

Supplementary Material

Forecasting realized oil-price volatility: The Role of financial stress and asymmetric loss

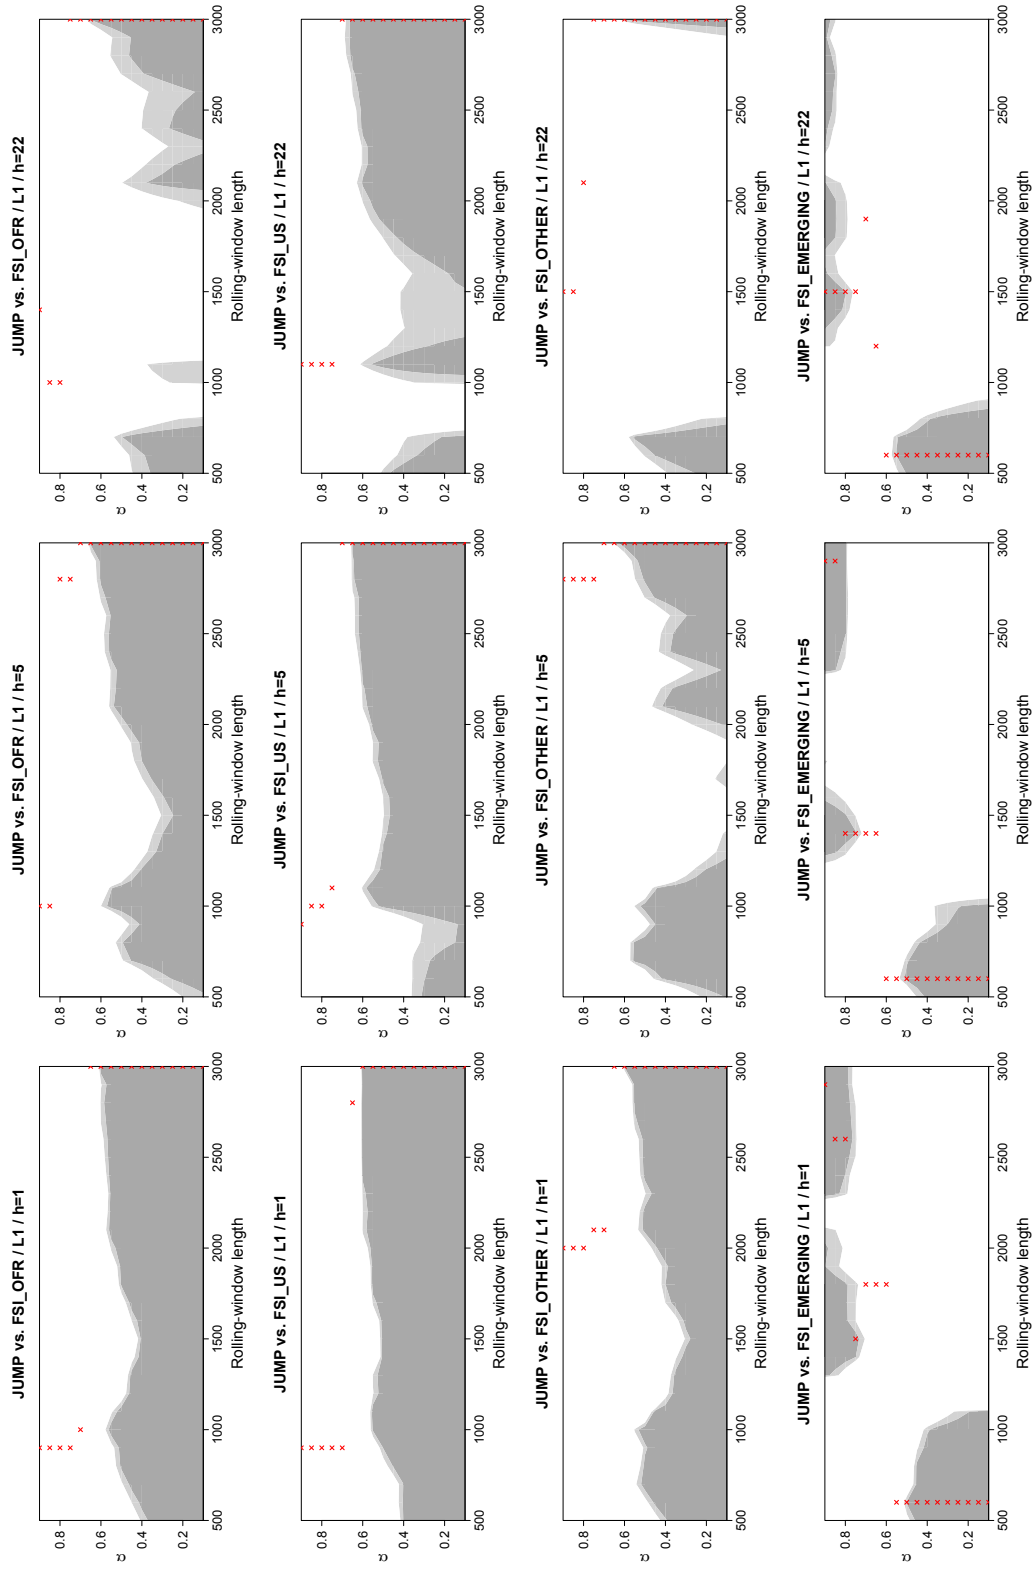
Konstantinos Gkillas, Rangan Gupta, Christian Pierdzioch

We present here additional results for the following scenarios:

- Figure **S1**: Out-of-Sample Results (Rival Model Features Jumps / L1 Loss)
- Figure **S2**: Out-of-Sample Results (Rival Model Features Jumps / L2 Loss)
- Figure **S3**: Out-of-Sample Results (Rival Model Features Higher-Order Moments / L1 Loss)
- Figure **S4**: Out-of-Sample Results (Rival Model Features Higher-Order Moments / L2 Loss)
- Figure **S5**: Out-of-Sample Results (Financial Stress Categorized Along Asset Classes / L1 Loss)
- Figure **S6**: Out-of-Sample Results (Financial Stress Categorized Along Asset Classes / L2 Loss)
- Figure **S7**: HAR-RV Estimated on $\ln(RV^{0.5})$ / L1 Loss)
- Figure **S8**: HAR-RV Estimated on $\ln(RV^{0.5})$ / L2 Loss)
- Figure **S9**: Out-of-Sample Results (Longer Forecasting Horizons / L1 Loss)
- Figure **S10**: Out-of-Sample Results (Longer Forecasting Horizons / L2 Loss)

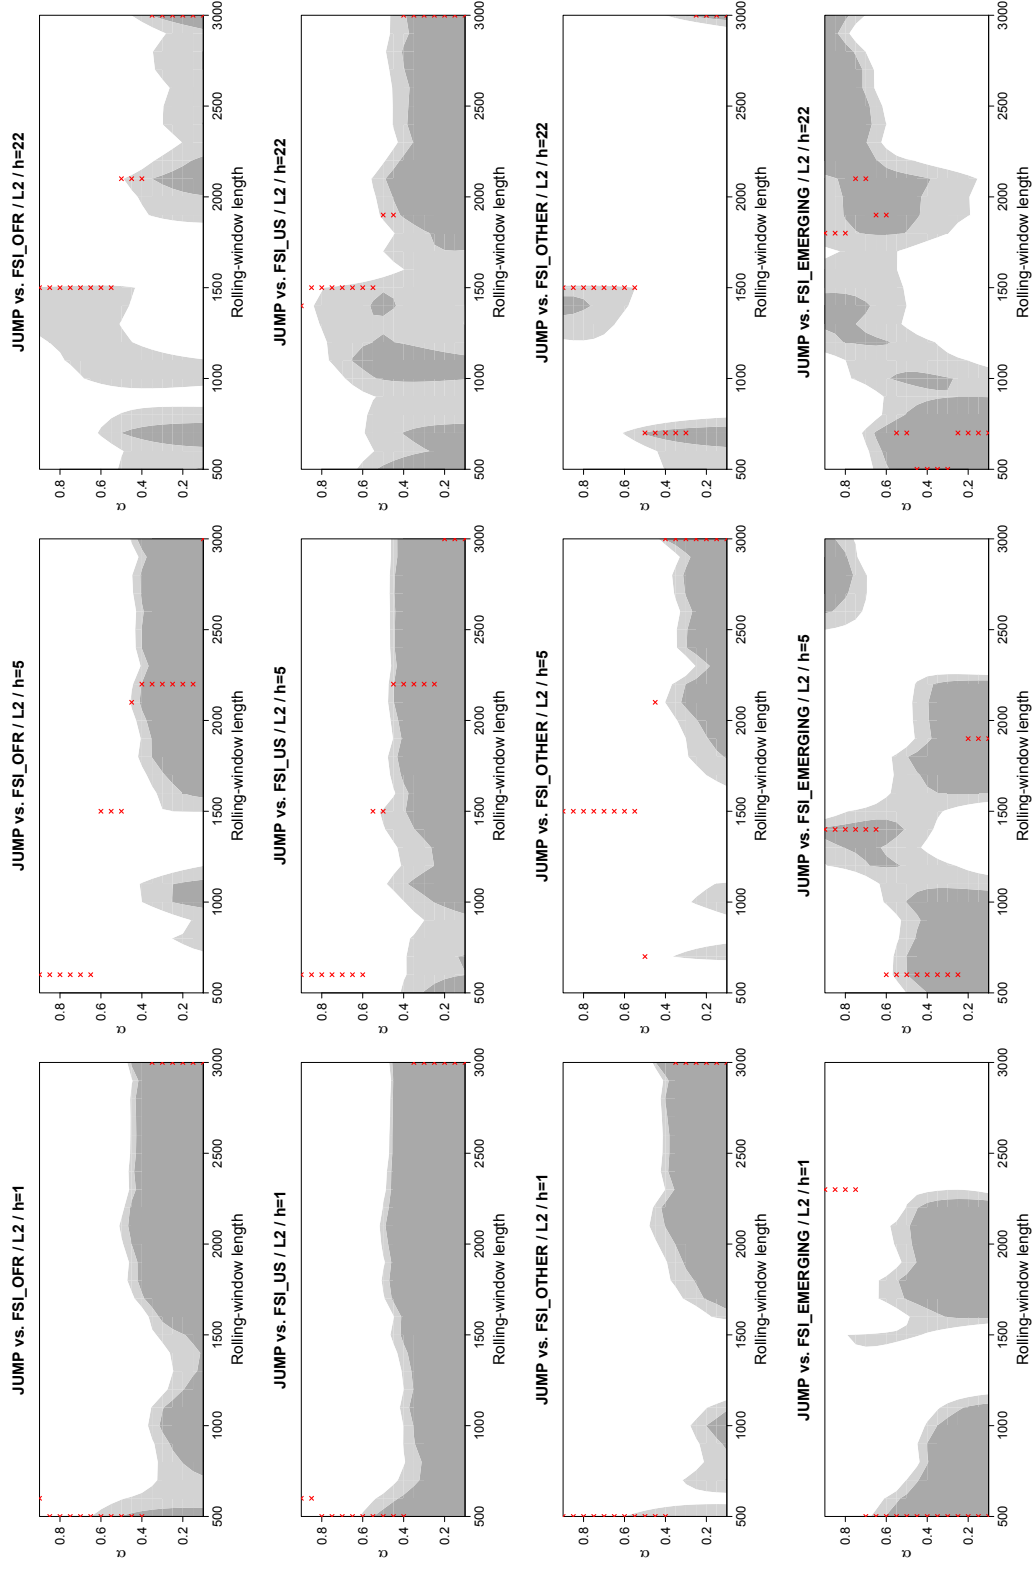
- Figure S11: Out-of-Sample Results (Models Feature Weekly and Monthly FSI / L1 Loss)
- Figure S12: Out-of-Sample Results (Models Feature Weekly and Monthly FSI / L2 Loss)
- Figure S13: Out-of-Sample Results (mean(MRV) / L1 Loss)
- Figure S14: Out-of-Sample Results (mean(MRV) / L2 Loss)
- Figure S15: Out-of-Sample Results (Recursive Estimation Window)

