The impact of the number and the size of clusters on prediction performance of the stratified and the conditional shared gamma frailty Cox proportional hazards models

Daniele Giardiello, PhD (1)

Edoardo Ratti, MSc (1)

Peter C. Austin, PhD (2,3,4)

Corresponding author:

Daniele Giardiello

Bicocca Bioinformatics Biostatistics and Bioimaging B4 Center, School of Medicine and Surgery, University of Milano-Bicocca

U28 Via Follereau 3

20854 Vedano al Lambro (MB), Lombardy

Italy

E-mail: daniele.giardiello1@gmail.com

Author affiliations:

(1) Bicocca Bioinformatics Biostatistics and Bioimaging B4 Center, School of Medicine and Surgery, University of Milano-Bicocca, Monza, Italy.

(2) ICES, Toronto, Ontario, Canada.

(3) Institute of Health Policy, Management and Evaluation, University of Toronto, Ontario, Canada.

(4) Sunnybrook Research Institute, Toronto, Ontario, Canada.

Word count: 4522 (excluding abstract)

Abstract (word count: 240/250**)**

Researchers in biomedical research often analyse data that are subject to clustering. Independence among observations are generally assumed to develop and validate risk prediction models. For survival outcomes, the Cox proportional hazards regression model is commonly used to estimate an individual's risk at fixed time horizons. The stratified Cox proportional hazards and the shared gamma frailty Cox proportional hazards regression models are two common approaches to account for the presence of clustering in the data. The accuracy of the predictions of these two approaches has not been examined. We conducted a set of Monte Carlo simulations to assess the impact of the number of clusters, the size of the clusters, and the within-cluster correlation in outcomes on the accuracy of the conditional predictions developed using the stratified and the shared gamma frailty Cox proportional hazards regression model. We compared the accuracy of the predictions in terms of discrimination, calibration and overall performance metrics. We found that the stratified and the shared gamma frailty model provided similar performance, especially for larger size and higher number of clusters. For small cluster size, we observed slightly better discrimination and overall performance for the stratified model and better calibration for the shared gamma frailty model especially at shorter prediction horizons. However, the practical applicability of the stratified Cox proportional hazards model to estimate predictions is limited especially for high within-cluster correlation and when clusters are small, and more likely at longer time prediction horizons.

Keywords (max 6): risk prediction models, clustered data, Cox regression, discrimination, calibration, Monte Carlo simulations

1. Introduction

Researchers in biomedical research commonly often analyse data that are subject to clustering. Over the last decades, international collaborations have become more frequent, allowing researchers to investigate disease diagnosis and prognosis in patients from different centres and countries. Multicentre clinical trial data, observational data with geographic clusters, family-based studies, and individual patient data meta-analysis are examples of individuals nested within different clusters such as centres, countries, genetically related individuals and studies, respectively. An important aim of clinical risk prediction models is to predict an individual's risk of the outcome using statistical regression models that typically assume independence among observations. When observations are nested within clusters, regression models should consider the presence of clustering in the data used to develop, to validate, and to deploy risk prediction models.

The Cox proportional hazards model is frequently used to develop prediction models for the analysis of time-to-event outcomes in biomedical and epidemiological research[1]. The Cox proportional hazards model allows one to estimate an individual's risk of experiencing the outcome at specific time horizons such as, for example, to estimate the 10-year risk of death. The stratified Cox proportional hazards models and frailty Cox models are the most common survival models that account for clustering[2]. The stratified Cox proportional hazards models assume that individuals in each cluster (namely, strata) have different baseline hazard functions. Using a different baseline hazard function for each cluster, the estimated risk predictions can be different in individuals with the same characteristics coming from different clusters. The frailty Cox proportional hazards models include random effects assuming that the clusters such as centres or countries are drawn from a population of possible clusters. This enables one to account for unobserved heterogeneity between clusters. For computational and mathematical reasons, the most common frailty Cox proportional hazards model is the shared gamma frailty distribution[2,3]. Risk predictions can be obtained in closed mathematical form using the shared gamma frailty Cox proportional hazards model. Both conditional and marginal predictions consider clustering in predicting individual's risk. In particular, the conditional frailty models estimate an individual's risk among patients from the same clusters that were used to develop the model. The marginal frailty models estimate predictions that are averaged across clusters. Thus, predictions can also be also estimated for patients coming from different clusters than those used to develop the model considering that data are nested within clusters.

