ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

COLLIN S. PHILIPPS
Department of Economics, 1407 W Gregory Dr., Urbana IL 61801
www.expectiles.com | cphili6@illinois.edu

Research Statement for Collin S. Philipps

I am an econometrician and empirical macroeconomist focusing on contingency prediction methods and their applications.

My work in econometrics focuses on the development and use of *generalized quantile regression*, with a focus on expectile regression—one type of generalized quantile. My work in that field is groundbreaking and advances the subject significantly. My econometrics job market paper takes expectiles seriously and asks whether there is a best linear unbiased estimator of a sample expectile. Then, I show that optimality results related to "classical" estimators—OLS and GLS—extend to expectiles under a generalized framework, and thus extend to quantile regression under special circumstances. In related work, I show that expectiles can be estimated using a unique quasi-likelihood, and I derive its moments, characteristic function, and other properties. A third paper rebuts recent comments by other authors in this field who had suggested a different quasi-likelihood for the same purpose. My work on quasi-likelihoods lays a foundation for future applications of these methods in empirical macroeconomics, where Bayesian estimation has been my preferred approach.

My research in empirical macroeconomics considers and re-considers the role of fiscal and monetary policy in combatting recessions. To answer whether the government spending multiplier is larger or smaller when monetary policy is hawkish or dovish towards inflation, my coauthor and I contribute the first nonlinear VAR developed for that purpose. We also contribute by modelling the monetary policymaker's decisions endogenously. We show that the government spending multiplier does not vary, except in the very short run, unless the monetary authority commits to hawkish or dovish policy for many years. We also model the Fed's behavior over time and provide a historical narrative. In related work, we ask whether the government spending multiplier varies between expansions and recessions. We contribute by simultaneously addressing three major complaints found in the nonlinear VAR literature: the way the model is estimated, the way impulse responses are calculated, and the way that government spending shocks are identified.

On the following pages, I will introduce these projects and others in more detail.

COLLIN S. PHILIPPS
Department of Economics, 1407 W Gregory Dr., Urbana IL 61801
www.expectiles.com | cphili6@illinois.edu

Main Theme: Contingency Prediction

The common theme in my research is an attention to how the data behave when something unusual takes place. To model special scenarios that do not represent "average" data, I have adopted (1) regression models other than the mean and (2) nonlinear models. The first allows the regression model to behave differently when the dependent variable is higher or lower than usual, and the second allow the regression model to behave differently when right-hand-side variables are higher or lower than usual.

For regression models other than the mean, I have made major theoretical advancements in the study of *generalized quantiles*, particularly with respect to expectile regression. Expectiles are weighted averages that put more (or less) weight on observations above (or below) the regression line. Expectile regression is often compared to quantile regression, but expectiles are simpler to use and can be used in some data environments where quantiles are degenerate.

For nonlinear models, I have contributed improved regime-changing vector autoregressions to study how the real economy reacts to shocks when monetary policy behavior varies, or when there is a recession. Using generalized impulse responses, my coauthor Sebastian Laumer (Economics, UIUC) and I can predict the future state of the economy with and without a shock. We also examine how economic behavior varies depending on the state of the economy when the shock occurs. With our methodological advancements, we have cast new light on the joint role of fiscal and monetary policy in combatting economic recessions.

Econometrics Research: Expectiles, Quantiles, and Likelihood

My job market paper, "*When is an Expectile the Best Linear Unbiased Estimator?*" answers a fundamental question that has not been addressed previously in the literature. It is well-known that ordinary least squares is the best linear unbiased estimator of the mean regression line under simple assumptions, including homoscedasticity. I contribute by pointing out conditions under which there *is* a best linear unbiased estimator of regression lines other than the mean, and that expectile regression is that estimator. I provide a similar result under heteroscedasticity and show what the BLUE will be when heteroscedasticity is present. Then, I provide several model-based motivations for these estimators and show conditions where they can also be the BLUE for quantile regression lines. Each of these estimators is simple, and their use is intuitive. Essentially, my work generalizes the Gauss-Markov theorem to include a generalized version of quantile regression.