Figure S1: Out-of-Sample Results (Rival Model Features Jumps / L1 Loss)



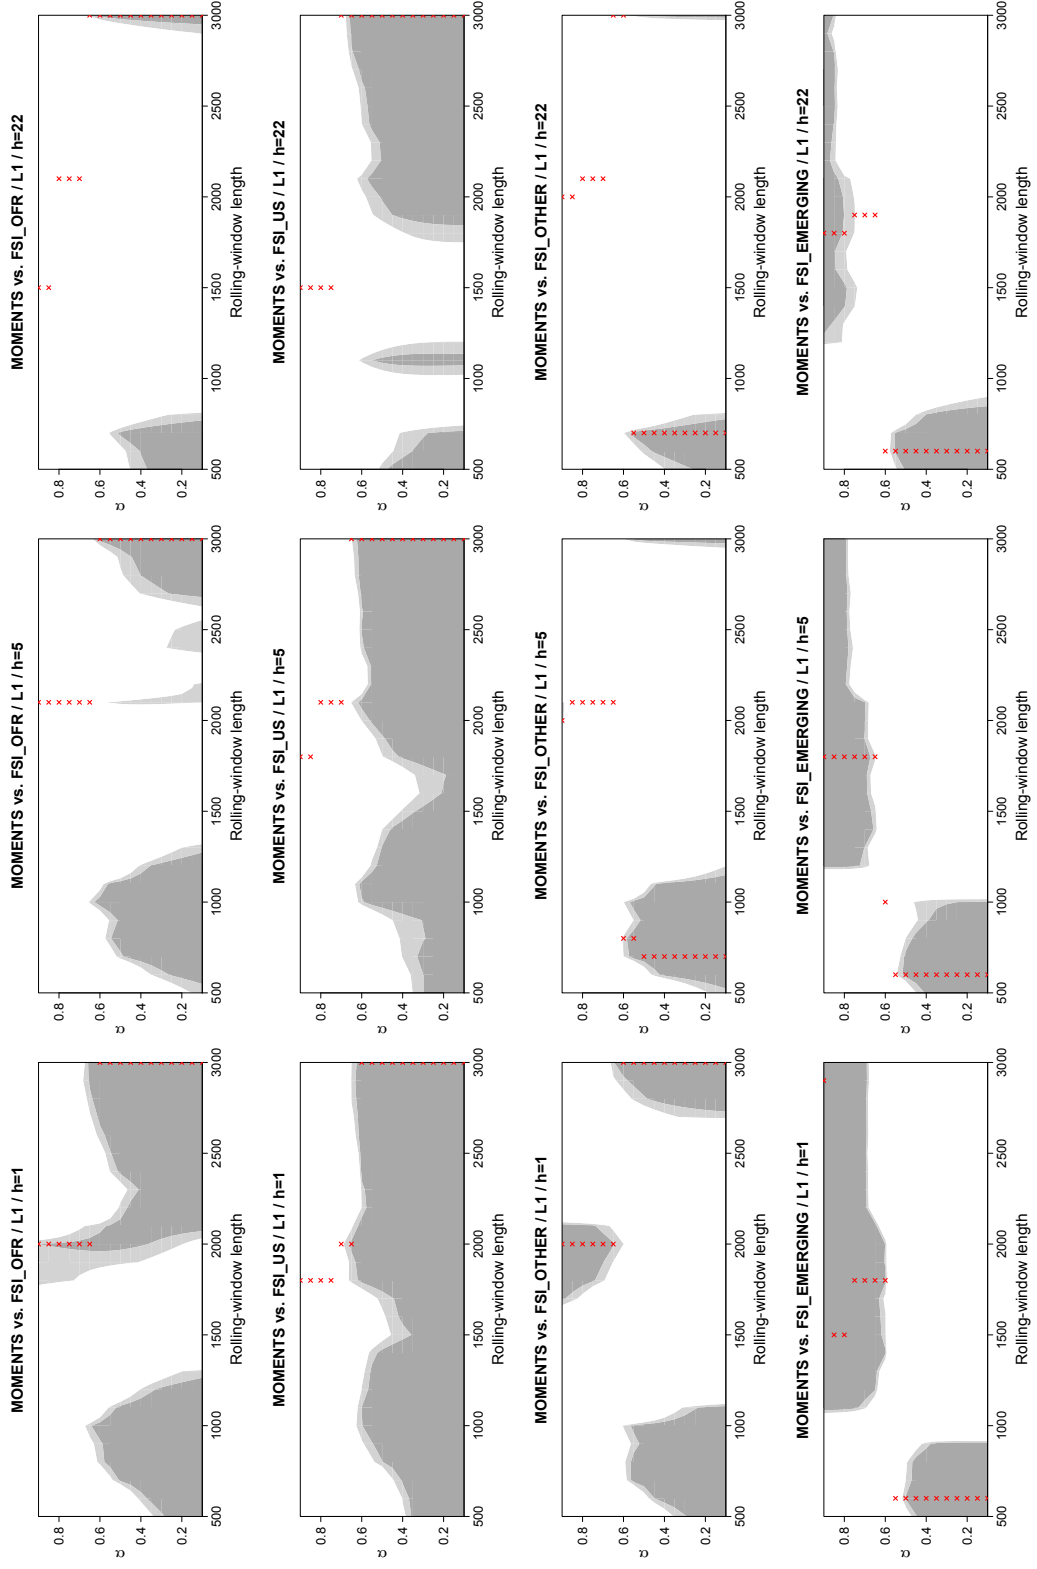
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L1: Lin—lin loss function.