To the best of our knowledge, there is a paucity of research on the impact of the size and the number of clusters on the accuracy of predictions obtained from the stratified and the conditional frailty Cox proportional hazards models. The objective of the current paper was to examine the impact of the number of clusters, the cluster size, and the within-cluster correlation in outcomes on the accuracy of the conditional prediction models developed using the stratified and frailty Cox proportional hazards regression model. The paper is structured as follows: in Section 2, we describe a series of Monte Carlo simulations that were designed to address this question. In Section 3, we report the results of these simulations. In Section 4, we provide a case study to illustrate the application of these methods. Finally, in Section 5, we summarize our findings and place them in the context of the existing literature.

2. Monte Carlo simulations methods

We conducted a set of Monte Carlo simulations to examine the impact of the size of clusters, the number of clusters, and within-cluster correlation on the accuracy of predictions obtained from the stratified and the shared gamma frailty Cox proportional hazards models using clusters as stratification and frailty term, respectively. We simulated clustered data with time-to-event outcomes related to subject-level covariates.

2.1 Factors in the Monte Carlo simulations

We allowed three factors to vary in the Monte Carlo simulations: (i) the number of clusters ($N_{cluster}$); (ii) the mean number of subjects per cluster ($N_{subjects}$); (iii) the within-cluster correlation in the time-to-event outcomes (τ). The first took on values from 2 to 10 in increments of 2, for a total of 5 values. The second took on four values: 25, 50, 100, 200, for a total of 4 values. The third took on four values: 0.01, 0.05, 0.10, and 0.20, for a total of 4 values. We used a full factorial design and thus examined 80 different scenarios. In each scenario we generated 1 000 datasets.

2.2 Data Generating Process

For each of the $N_{cluster}$ clusters we generated a random variable denoting the size of the cluster ($N_{subjects}$). The size of each cluster was drawn from a uniform distribution with the lower and upper limit of $0.5*N_{subjects}$ and $1.5*N_{subjects}$, respectively. Thus, the average size of each cluster was $N_{subjects}$, rounded to the nearest integer. In every scenario, we generated two independent subject-level baseline covariates: X1 and X2. The first was

continuous and was generated from a standard normal distribution, while the second was binary and was generated from a Bernoulli distribution with parameter 0.5. Thus, in expectation, Pr(X2 = 1) = 0.5.

For each subject we simulated a time-to-event outcome and a censoring time. First, we simulated correlated individual-level time-to-event outcomes using a gamma frailty model combined with Bender and colleagues' approach for simulating time-to-event outcomes using a Cox–Weibull model[4]. Define LP_{ij} to be the log-hazard of the outcome, which is a linear combination of the two covariates (i.e., X1 and X2) and a cluster-specific random effect μ:

$$LP_{ij} = \log(1.5)x_{1ij} + \log(1.5)x_{2ij} + \mu_j$$

where the subscript *ij* denotes the ith subject in the jth cluster. The cluster-specific random effects $z_j = exp(\mu_j)$ are randomly drawn from a gamma distribution with shape parameter $\frac{1}{\theta}$ and scale parameter θ where $\theta = \frac{2\tau}{1-\tau}$. The rationale for this choice of shape and scale parameter is that Kendall's τ is used as a measure of dependence for the gamma frailty model assuming values between zero (no within-cluster correlations) and one (the highest within-cluster correlation)[2,3,5]. For a specified value of θ, we have $\tau = \frac{\theta}{\theta+2}$. Thus, we can induce data with a specific degree of within-cluster dependence for a specified value of θ. We then generated time-to-event outcomes under the Cox-Weibull model: $T_{ij} = \left(-\frac{\log(u)}{\lambda exp(LP_{ij})}\right)^{\upsilon}$ where u is drawn from a standard uniform distribution, and $\lambda = 0.01$ and $\upsilon = 0.75$. For each subject we also simulated independent censoring times from an exponential distribution. The observed survival time was defined to be the minimum of the event time and the censoring time. The rate parameter of the exponential distribution was chosen using an iterative bisection procedure so that 20% of the simulated subjects would be censored[6]. We repeated this entire process with the same cluster size and frailty terms to generate validation samples (which is equivalent to bisect a sample of $2*N_{subjects}$ per cluster for every scenario). Thus, we evaluated the predictive performances of the stratified and the frailty Cox proportional hazards models fitted using the derivation samples and evaluated in the validation samples (see the subsequent section).

After we had the simulated an event time for every scenario for each individual, we generated a super-population with 1 000 clusters and 1 000 subjects per cluster for a total of 1 000 000 subjects for each τ for a total of 4 000 000 subjects. We determined the *25th*,

50th, 75th percentiles of the event times. These will be used as the prediction horizons at which predicted event probabilities will be estimated from the stratified and the frailty Cox proportional hazards models (see the next section). We refer to these three quantities of event time as T_{25}, T_{50} and T_{75}.