In an application, I demonstrate that expectiles are useful in cases where standard quantile regression is not applicable—such as binary response models. With mortgage application denial

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN

COLLIN S. PHILIPPS
Department of Economics, 1407 W Gregory Dr., Urbana IL 61801
www.expectiles.com | cphili6@illinois.edu

data, I show that applicants who are more likely to be denied (than average, conditional on their covariates) also show greater racial disparities than average. Individuals who are about five times as likely to be denied see twice the average racial difference—a black applicant is 32% more likely to be denied than a non-black applicant. These differences diminish when explanatory variables are added, but the ratio remains the same—these relatively disadvantaged applicants have twice the average racial premium. Others have a smaller premium. My regression works in data environments where quantile regression does not, and the results look just as good. For linear probability models, my approach is the paragon of performance.

Following that result, it should be no surprise that expectiles have recently become very popular in several branches of the literature—especially quantitative finance, where expectiles are the only coherent and elicitable risk measure. Recent banking regulations have required banks to report Expected Shortfall, which is a function of expectiles, on their periodic financial statements. But there have been growing pains. In "*The MLE of Aigner, Amemiya, and Poirier is not the Expectile MLE*" (*R&R, Econometric Reviews*), I identify a bubble of misinformation that has spread through the literature. It has somehow become popular to cite an antique paper by Aigner, Amemiya, and Poirier (henceforth AAP) and claim that the likelihood estimator from that paper was the first appearance of expectiles in the literature. But expectiles make no appearance *as such* in that paper, and the likelihood estimator discussed therein does not elicit them.

There *is* a likelihood estimator that elicits expectiles, and it is related to the AAP model. Both are asymmetrically-weighted versions of the Gaussian distribution, and I describe each in detail. I also point out that the asymptotic behavior of the AAP likelihood estimator is nonstandard. Ironically, expectile regression can be used to estimate the AAP location parameter with consistency and asymptotic normality, which the AAP estimator does not have!

In "*Quasi-Maximum Likelihood Estimation for Conditional Expectiles*," we discuss quasi-likelihood estimation of expectiles more formally and more extensively. First, we demonstrate that the likelihood model that elicits expectiles is unique under general conditions. The maximum likelihood estimator thus obtained is consistent, asymptotically normal under standard moment conditions (which can be relaxed further), and it achieves a lower bound for the variance of an expectile regression in at least one special case. We also extend these results to the heteroscedastic regression model, where the estimator is equivalent to the estimator discussed in "*When is an Expectile the Best Linear Unbiased Estimator?*". The consistency, asymptotic normality, and efficiency results apply to that estimator as well.

"*Quasi-Maximum Likelihood Estimation for Conditional Expectiles*" proceeds to analyze the quasi-likelihood model that elicits expectiles, and I contribute probabilistic and information-theoretic results. First, I derive each of the moments, the characteristic function, the entropy, and the Kullback-Leibler divergence from a Gaussian distribution. As it turns out, the quasi-likelihood that elicits expectiles also has the greatest entropy of any distribution with asymmetric variance, conditional on whether the observation is above or below the location parameter (which is true for expectiles). According to the principle of maximum entropy, this model would therefore be the best model to elicit expectiles.

"*Are Forecasters' Loss Functions Observable?*" is a major work in progress with Minchul Shin (Federal Reserve Bank of Philadelphia). This work uses an unprecedented methodology to arrive at simple and clear results. Using the fact that expectiles, quantiles, and other "generalized quantiles" are the same sets of statistics (under some conditions), we note that a forecaster's preferences for summary statistics or predictors can be calculated directly when the forecaster can be observed producing forecasts for different known densities. Using data from the Survey of Professional Forecasters, which includes density forecasts, we calculate which loss functions would most closely replicate each forecaster's point forecasts. This constitutes a direct measure of forecasters' preferences, from which we can determine whether they behave rationally. Some of the forecasters follow no identifiable loss function at all, but others are very close. This work also allows us to determine whether professional forecasters prefer expectiles, quantiles, the mean or the median. They do not prefer any of these! Their point forecasts are usually closest to the *mode* of the density.