Figure S2: Out-of-Sample Results (Rival Model Features Jumps / L2 Loss)



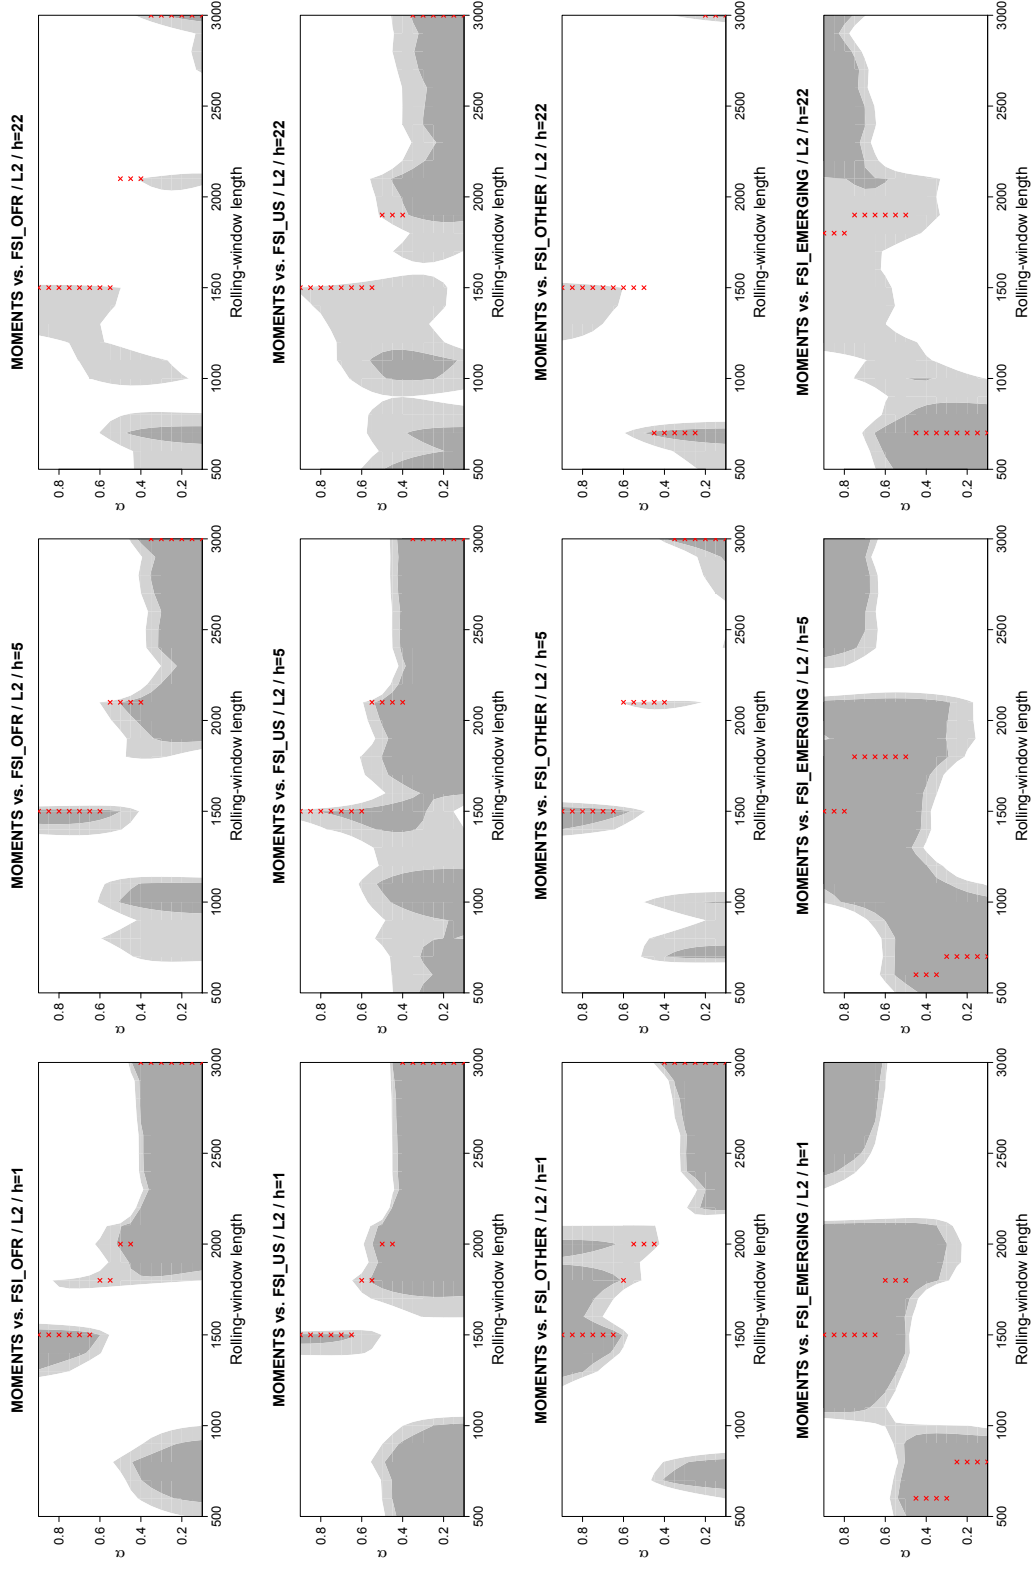
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L2: Quad-quad loss function.

Figure S3: Out-of-Sample Results (Rival Model Features Higher-Order Moments / L1 Loss)



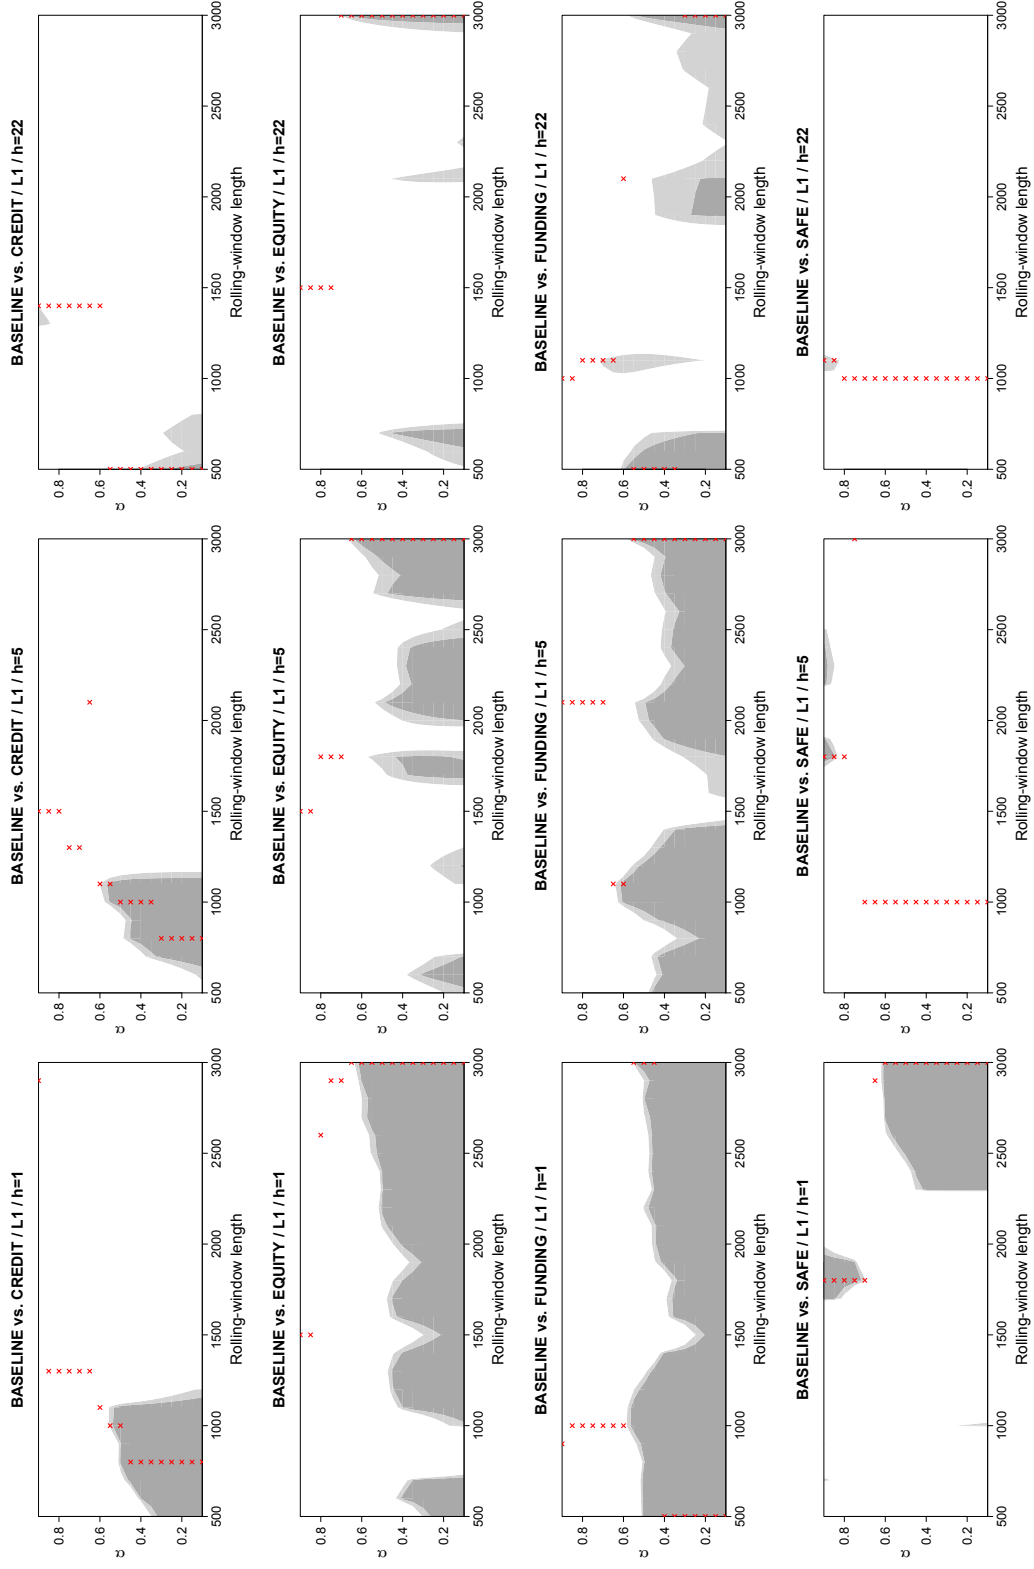
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L1: Lin—lin loss function. Higher-Order moments: Realized skewness and realized kurtosis.