For each of the 80 different scenarios, we generated 2 000 random samples (1 000 for derivation sample and 1 000 for validation sample), for a total of 160 000 random samples.

2.3 Statistical analyses

In each of the 80 000 random derivation samples (1 000 derivation samples for 80 different scenarios), we fit a stratified and a shared gamma frailty Cox proportional hazard model in which the hazard of the outcome was regressed on the two subject-level covariates including clusters as a stratified or as a frailty term, respectively. The fitted models were then applied to the validation samples and a predicted probability of the outcome at T_{25}, T_{50} and T_{75} was obtained for each individual in the corresponding validation sample.

We used the following quantitative metrics to evaluate the prediction performances of the models: for discrimination performances, the time-varying Area Under the receiver operating characteristic (ROC) Curve, which we refer to as time-varying AUC; for calibration performances, the integrated calibration index (ICI), E50, E90, and the ratio of observed-to-predicted risk[7–9]. We also computed the Index of Prediction Accuracy (IPA), which is equivalent to the scaled Brier score, at each of the prediction horizons (strictly speaking the IPA is not a measure of only discrimination, but combines discrimination and calibration; we include it here for the sake of completeness)[10]. The ICI, E50, and E90 are the mean, median, and 90th percentile, respectively, of the absolute difference between predicted survival probabilities and smoothed survival frequencies[11]. Smoothed survival frequencies were obtained using Kooperberg's flexible adaptive hazard regression model with the complementary log–log transformation of the predicted probabilities as the sole predictor variable[12]. The ratio of observed-to-expected risk was computed by dividing the observed risk of the outcome at the specified time horizons by the mean predicted risk of the outcome derived from the fitted Cox regression models. The observed risk of the outcome was determined using the Kaplan–Meier estimate of risk at the given duration of time. Thus, in each simulated validation sample, we computed the time-varying AUC, time-varying IPA, ICI, E50, E90, the ratio of observed-to-expected risk. The mean and standard

deviation of each of these measures were then computed across the 1 000 simulated validation samples for each simulation scenario. These measures represented the Monte Carlo estimates.

2.4 Software

The simulations were conducted using the R statistical programming language (version 4.3.3)[13]. The Cox model was fit using the coxph function from the survival package (version 3.8-3). The penalized partial likelihood estimation is used to estimate coefficients for the gamma frailty model through the coxph function from the survival package including the frailty term[14,15]. Observed risk for the observed-to-expected ratio was computed using the survfit function from the survival package. The predicted risk from the fitted stratified Cox model was estimated using the predictSurvProb function from the pec package (version 2025.06.24). The time-varying AUC and IPA were computed using the Score function from the riskRegression package (version 2025.05.20). The ICI, E50, and E90 were computed using the smoothed survival frequencies generated through the hare function from the polspline package (version 1.1.25). Simulation results were summarized using the simsum function from the rsimsum package (version 0.13.0).

3. Monte Carlo simulations results

We report our results separately for different time horizons T_{25}, T_{50} and T_{75} (Q1/Q2/Q3 for the three quantiles of event times in the super-population).

The simulation results are summarized in Figure 1 for T_{25}, Figure 2 for T_{50}, Figure 3 for T_{75}. Each figure consists of six panels, one for each of the performance metrics. Each panel uses a nested loop plot to report the performance metric across the different scenarios[16]. The value of the performance metric is reported on the vertical axis. At the top of each panel are three black step functions, one for each of the factors: the average number of subjects per cluster (i.e., the size of the cluster, $N_{subjects}$), the number of clusters ($N_{clusters}$), and the homogeneity within clusters τ (tau). Below these three black step functions are two step functions (in red for the stratified Cox proportional hazard model and in blue for frailty Cox proportional hazard model). For a given combination of size, number of clusters and homogeneity within clusters, one determines the vertical line that bisects each of the top three black step functions at the given combination of values. The vertical location where this vertical line bisects the step functions for the stratified and the frailty Cox model is the value of the performance metric for that model and for that

scenario. In interpreting the results, recall that higher values of AUC and IPA denote improved model performance, while lower values of ICI, E50, and E90 denote improved model performance. For the observed-to-expected ratio, a better performing model will have a ratio closer to one.