Empirical Macroeconomics Research: Recession-busting

"*Government Spending between Active and Passive Monetary Policy*" is joint work with Sebastian Laumer (Economics, UIUC). We contribute the first purely empirical analysis of how government spending multipliers vary, contingent upon the behavior of the monetary policymaker. It is widely believed that the multiplier should be larger when monetary policy is "passive" and allows inflation to increase after the government spending shock. However, theoretical models pointing to that conclusion rely on medium or long-run multipliers with the Fed's policy rule fixed for the entire forecast horizon. We estimate a time-varying Taylor rule and show that the Fed's response to inflation does *not* remain constant over time—the Fed changes its behavior frequently in our historical sample. We estimate a smooth-transition VAR where the behavior of the economy varies along with the Taylor rule response to inflation and find the same result as theorists have found previously, if we calculate impulse responses with the Fed's behavior fixed. But when we allow the Fed to behave as it does in our sample, our nonlinear impulse responses do *not* show different multipliers contingent on the Fed's initial "active" or "passive" stance. This is because the Fed's policy converges quickly to a weakly "active" stance, regardless of its initial condition. Our results suggest that the Fed only deviates from its average policy rule for very short periods of time. Thus, we find that the result found in the theory is *not realistic.*

There is one case where the Fed's monetary policy rule has been fixed for a long period of time: the zero-lower-bound era between 2008-2015. We make another methodological contribution by pointing out that Cholesky identification is not reliable during the zero-lower-bound era: there are several major examples of government spending packages that were passed by congress and

COLLIN S. PHILIPPS
Department of Economics, 1407 W Gregory Dr., Urbana IL 61801
www.expectiles.com | cphili6@illinois.edu

implemented in less than one quarter, which is impossible under the Cholesky restrictions. The TARP program of 2008 was one such example, passed by Congress in October and spending 250 billion dollars before Christmas. The CARES act of 2020 is a more shocking example: it represents the largest counter-cyclical fiscal policy package of any kind in U.S. history and was passed by congress in only a few weeks after the coronavirus was introduced to the United States. The Paycheck Protection Program, for example, disbursed its initial endowment of 349 billion dollars in less than two weeks. In light of these shocking events, we abandon Cholesky identification and utilize a weaker sign-restriction strategy to identify a government spending shock. Our estimated multipliers—from generalized impulse response functions with sign restrictions—are larger than previous estimates from the ZLB period. Importantly, our multiplier from the ZLB era is similar to the "passive" regime multiplier from before 2008. This evidence suggests that government spending shocks may effective to boost the economy at the zero lower bound, and that monetary policy inflexibility is one mechanism for that effect.

In "*Does the Government Spending Multiplier Depend on the Business Cycle?*" with Sebastian Laumer (Economics, UIUC) we adapt similar methodological advancements to study how the effect of government spending varies between expansions and recessions. We note that a body of literature has been critical of previous attempts to answer this question by estimating Bayesian smooth-transition VARs. Others have used partially-calibrated ST-VARs, linear impulse response functions, and Cholesky identification. We contribute by developing a new multi-move Gibbs sampler to fully estimate the model, which is a major step forward. Then, we use our generalized impulse response functions to predict the future economy with and without government spending shocks—without forcing the economy to stay in an expansion or recession for 5 years. We find that the economy reverts to a similar average level of expansion, on average, after about 2 years. We also show that government spending shocks of feasible size do *not* significantly affect the future path of the business cycle. Lastly, we utilize sign restrictions to identify government spending shocks, without imposing Cholesky restrictions. Our multipliers are again slightly larger than others have reported. Importantly, we do find small differences in multipliers in the short run (less than one year), with the multiplier slightly larger during expansions than during recessions. This difference vanishes rapidly over time.

To explore this variation further, we consider several types of government spending. We show that (i) government investment has a larger multiplier than government consumption, (ii) non-defense spending has a higher multiplier than defense spending, (iii) state and local spending have a larger multiplier than federal spending. We also find that none of these multipliers vary across the business cycle. To explain why our results are different from those found by other authors, we also decompose our methodology step-by-step to determine which of our contributions affected the multipliers significantly. As it turns out, neither our fully-estimated model nor our sign-restricted identification strategy were game changers on their own. The generalized impulse response functions, which allow the business cycle to evolve freely within our simulated economy, were the most pivotal contribution.