Figure S4: Out-of-Sample Results (Rival Model Features Higher-Order Moments / L2 Loss)



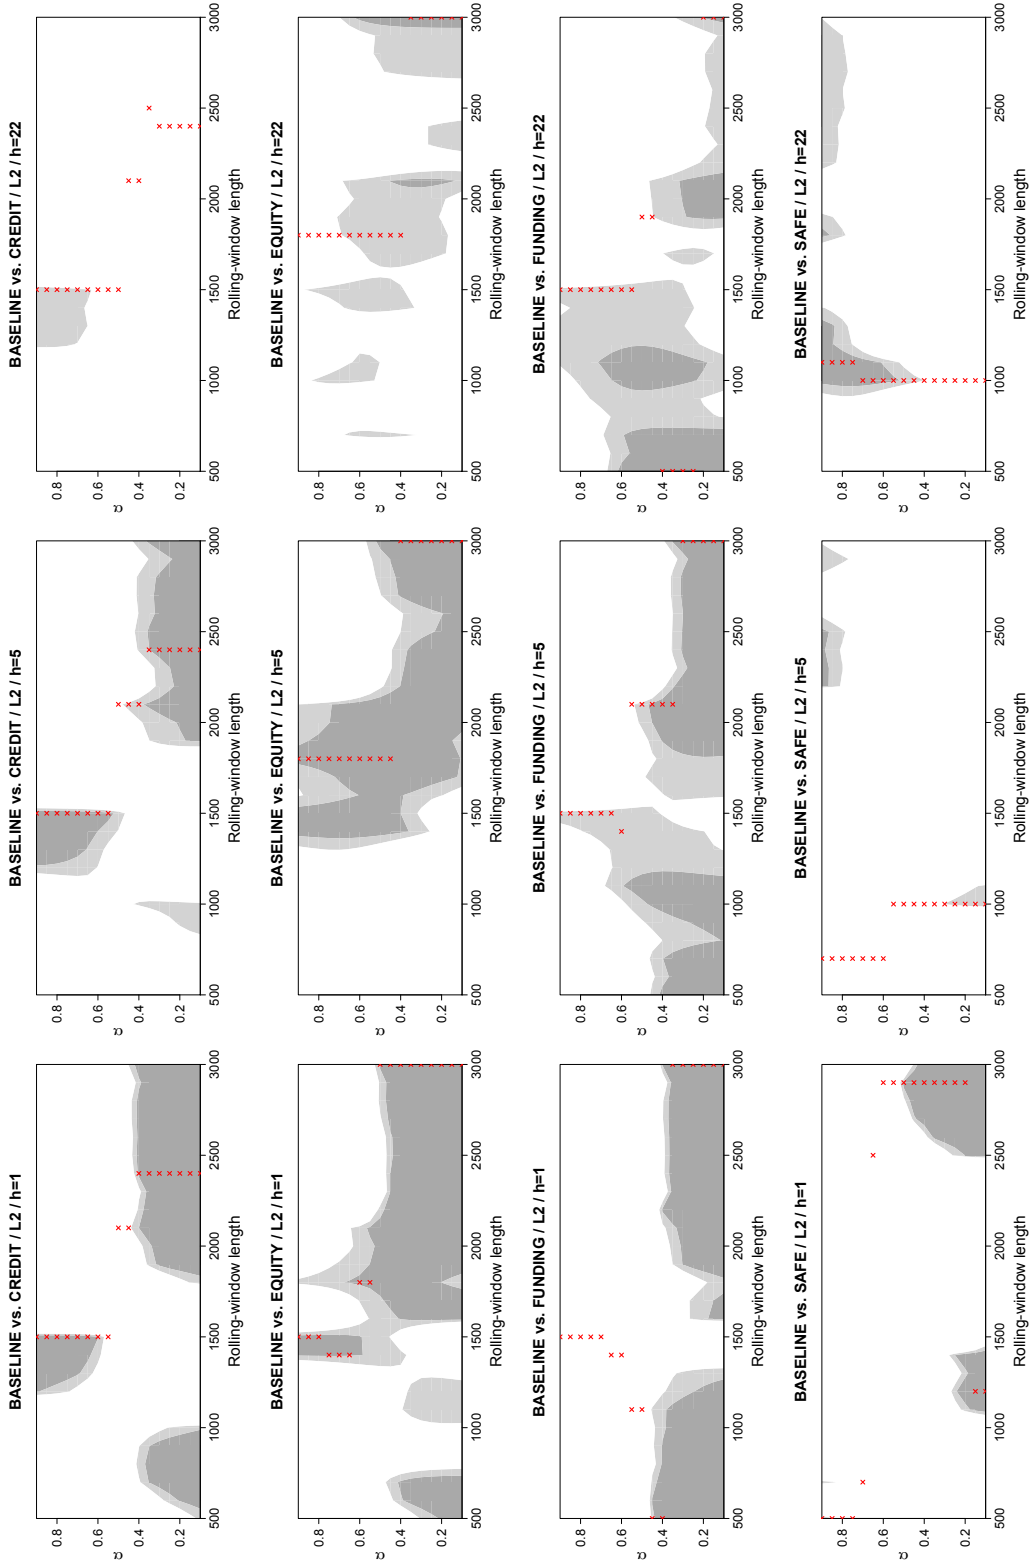
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L2: Quad-quad loss function. Higher-Order moments: Realized skewness and realized kurtosis.

Figure S5: Out-of-Sample Results (Financial Stress Categorized Along Asset Classes / L1 Loss)



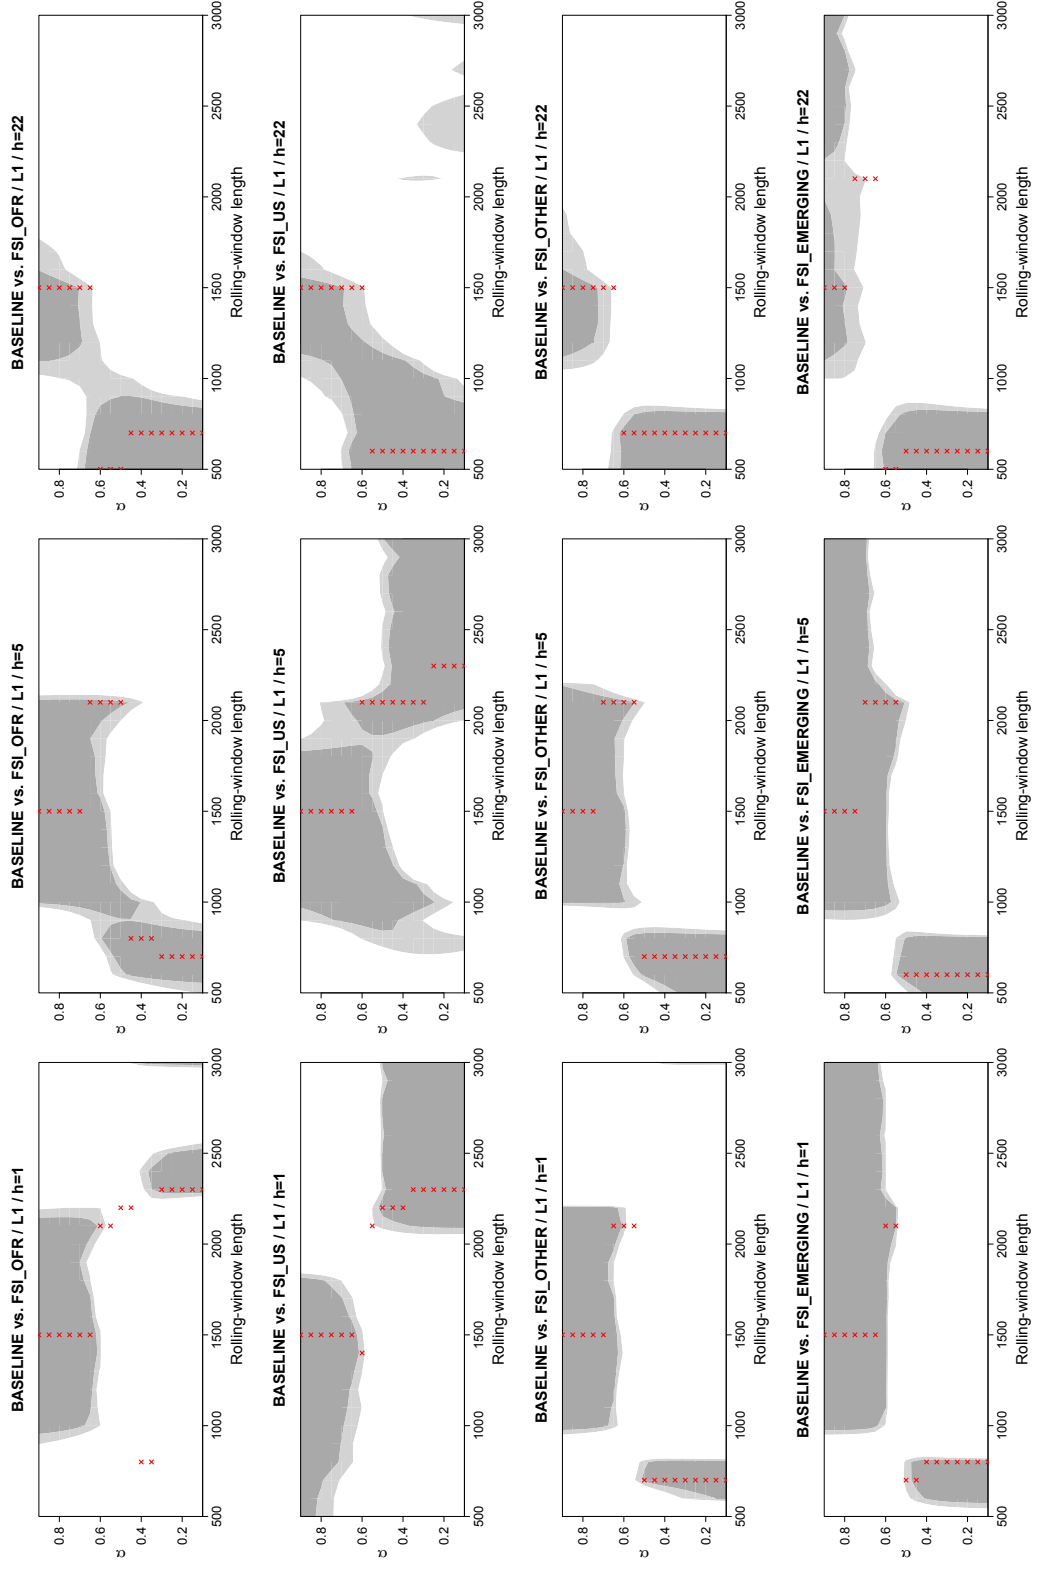
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L1: Lin—lin loss function.

