Dingli Wang

Research Statement

Department of Economics, University of California, Riverside

dingli.wang@email.ucr.edu | dingliwang.github.io

My research is in econometrics and financial econometrics, with related work in forecasting, time series, and machine learning. I study how statistical decisions and forecasts perform when models are estimated from limited data or applied under different distributions. My two job market papers are on human–algorithm replacement under distribution shift and on the adjustment of financial tail-risk forecasts. Other projects examine adversarial density forecasting and the costs of estimating forecast-combination weights.

Human–algorithm replacement under distribution shift

My first job market paper, *Robust Replacement of Human Decision Makers under Distribution Shift*, is a revise-and-resubmit at the *Journal of Econometrics*. It asks when an algorithm should replace an individual decision maker, such as a bail judge, a loan officer, or a doctor, if the data seen by the algorithm may change after deployment. I infer each decision maker's error rates from past decisions and assume that they stay the same on future cases.

The data seen by the algorithm may shift within a label-preserving Wasserstein ball around the training distribution. For each threshold, I compute the largest false positive rate and the smallest true positive rate over the ball, which gives a worst-case ROC curve. The rule replaces a decision maker only if some point on this curve is, with high posterior probability, at least as good as the decision maker in both rates. I show that the estimated curve converges uniformly at the root-n rate and that, with probability approaching one, the rule replaces exactly the decision makers below the true curve. Under every distribution in the ball, the possible welfare loss relative to retaining all decision makers shrinks with more data. Relative to a benchmark that knows the deployment distribution and each decision maker's error rates, however, choosing the threshold with the highest posterior probability can leave a welfare loss that more data do not remove.

The ball's radius is an average budget: it can be spent on the few cases closest to a threshold. Even a small one moves the worst-case rates much more than a pointwise budget of the same size, which limits how far each case can move. In simulations, average budgets from 0.05 to 0.3 retain every decision maker, while pointwise budgets of 0.05 and 0.1 replace some and raise mean welfare even under the worst shift within the budget. In an illustration with 2018 American Community Survey income data, the algorithm is trained on three large states and used in five others, with synthetic decision makers.

Adjustment of joint tail-risk forecasts

My second job market paper, *Median-Anchored Adjustment of Joint VaR–ES Forecasts* (with Tae-Hwy Lee), is under review at the *Journal of Business & Economic Statistics*. Risk managers may need to adjust a fitted model's tail forecasts without rebuilding the model. We keep its reported median fixed and multiply the distances from the median to the Value-at-Risk (VaR) and Expected Shortfall (ES) by the same positive factor. This preserves the VaR–ES ordering without a separate constraint.

We derive a closed-form weighted-quantile estimator for the adjustment factor and establish its consistency and asymptotic distribution, accounting for the estimation of the baseline model. We also give a bound on the VaR coverage error. When the same factor correctly adjusts both VaR and ES, we show that the adjustment reduces the conditional joint VaR–ES forecast risk. In rolling S&P 500 forecasts, the adjustment lowers the FZ0 joint VaR–ES loss of the GARCH baselines, most (by 10.4%) for Gaussian GARCH at the 1% tail. The benefit is limited when the baseline already models the tail flexibly.

Adversarial density forecasts in short samples

In *Adversarial Density Forecast for Macro-financial Risks*, Tae-Hwy Lee and I study conditional density forecasts estimated by Wasserstein generative methods. Flexible models can fit features of a distribution that parametric forecasts miss, but short macroeconomic samples make their estimates sensitive to training choices. We add a penalty that matches the variance and selected quantiles of generated outcomes to those in the observed data.

We relate the penalty to bounds on the Wasserstein forecast error and the expected continuous ranked probability score, allowing for a finite number of simulation draws. In simulations, the penalized estimator improves on the unpenalized one and varies less across fits. We also evaluate monthly U.S. unemployment, inflation, and industrial-production forecasts. The gains over the unpenalized estimator differ across series and are clearest for unemployment.

Forecast combination and the cost of estimating weights

I study the choice between equal weights and estimated weights in *When Should Forecasts Be Equally Weighted? Signal Recovery and Estimation Costs*. I express the expected squared-error gain from estimated weights as recovered predictive signal minus the costs of estimating the component forecasts and combination weights. A tuning parameter selected from the data must also be treated as part of the forecasting procedure. Ignoring refitting after selection can even reverse the estimated gain.

For implementation, I compare the complete procedures using past out-of-sample forecast losses. Under stationarity and weak dependence, the selection rule favors the procedure with lower expected loss as the evaluation sample grows, provided the expected losses differ. The empirical analysis uses U.S. macroeconomic forecasts. In some settings, the simple average remains the better choice once weight estimation and tuning are included.

Future research

I plan to extend the replacement problem to settings in which the decision makers' error rates can also change, for example when loan officers read the same shifted credit reports as the algorithm. Updating replacement decisions as deployment data arrive would require estimating these changes while keeping track of the remaining uncertainty about both the decision makers and the algorithm.

My forecasting work suggests a second direction: allowing a small amount of model flexibility to respond to observable economic conditions. This includes state-dependent VaR–ES adjustment factors, forecast-combination choices that vary with predictive conditions, and distributional anchors chosen for the forecast horizon or tail outcome. These extensions share an estimation problem: a method that reacts to changing conditions may improve forecasts, but only if there are enough observations to estimate that response reliably.