The use of the stratified Cox model resulted in predictions at T_{25} with slightly higher AUC compared to the frailty model in settings with smaller cluster size and especially low homogeneity within clusters (Figure 1 top left panel). The difference in AUC was attenuated for larger cluster size. Differences between the two models tended to be negligible with higher homogeneity within clusters and longer time horizons (Figure 1-2-3 top left panels). However, comparisons between the two models were limited since predictions from the stratified Cox model could not be estimated in scenarios with longer time horizons, higher within-cluster homogeneity, and particularly for small cluster sizes and higher number of clusters. Among the evaluable comparisons, the AUC difference between the stratified versus frailty models across all scenarios ranged from -0.002 to 0.017 with a median of 0.002 (25th and 75th percentiles: 0.001 and 0.006) (Table 1).

Both models resulted in estimates with similar IPA (Figure 1-2-3 top centre panels). The stratified model resulted in slightly higher IPA for the at T_{25} especially for small cluster size (Figure 1 top centre panel). For every scenario, the number of comparisons between the two models decreased with longer time horizons and higher homogeneity within clusters (especially for higher number of clusters with smaller size). Among the evaluable comparisons, the IPA difference between the stratified versus frailty models across all scenarios ranged from -0.011 to 0.013 with a median of 0.001 (25th and 75th percentiles: 0.000 and 0.004) (Table 1).

The frailty Cox model tended to result in the observed-to-expected ratio slightly closer to one compared with the stratified Cox model when clusters were simultaneously characterized by small size, low number of clusters, and low homogeneity especially at T_{25} (Figure 1-2-3 top right panels). The frailty Cox model tended to results in the observed-to-expected ratio closer to one compared to the stratified Cox model for most of the scenarios with small cluster size at longer time horizons (Figure 2-3 top right panels).The difference in observed-to-expected ratio was attenuated for larger cluster size For every scenario, the number of comparisons between the two models decreased with longer time horizons and higher homogeneity within clusters (especially for higher number of clusters with smaller size). Among the evaluable comparisons, for each observed-to-expected ratio we computed the absolute value of its difference from 1 and then computed the difference

between these absolute differences. The observed-to-expected differences between the stratified versus frailty models across all scenarios ranged from -0.019 to 0.028 with a median of 0.003 (25th and 75th percentiles: -0.002 and 0.012) (Table 1).

The use of the frailty Cox model tended to result in predictions that provided better calibration measures (ICI, E50, E590) compared to the stratified Cox in several scenarios with either smaller cluster size, low homogeneity within clusters, or short time horizon (Figure 1-2-3 bottom panels). The difference of the two models in terms of calibration attenuated as size and number of clusters increased especially at longer time horizons (Figure 1-2-3 bottom panels). However, for every scenario, the number of comparisons between the two models decreased with longer time horizons and higher homogeneity within clusters (especially for higher number of clusters with smaller size). Among the evaluable comparisons, the ICI difference between the stratified versus frailty models across all scenarios ranged from -0.069 to 0.050 with a median of 0.002 (25th and 75th percentile: 0.000 and 0.009). The corresponding E50 and E90 differences between the two models ranged between -0.081 and 0.046 and between -0.096 and 0.106, respectively (Table 1). The frailty Cox model tented to be more efficient than the stratified Cox at short time horizons especially for calibration measures (see supplementary). Details about comparisons between the prediction performances of the two models across all scenarios by different time horizons are provided in the supplementary.

4. Case study

We provide an empirical analysis comparing the performance of the stratified and the shared gamma frailty Cox proportional hazards models for predicting time to first insemination in dairy heifer cows coming from different herds.

4.1 Methods

We used data insem from the R package parmf described in example 1.8 by Duchateau and Janssen[2]. The dataset consisted of 10 513 dairy cows from 181 herds. In a dairy farm, the calving interval (the time between two calvings) should be optimally between 12 and 13 months. One of the main factors determining the length of the calving interval is the time from parturition to the time of first insemination. As no inseminations take place in the first 29 days after calving, authors subtracted 29 days (and not 30 days as the first event would then have first insemination time zero) from the time to first insemination since at risk time starts only then. Cows which were culled without being

inseminated were censored at their culling time. Furthermore, cows that were not yet inseminated 329 days after calving were censored at that time. A total of 9 939 cows were inseminated at the end of the follow-up during a median follow-up of 311 days.