Figure S6: Out-of-Sample Results (Financial Stress Categorized Along Asset Classes / L2 Loss)



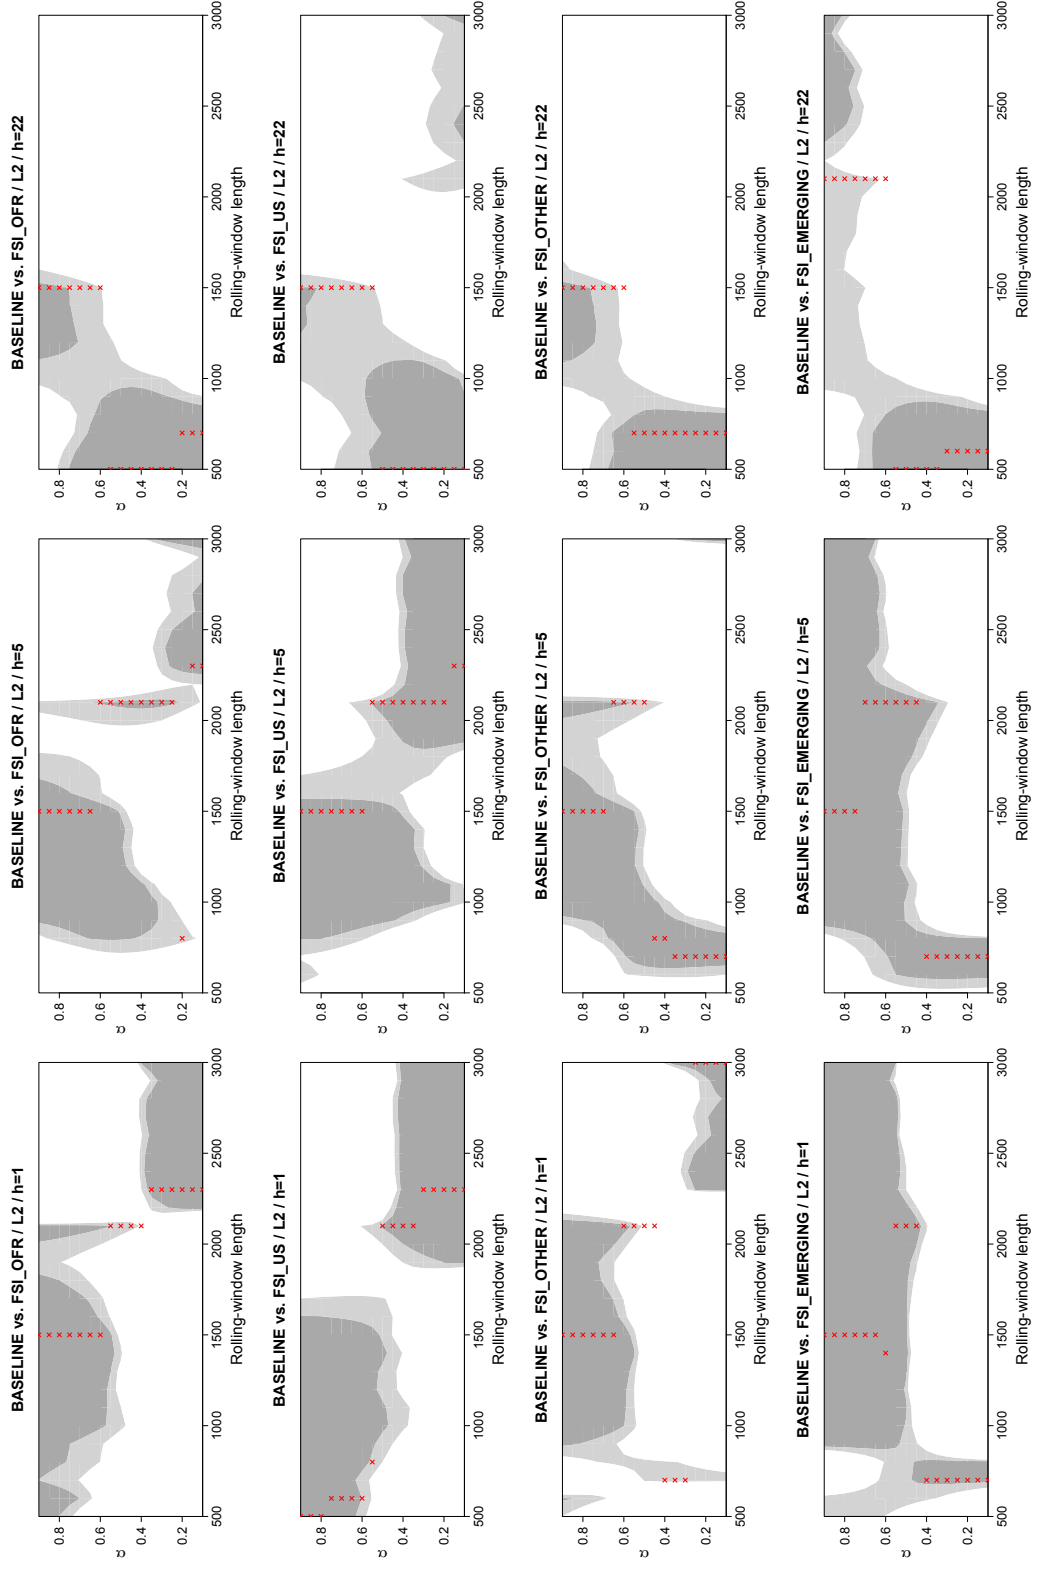
Note: Black area: Diebold-Mariano test is significant at the 10% level. Gray area: Diebold-Mariano test is significant at the 5% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L2: Quad-quad loss function.

Figure S7: Out-of-Sample Results (HAR-RV Estimated on $\ln(RV^{0.5}) / L1$ Loss)



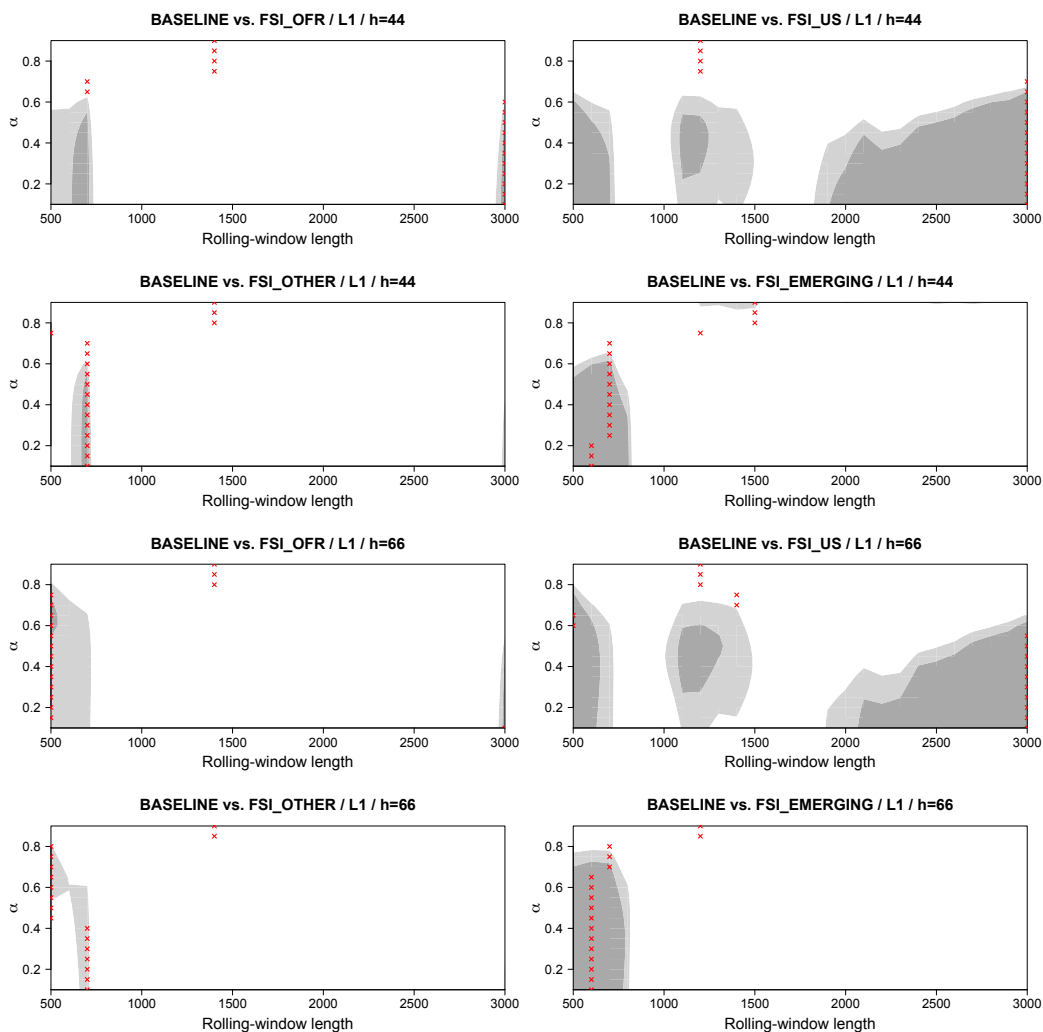
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L1: Lin—lin loss function.

