making for a total cross-section size of 200 observations. To get a sense of how high the dimensions are, if we modeled this data with a static Gaussian copula there would be a total of 19900 parameters to be estimated in the correlation matrix. The number of parameters vastly exceeds the number of data points making estimation infeasible unless we impose some form of shrinkage. Despite these challenges, we illustrate how our copula models have sufficient structure to make computation feasible even in high dimensions.

References

Baum, L. E. and T. P. Petrie (1966). Statistical inference for probabilistic functions of finite state Markov chains. *The Annals of Mathematical Statistics 37*(1), 1554–1563.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica 77*(3), 623–685.

Bloom, N. (2014). Fluctuations in uncertainty. *Journal of Economic Perspectives 28*(2), 153–176.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics 31*(3), 307–327.

Bollerslev, T. (1990). Modeling the coherence in short-run nominal exchange-rates: a multivariate generalized ARCH model. *The Review of Economics and Statistics 72*, 498–505.

Calvori, F., D. Creal, S. J. Koopman, and A. Lucas (2017). Testing for parameter instability across different modelling frameworks. *Journal of Financial Econometrics 15*(2), 223–246.

Cox, D. R. (1958). The regression analysis of binary sequences (with discussion). *Journal of the Royal Statistical Society, Series B 20*(2), 215–242.

Cox, J. C., J. E. Ingersoll, and S. A. Ross (1985). A theory of the term structure of interest rates. *Econometrica 53*(2), 385–407.

Creal, D. D. (2008). Analysis of filtering and smoothing algorithms for Lévy-driven stochastic volatility models. *Computational Statistics and Data Analysis 52*, 2863–2876.

Creal, D. D. (2012). A survey of sequential Monte Carlo methods for economics and finance. *Econometric Reviews 31*(3), 245–296.

Creal, D. D. (2017). A class of non-Gaussian state space models with exact likelihood inference. *Journal of Business and Economic Statistics 35*(4), 585–597.

Creal, D. D., R. B. Gramacy, and R. Tsay (2014). Market-based credit ratings. *Journal of Business and Economic Statistics 32*(3), 430–444.

Creal, D. D., S. J. Koopman, and A. Lucas (2011). A dynamic multivariate heavy-tailed model for time-varying volatilities and correlations. *Journal of Business and Economic Statistics 29*(4), 552–563.

Creal, D. D., S. J. Koopman, and A. Lucas (2013). Generalized autoregressive score models with applications. *Journal of Applied Econometrics 28*(5), 777–795.

Creal, D. D., S. J. Koopman, and E. Zivot (2010). Extracting a robust U.S. business cycle using a time-varying multivariate model-based bandpass filter. *Journal of Applied Econometrics 25*(4), 695–719.

Creal, D. D., B. Schwaab, S. J. Koopman, and A. Lucas (2014). Observation driven mixed measurement dynamic factor models with an application to credit risk. *The Review of Economics and Statistics 96*(5), 898–915.

Creal, D. D. and R. S. Tsay (2015). High-dimensional dynamic stochastic copula models. *Journal of Econometrics 189*(2), 335–345.

Creal, D. D. and J. C. Wu (2015). Estimation of affine term structure models with spanned or unspanned stochastic volatility. *Journal of Econometrics 185*(1), 60–81.

Creal, D. D. and J. C. Wu (2017). Monetary policy uncertainty and economic fluctuations. *International Economic Review 58*(4), 1317–1354.

Davis, R. A., W. T. M. Dunsmuir, and S. Streett (2003). Observation driven models for Poisson counts. *Biometrika 90*(4), 777–790.

Duffie, D., A. Eckner, G. Horel, and L. Saita (2009). Frailty correlated default. *The Journal of Finance 64*(5), 2089–2123.

Engle, R. F. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom inflation. *Econometrica 50*(4), 987–1007.

Engle, R. F. (2002). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models. *Journal of Business and Economic Statistics 20*(3), 339–350.

Engle, R. F. and G. M. Gallo (2006). A multiple indicators model for volatility using intra-daily data. *Journal of Econometrics 131*(1-2), 3–27.

Engle, R. F. and K. F. Kroner (1995). Multivariate simultaneous generalized ARCH. *Econometric Theory 11*, 122–150.

Engle, R. F. and J. R. Russell (1998). Autoregressive conditional duration: a new model for irregularly spaced transaction data. *Econometrica 66*(5), 1127–1162.

Hamilton, J. D. (1989). A new approach to the econometric analysis of nonstationary time series and the business cycle. *Econometrica 57*(2), 357–384.

Hamilton, J. D. and J. C. Wu (2012). Identification and estimation of Gaussian affine term structure models. *Journal of Econometrics 168*(2), 315–331.

Joslin, S., K. J. Singleton, and H. Zhu (2011). A new perspective on Gaussian affine term structure models. *The Review of Financial Studies 27*(2), 926–970.

Jurado, K., S. C. Ludvigson, and S. Ng (2015). Measuring uncertainty. *American Economic Review 105*(3), 1177–1216.

Kalman, R. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering, Transactions of the ASME 82*(1), 35–46.

Koopman, S. J., A. Lucas, and A. Monteiro (2008). The multi-state latent factor intensity model for credit rating transitions. *Journal of Econometrics 142*(1), 399–424.

McNeil, A. J. and J. Wendin (2007). Bayesian inference for generalized linear mixed models of portfolio credit risk. *Journal of Empirical Finance 14*(2), 131–149.

Patton, A. J. (2006). Modelling asymmetric exchange rate dependence. *International Economic Review 47*(2), 527–556.

Patton, A. J. (2009). Copula-based models for financial time series. In T. G. Anderson, R. A. Davis, J. P. Kreiss, and T. Mikosch (Eds.), *Handbook of Financial Time Series*, pp. 767–786. New York: Springer Verlag.

Russell, J. R. and R. F. Engle (2005). A discrete-state continuous-time model of financial transactions prices and times: the autoregressive conditional multinomial-autoregressive conditional duration model. *Journal of Business and Economic Statistics 23*(2), 166–180.

Shephard, N. (1995). Generalized linear autoregressions. Nuffield College, Oxford.

Zhang, X., D. D. Creal, S. J. Koopman, and A. Lucas (2012). Modeling dynamic volatilities and correlations under skewness and fat tails. Working paper, University of Chicago, Booth School of Business.