We used the Monte Carlo cross-validation procedure to assess prediction performances of the stratified and frailty Cox model. One hundred repeated random splits were used to define the derivation and validation samples. We stratified the repeated random splits by herd (i.e., the cluster) to obtain the same herds in both the derivation and validation samples. Thus, in each herd, we split the observations in the derivation and validation samples to preserve the homogeneity within cluster on both derivation and validation samples. The splitting proportions were fixed at 70% / 30% to generate the derivation and validation samples in every cluster, respectively. We excluded herds with less than 10 cows since it was not possible to obtain sufficient cows in both derivation and validation sample. We considered two scenarios: the “small size cluster scenario” using herds with less than 50 cows per herd and the “large size cluster scenario” using herds with at least 50 cows per herd. We developed the prediction models using the stratified and the shared gamma frailty Cox models in each scenario across all 100 random derivation samples using predictors available at last parturition such as number of calvings, multiparity, milk urea and protein concentration (%) at the start of the lactation period. We predicted the insemination probability at fixed time horizons in the validation samples using the stratify and frailty models developed using the derivation samples. The time horizons were defined by the 25th, 50th and 75th percentile of the marginal insemination times using the Kaplan-Meier estimator which corresponded to 51, 91 and 150 days, respectively. We refer to these three quantities of event time as T_{25}, T_{50} and T_{75}. We assessed the prediction performances of each model in two scenarios in the 100 random validation samples at the three time horizons using the following metrics: time-varying AUC, time-varying IPA, observed-to-expected ratio, ICI, E50 and E90. The mean of the performances metrics across the 100 random validation samples with the corresponding standard deviation were reported for both scenarios and at each prediction time horizon to assess the prediction performances of the stratified and the frailty Cox models. The accompanying R code for the case study is available in the supplementary and at https://github.com/danielegiardiello/ClusterSurvPred

4.2 Results

Results are reported in Table 2. The results suggest very similar prediction performances in both scenarios. All performance measures were on average slightly better for the frailty model. Not all the repeated random splits were usable for the stratified models: no predictions from the stratified model were possible at T_{75} for both small and large size scenarios. For the small size cluster scenario, 95% and 78% repeated random splits were usable at T_{25} and T_{50}, respectively, while 97 out of 100 random splits were usable for the large size cluster scenario at T_{50}.

5. Discussion

Large electronic health records from different data sources have become more common in the last decades. The use of individual data from multiple sources, also better known in medical research as individual patient meta-analysis (IPD-MA), are more commonly encouraged and used to develop and evaluate risk prediction models considering data from different studies[17–19]. In this work we have investigated, using simulations, the impact of the size, number of clusters and within-cluster homogeneity on risk prediction performances of the Cox proportional hazards models for observations nested within a cluster such as a study centre. We summarize our findings as follows: first, the stratified Cox model showed slightly better discrimination compared to the frailty Cox model where the number and the size of the clusters and homogeneity within clusters were small. AUC were very similar as the number and size of clusters increased. Second, IPA very similar in most of the scenarios with a slightly better overall performance for small cluster size at short time horizons for the stratified model. Third, the frailty Cox model tended to have better calibration than the stratified Cox for low number of clusters with small size and low homogeneity within clusters at short time horizons. The calibration tented to be similar in both models as size and number of clusters increased especially with longer time horizons. Across all scenarios, the stratified and frailty Cox had generally similar performances.

To the best of our knowledge, this is the first study that investigates the impact of number, size and homogeneity within clusters on the prediction performance of Cox proportional hazards models for censored time-to-event clustered data using simulations. The Cox proportional hazards model is the most popular regression model for time-to-event outcomes. The Cox proportional hazards model may manage clusters non parametrically using stratification or parametrically using the frailty (i.e., random effect) term. Thus, the stratified approach makes less assumptions compared to the frailty

approach, although the former approach can be less efficient and less useful when the size of the cluster is small[20]. To our knowledge, other works primarily focused on the effect of the intervention in multicentre (randomized) clinical trials[21–23]. Glidden and Vittinghoff investigated how to incorporate centres effects into the analysis of multicentre clinical trials for censored time-to-event outcomes[21]. They found substantial advantages of the gamma frailty Cox model compared to the fixed and stratified Cox model. Munda and Legrand extended the findings of Glidden and Vittinghoff also when the frailty distribution is misspecified[22]. de Jong et al. focused on the one stage approach (i.e., all subjects from different studies are analysed in a single step) and two stage approach (i.e., subjects are analysed separately in each study and then results are aggregated) for IPD-MA of intervention studies with time-to-event outcomes[23–25]. The authors suggested similar conclusions with previous works. However, less is known about the impact of prediction performances of risk prediction models for survival outcomes for clustered data. In the one stage approach of IPD-MA for intervention and prognostic studies with time-to-event outcomes, Riley and colleagues stated "It is fundamental to account for clustering using either stratified or frailty models […]. Frailty models are perhaps best when the number of participants and outcome events per trial are small, given the estimation issues may arise with separate baseline hazards per trial"[19,26]. With this simulation study, we aimed to give further insights for applied researchers. We essentially confirmed the statement of Riley and colleagues and the literature's findings about the multicentre interventional studies[19,21–23]. Our simulation study suggested that prediction performance was generally similar across most of the investigated scenarios. For small cluster size with higher number of clusters and moderate to high homogeneity within clusters, the frailty and stratified model tended to results with similar prediction performances with slightly better discrimination and overall measures for the stratified model. The frailty model tended to result in better calibration for small cluster size especially for mild within-cluster homogeneity and at short time horizons. However, the practical utility of the stratified Cox proportional hazards model is limited since the baseline hazard in every cluster and the corresponding estimated risk predictions cannot be always estimated for a considerable heterogeneous number of clusters with small size and in particular at longer time horizons.