Figure S8: Out-of-Sample Results (HAR-RV Estimated on $\ln(RV^{0.5}) / L2$ Loss)



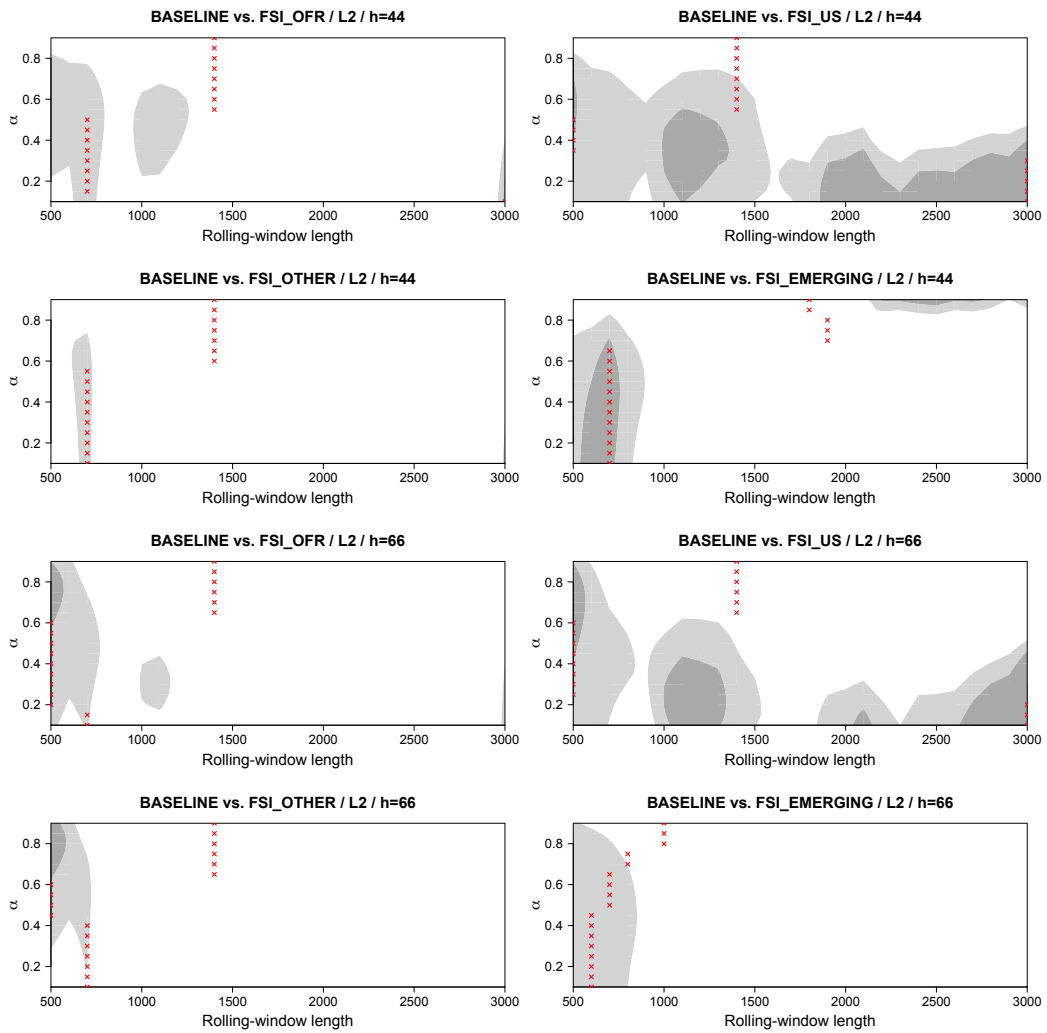
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L2: Quad-quad loss function.

Figure S9: Out-of-Sample Results (Longer Forecasting Horizons / L1 Loss)



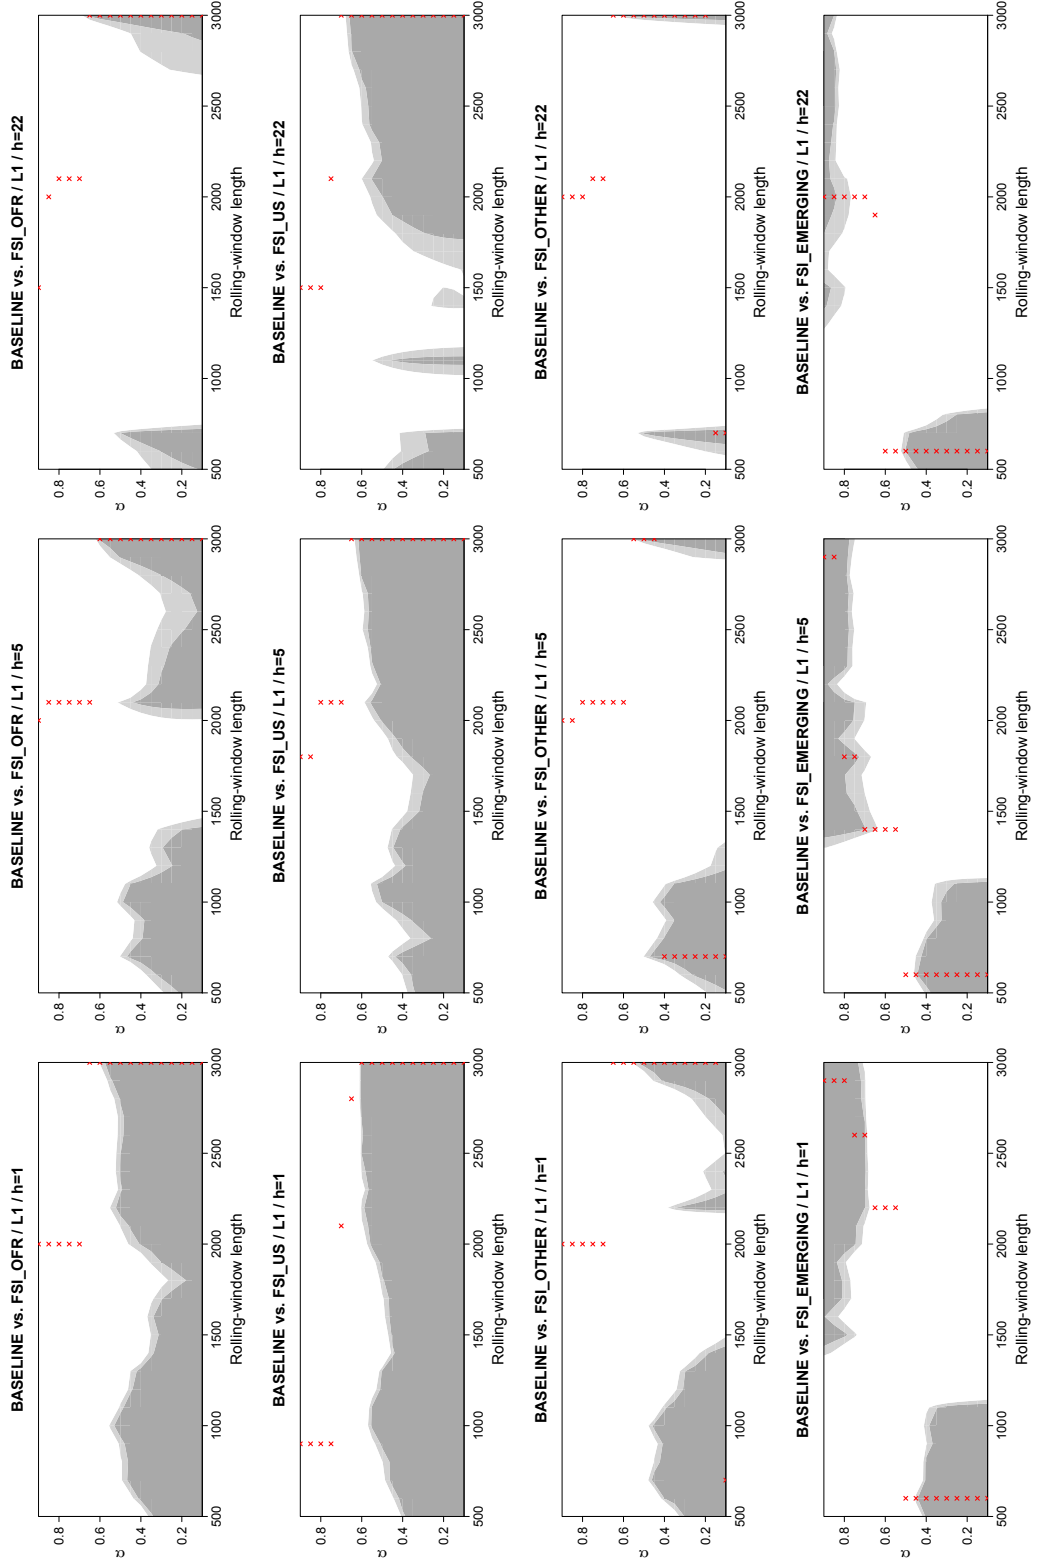
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L1: Lin–lin loss function.

Figure S10: Out-of-Sample Results (Longer Forecasting Horizons / L2 Loss)



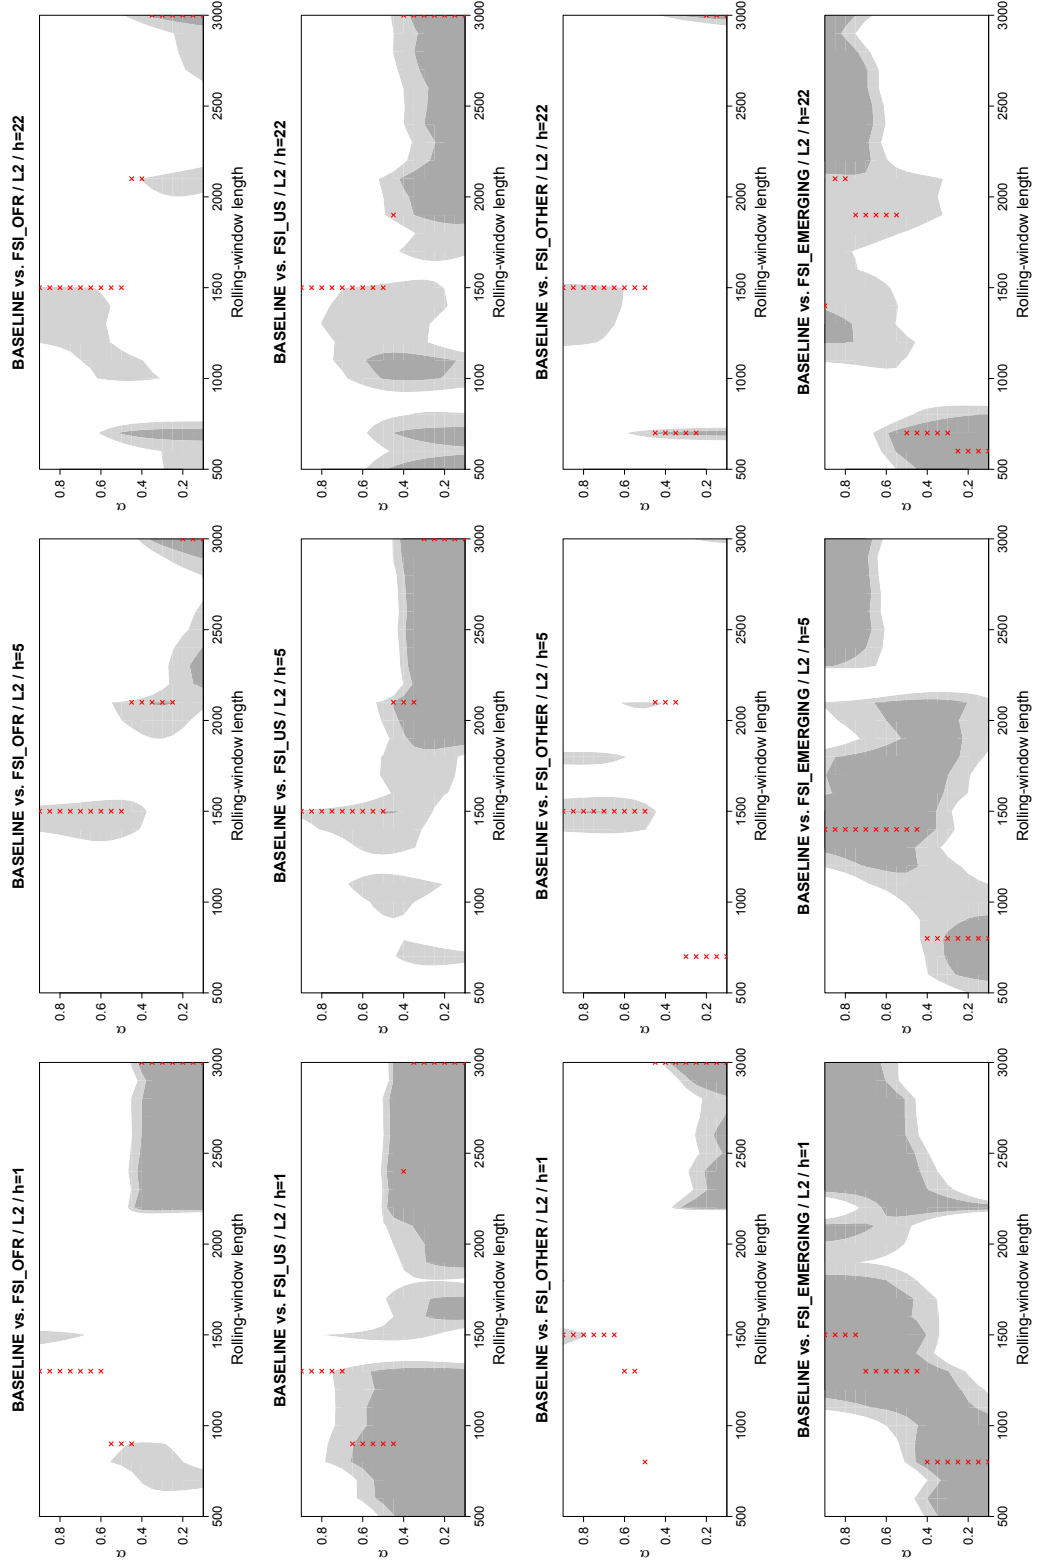
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L2: Quad-quad loss function.

Figure S11: Out-of-Sample Results (Models Feature Weekly and Monthly FSI / L1 Loss)



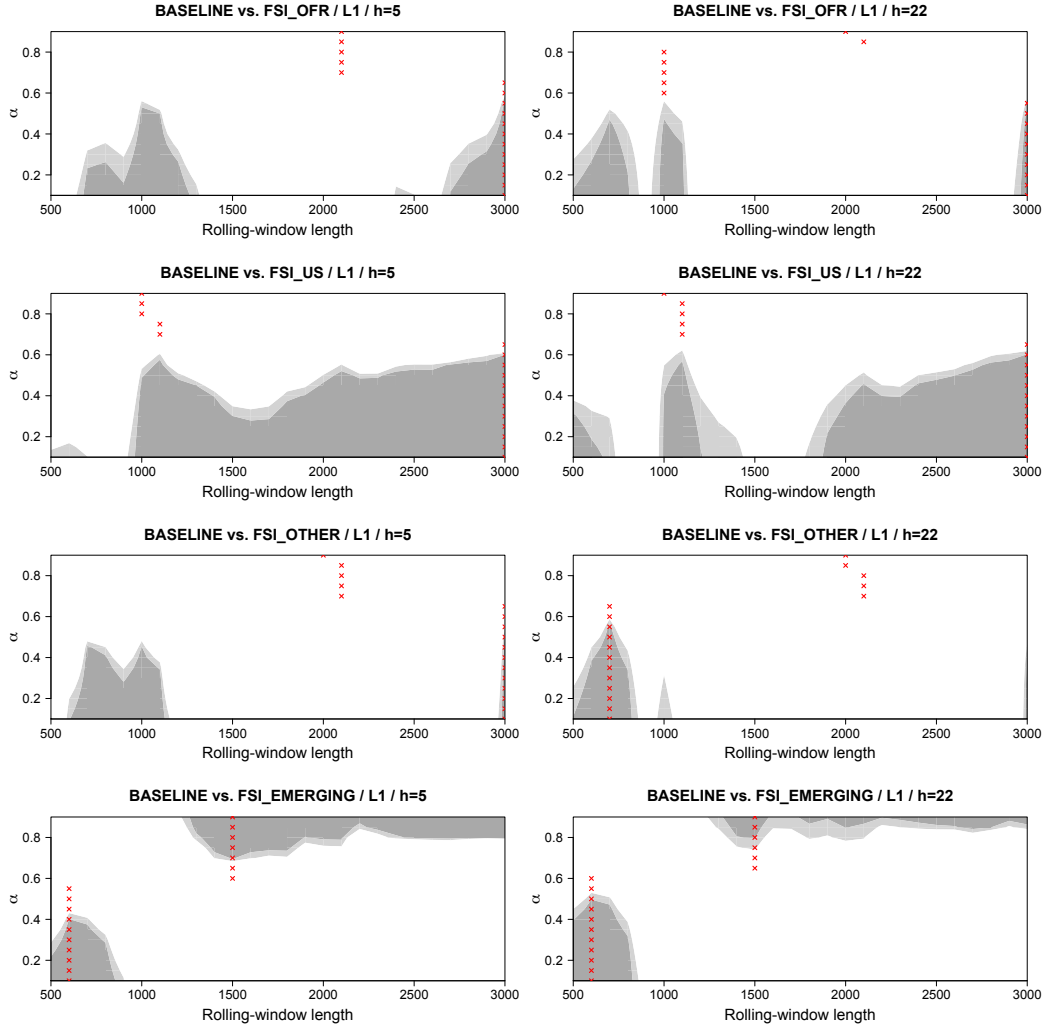
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L1: Lin—lin loss function.