Our study has also several limitations. Firstly, our study relies on a limited Monte Carlo simulations' scenarios mainly due to computationally intensive simulations. On the other hand, we provided results from 80 scenarios for three time horizons for a total of 240 scenarios. Secondly, we generated times using the Cox-Weibull framework[4]. Different

models could have been investigated to generate outcomes. Thirdly, our work focused exclusively on predictions in the same clusters used to develop the risk prediction model. Therefore, our study did not consider marginal prediction performance assessment of the marginal gamma frailty Cox model. It is formally not possible to provide marginal predictions from the stratified Cox model, unless an appropriate baseline hazard could then be selected from existing studies in the meta-analysis using the incidence in the new study population[18]. In the context of IPD-MA, the prediction performance assessment is generally evaluated using the internal-external cross validation[18,27].

An important limitation of the stratified Cox model was the impossibility to estimate the predictions for all clusters. Predictions from the stratified Cox model could not be estimated for clusters in which subjects with times in the validation sample were greater than the maximum event time in the derivation sample. For the subjects across those clusters, the corresponding baseline hazard for that cluster could not be estimated. Therefore, this substantially limits the utility of the stratified Cox compared to the frailty Cox model. In our simulations, the number of usable comparisons between the two models decreased substantially for longer time horizons and higher homogeneity within clusters (especially among scenarios with higher number of clusters with small size). The scenario with the least relative utility of the comparisons between the two models was observed for a homogeneity within clusters of 0.10 with 8 clusters and with an average number of 25 patients per cluster where only 14.1% of comparisons between the two models were possible (details available in the supplementary). In the frailty Cox model, a shared baseline hazard is estimated across clusters which enables to provide predictions over a considerably larger range of reasonable potential prediction time horizons.

In conclusion, the stratified and frailty Cox proportional hazards models generally provide similar predictions performances in most of the simulated scenarios with different number and sizes of clusters. When the cluster size is small, the stratified Cox proportional hazard model provides slightly better discrimination but worse calibration compared to the shared gamma frailty Cox proportional hazards model especially at short time horizons and when homogeneity within clusters is mild. The practical applicability of the stratified Cox proportional hazards model to estimate predictions is limited especially for high heterogeneity between clusters and when clusters are small, and more likely at longer time prediction horizons. The shared gamma frailty Cox proportional hazards model can provide predictions in case of small and heterogeneous clusters. It also enables marginal predictions for individuals from new or previously unobserved clusters.

Acknowledgments

We thank Stefano Bottelli for his suggestions on this manuscript.

Funding

Daniele Giardiello is funded by the National Plan for NRRP Complementary Investments (PNC, established with the decree-law 6 May 2021, n. 59, converted by law n. 101 of 2021) in the call for the funding of research initiatives for technologies and innovative trajectories in the health and care sectors (Directorial Decree n. 931 of 06-06-2022) - project n. PNC0000003 - AdvaNced Technologies for Human-centrEd Medicine (project acronym: ANTHEM).

Edoardo Ratti is partially supported by the grant: Italian MUR Dipartimenti di Eccellenza 2023-2027 (l.232/2016, art. 1, commi 314-337).

ICES is an independent, non-profit research institute funded by an annual grant from the Ontario Ministry of Health (MOH) and the Ministry of Long-Term Care (MLTC). As a prescribed entity under Ontario's privacy legislation, ICES is authorized to collect and use health care data for the purposes of health system analysis, evaluation and decision support. Secure access to these data is governed by policies and procedures that are approved by the Information and Privacy Commissioner of Ontario. The use of the data in this project is authorized under section 45 of Ontario's Personal Health Information Protection Act (PHIPA) and does not require review by a Research Ethics Board. This study was supported by ICES, which is funded by an annual grant from the Ontario Ministry of Health (MOH) and the Ministry of Long-Term Care (MLTC). This study also received funding from the Canadian Institutes of Health Research (CIHR) (PJT 166161). This document used data adapted from the Statistics Canada Postal CodeOM Conversion File, which is based on data licensed from Canada Post Corporation, and/or data adapted from the Ontario Ministry of Health Postal Code Conversion File, which contains data copied under license from ©Canada Post Corporation and Statistics Canada. Parts of this material are based on data and/or information compiled and provided by CIHI and the Ontario Ministry of Health. The analyses, conclusions, opinions and statements expressed herein are solely those of the authors and do not reflect those of the funding or data sources; no endorsement is intended or should be inferred.

Conflicts of Interest

The authors declare no conflicts of interest.

References

1. Cox, DR. Regression Models and Life Tables (With Discussion). *Journal of the Royal Statistical Society, Series B (Statistical Methodology)*. 1972;34:187-202.

2. Duchateau, L, Janssen, P. *The Frailty Model*. First. Springer New York, NY; 2008.

3. Wienke A. *Frailty Model in Survival Analysis*. First. Chapman & Hall/CRC Press; 2011.

4. Bender R, Augustin T, Blettner M. Generating survival times to simulate Cox proportional hazards models. *Statistics in Medicine*. 2005;24.

5. Hougaard, P. *Analysis of Multivariate Survival Data*. 1st ed. Springer New York, NY; 2000.

6. Austin PC. The iterative bisection procedure: a useful tool for determining parameter values in data-generating processes in Monte Carlo simulations. *BMC Medical Research Methodology*. 2023;23(1):45.

7. Gerds TA, Kattan MW. *Medical Risk Prediction Models: With Ties to Machine Learning*. 1st ed. Chapman & Hall/CRC Press; 2022.

8. McLernon DJ, Giardiello D, Van Calster B, et al. Assessing Performance and Clinical Usefulness in Prediction Models With Survival Outcomes: Practical Guidance for Cox Proportional Hazards Models. *Ann Intern Med*. 2023;176(1):105-114.

9. Blanche P, Kattan MW, Gerds TA. The c-index is not proper for the evaluation of t-year predicted risks. *Biostatistics*. 2019;20(2):347-357.

10. Kattan MW, Gerds TA. The index of prediction accuracy: an intuitive measure useful for evaluating risk prediction models. *Diagnostic and Prognostic Research*. 2018;2(1):7.

11. Austin, PC, Harrell, FE, van Klaveren D. Graphical calibration curves and the integrated calibration index (ICI) for survival models. *Stat Med*. 2020;39(21):2714-2742.

12. Kooperberg C, Stone CJ, Truong YK. Hazard Regression. *Journal of the American Statistical Association*. Published online 1995.

13. R Core Team. R: A language and environment for statistical computing.

14. Therneau, MT, Grambsch, PM. Frailty Models (chapter 9). In: *Modeling Survival Data - Extending the Cox Model*. Statistics for Biology and Health. Springer New York, NY.

15. Therneau TM, Grambsch PM, Pankratz VS. Penalized Survival Models and Frailty. *Journal of Computational and Graphical Statistics*. 2003;12(1):156-175.

16. Rücker G, Schwarzer G. Presenting simulation results in a nested loop plot. *BMC Medical Research Methodology*. 2014;14(1):129.

17. Ahmed I, Debray TP, Moons KG, Riley RD. Developing and validating risk prediction models in an individual participant data meta-analysis. *BMC Medical Research Methodology*. 2014;14(1):3.

18. Debray TPA, Moons KGM, Ahmed I, Koffijberg H, Riley RD. A framework for developing, implementing, and evaluating clinical prediction models in an individual participant data meta-analysis. *Stat Med*. 2013;32(18):3158-3180.

19. Riley, RD, Tierney, J, Stewart, LA (Eds). *Individual Participant Data Meta-Analysis: A Handbook for Healthcare Research*. John Wiley & Sons Ltd; 2021.

20. O'Quigley, J. *Proportional Hazards Regression*. Springer New York, NY; 2008.

21. Glidden DV, Vittinghoff E. Modelling clustered survival data from multicentre clinical trials. *Stat Med*. 2004;23(3):369-388.

22. Munda M, Legrand C. Adjusting for centre heterogeneity in multicentre clinical trials with a time to event outcome. *Pharmaceutical Statistics*. 2014;13(2):145-152.