Figure S12: Out-of-Sample Results (Models Feature Weekly and Monthly FSI / L2 Loss)



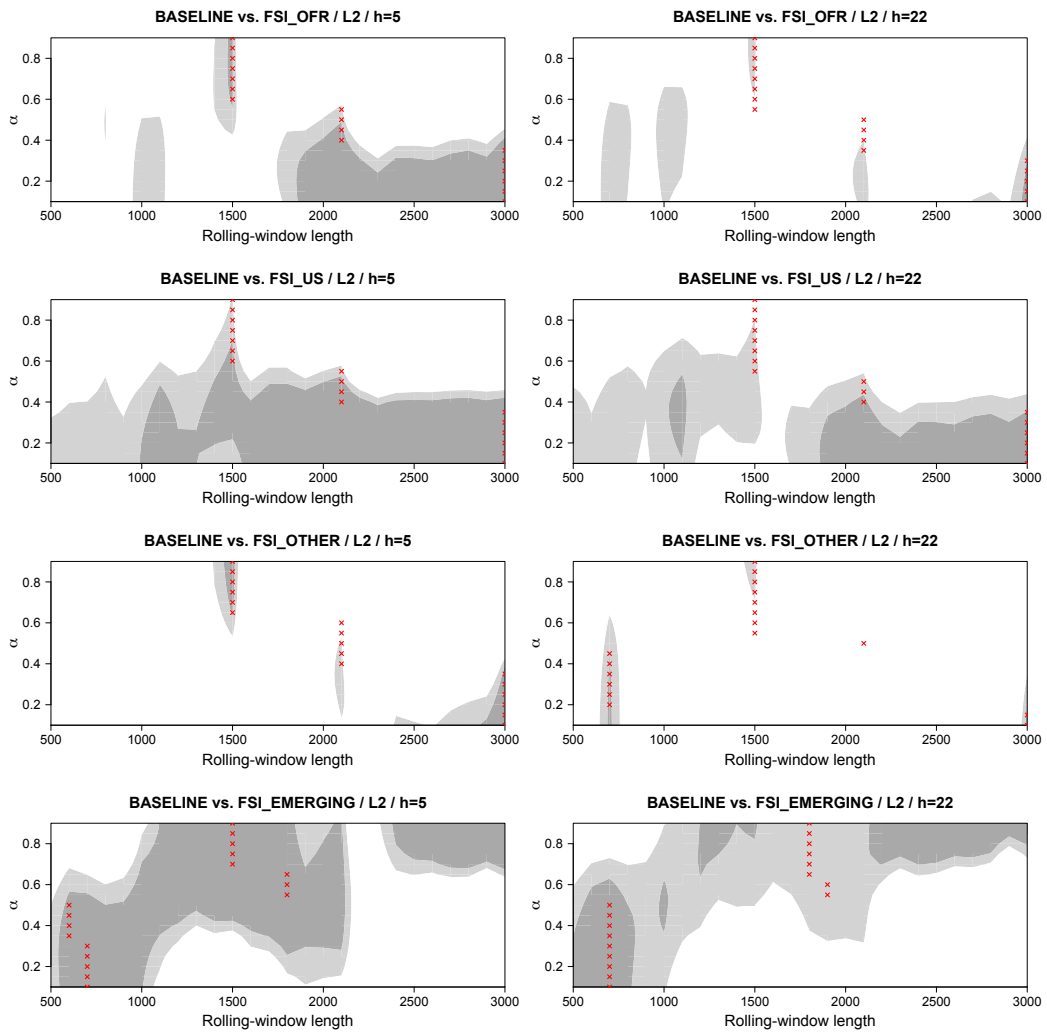
Note: Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L2: Quad-quad loss function.

Figure S13: Out-of-Sample Results (mean(MRV) / L1 Loss)



Note: Forecasts for $h = 5, 22$ are for mean(MRV). Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L1: Lin-lin loss function.

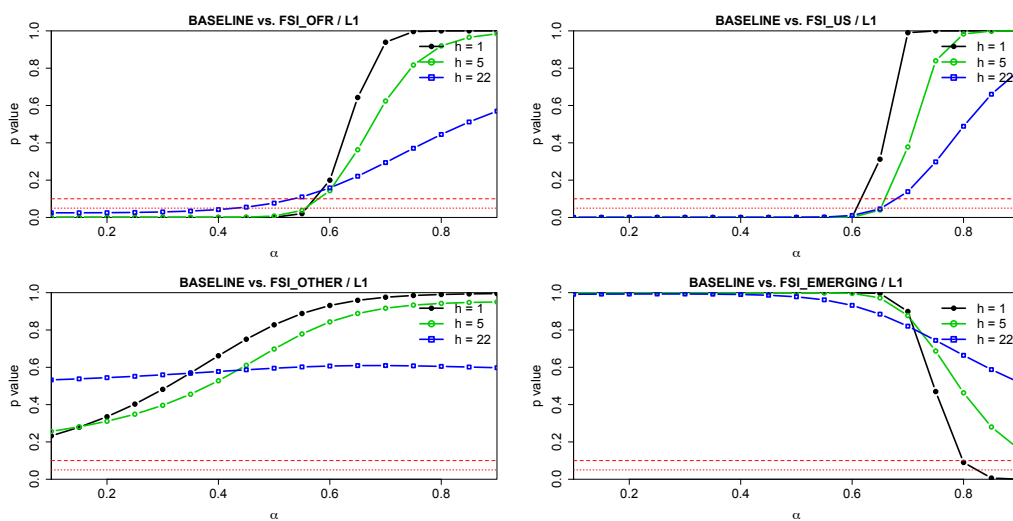
Figure S14: Out-of-Sample Results (mean(MRV) / L2 Loss)



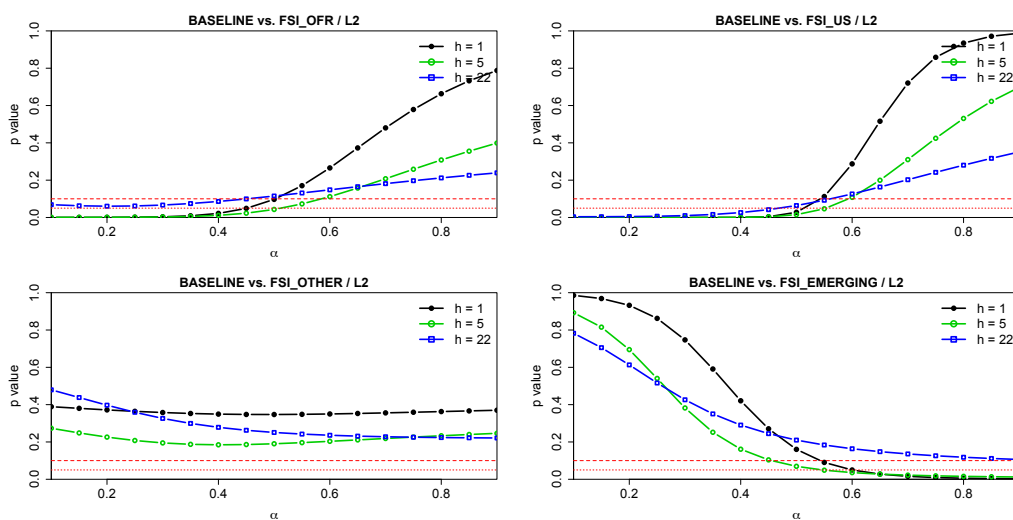
Note: Forecasts for $h = 5, 22$ are for mean(MRV). Black area: Diebold-Mariano test is significant at the 5% level. Gray area: Diebold-Mariano test is significant at the 10% level. Null hypothesis: the two series of forecasts are equally accurate. Alternative hypothesis: the forecasts from the model extended to include financial stress is more accurate. Results are based on rolling-window estimates. The horizontal axis displays the length of a rolling window. The vertical axis displays the asymmetry parameter of the loss function. L2: Quad-quad loss function.

Figure S15: Out-of-Sample Results (Recursive Estimation Window)

Panel A: L1 loss



Panel B: L2 loss



Note: This figure shows p-values of the Diebold-Mariano test. Results are based on a recursive window estimates. Training period: 500 observations. L1: Lin-lin loss function. L2: Quad-quad loss function. Red horizontal lines: 10% and 5% levels of significance.