23. de Jong VMT, Moons KGM, Riley RD, et al. Individual participant data meta-analysis of intervention studies with time-to-event outcomes: A review of the methodology and an applied example. *Research Synthesis Methods*. 2020;11(2):148-168.

24. Burke DL, Ensor J, Riley RD. Meta-analysis using individual participant data: one-stage and two-stage approaches, and why they may differ. *Stat Med*. 2017;36(5):855-875.

25. Riley RD, Ensor J, Hattle M, Papadimitropoulou K, Morris TP. Two-stage or not two-stage? That is the question for IPD meta-analysis projects. *Research Synthesis Methods*. 2023;14(6):903-910.

26. Riley, RD, Debray, TPA. The One-stage Approach to IPD Meta-analysis (chapter 6, section 6.3). In: *Individual Participant Data Meta-Analysis: A Handbook for Healthcare Research*. John Wiley & Sons Ltd; 2021.

27. Steyerberg EW, Harrell FE. Prediction models need appropriate internal, internal-external, and external validation. *J Clin Epidemiol*. 2016;69:245-247.

Table 1: comparison of stratified versus frailty Cox proportional hazards models on prediction performance metrics across all comparable simulation scenarios and across all time horizons

Performance metric	Comparison	Minimum	25th percentile	50th percentile	75th percentile	Maximum
AUC	Stratify vs frailty	-0.002	0.001	0.002	0.006	0.017
IPA	Stratify vs frailty	-0.011	0.000	0.001	0.004	0.013
O-E ratio	Stratify vs frailty	-0.019	-0.002	0.003	0.012	0.028
ICI	Stratify vs frailty	-0.069	0.000	0.002	0.009	0.050
E50	Stratify vs frailty	-0.081	0.000	0.003	0.008	0.046
E90	Stratify vs frailty	-0.096	-0.001	0.005	0.017	0.106

AUC: area Under the receiver operating characteristic curve Curve; O-E: observed-to-expected; ICI: Integrated Calibration Index

Table 2: accuracy of the predictions for the case study by scenario, model and time horizon

Scenario	Model	Time horizon	AUC (sd)	IPA (sd)	O-E ratio (sd)	ICI (sd)	E50 (sd)	E90 (sd)	% usable splits
Small size cluster	Stratified	T_{25}	0.658 (0.015)	0.039 (0.021)	1.024 (0.087)	0.037 (0.011)	0.032 (0.013)	0.071 (0.013)	95%
	Frailty	T_{25}	0.688 (0.015)	0.075 (0.016)	0.994 (0.055)	0.025 (0.013)	0.023 (0.013)	0.046 (0.013)	100%
	Stratified	T_{50}	0.723 (0.016)	0.146 (0.026)	1.023 (0.045)	0.037 (0.014)	0.033 (0.014)	0.064 (0.014)	78%
	Frailty	T_{50}	0.730 (0.016)	0.161 (0.024)	0.999 (0.059)	0.022 (0.011)	0.021 (0.011)	0.038 (0.011)	100%
	Stratified	T_{75}	-	-	-	-	-	-	0%
	Frailty	T_{75}	0.813 (0.016)	0.262 (0.033)	0.987 (0.056)	0.022 (0.011)	0.020 (0.011)	0.039 (0.011)	100%
Large size cluster	Stratified	T_{25}	0.656 (0.009)	0.055 (0.009)	1.015 (0.045)	0.026 (0.047)	0.023 (0.039)	0.048 (0.087)	100%
	Frailty	T_{25}	0.670 (0.010)	0.067 (0.009)	1.005 (0.028)	0.013 (0.014)	0.012 (0.008)	0.023 (0.021)	100%
	Stratified	T_{50}	0.701 (0.008)	0.124 (0.011)	1.011 (0.024)	0.020 (0.009)	0.019 (0.010)	0.036 (0.018)	97%
	Frailty	T_{50}	0.706 (0.008)	0.130 (0.011)	1.001 (0.029)	0.013 (0.018)	0.011 (0.006)	0.024 (0.041)	100%
	Stratified	T_{75}	-	-	-	-	-	-	0%
	Frailty	T_{75}	0.776 (0.010)	0.201 (0.017)	0.998 (0.030)	0.015 (0.023)	0.012 (0.007)	0.028 (0.069)	100%

sd: standard deviation; AUC: area Under the receiver operating characteristic curve Curve; O-E: observed-to-expected; ICI: Integrated Calibration Index

Figure 1. Prediction performances at T_{25}

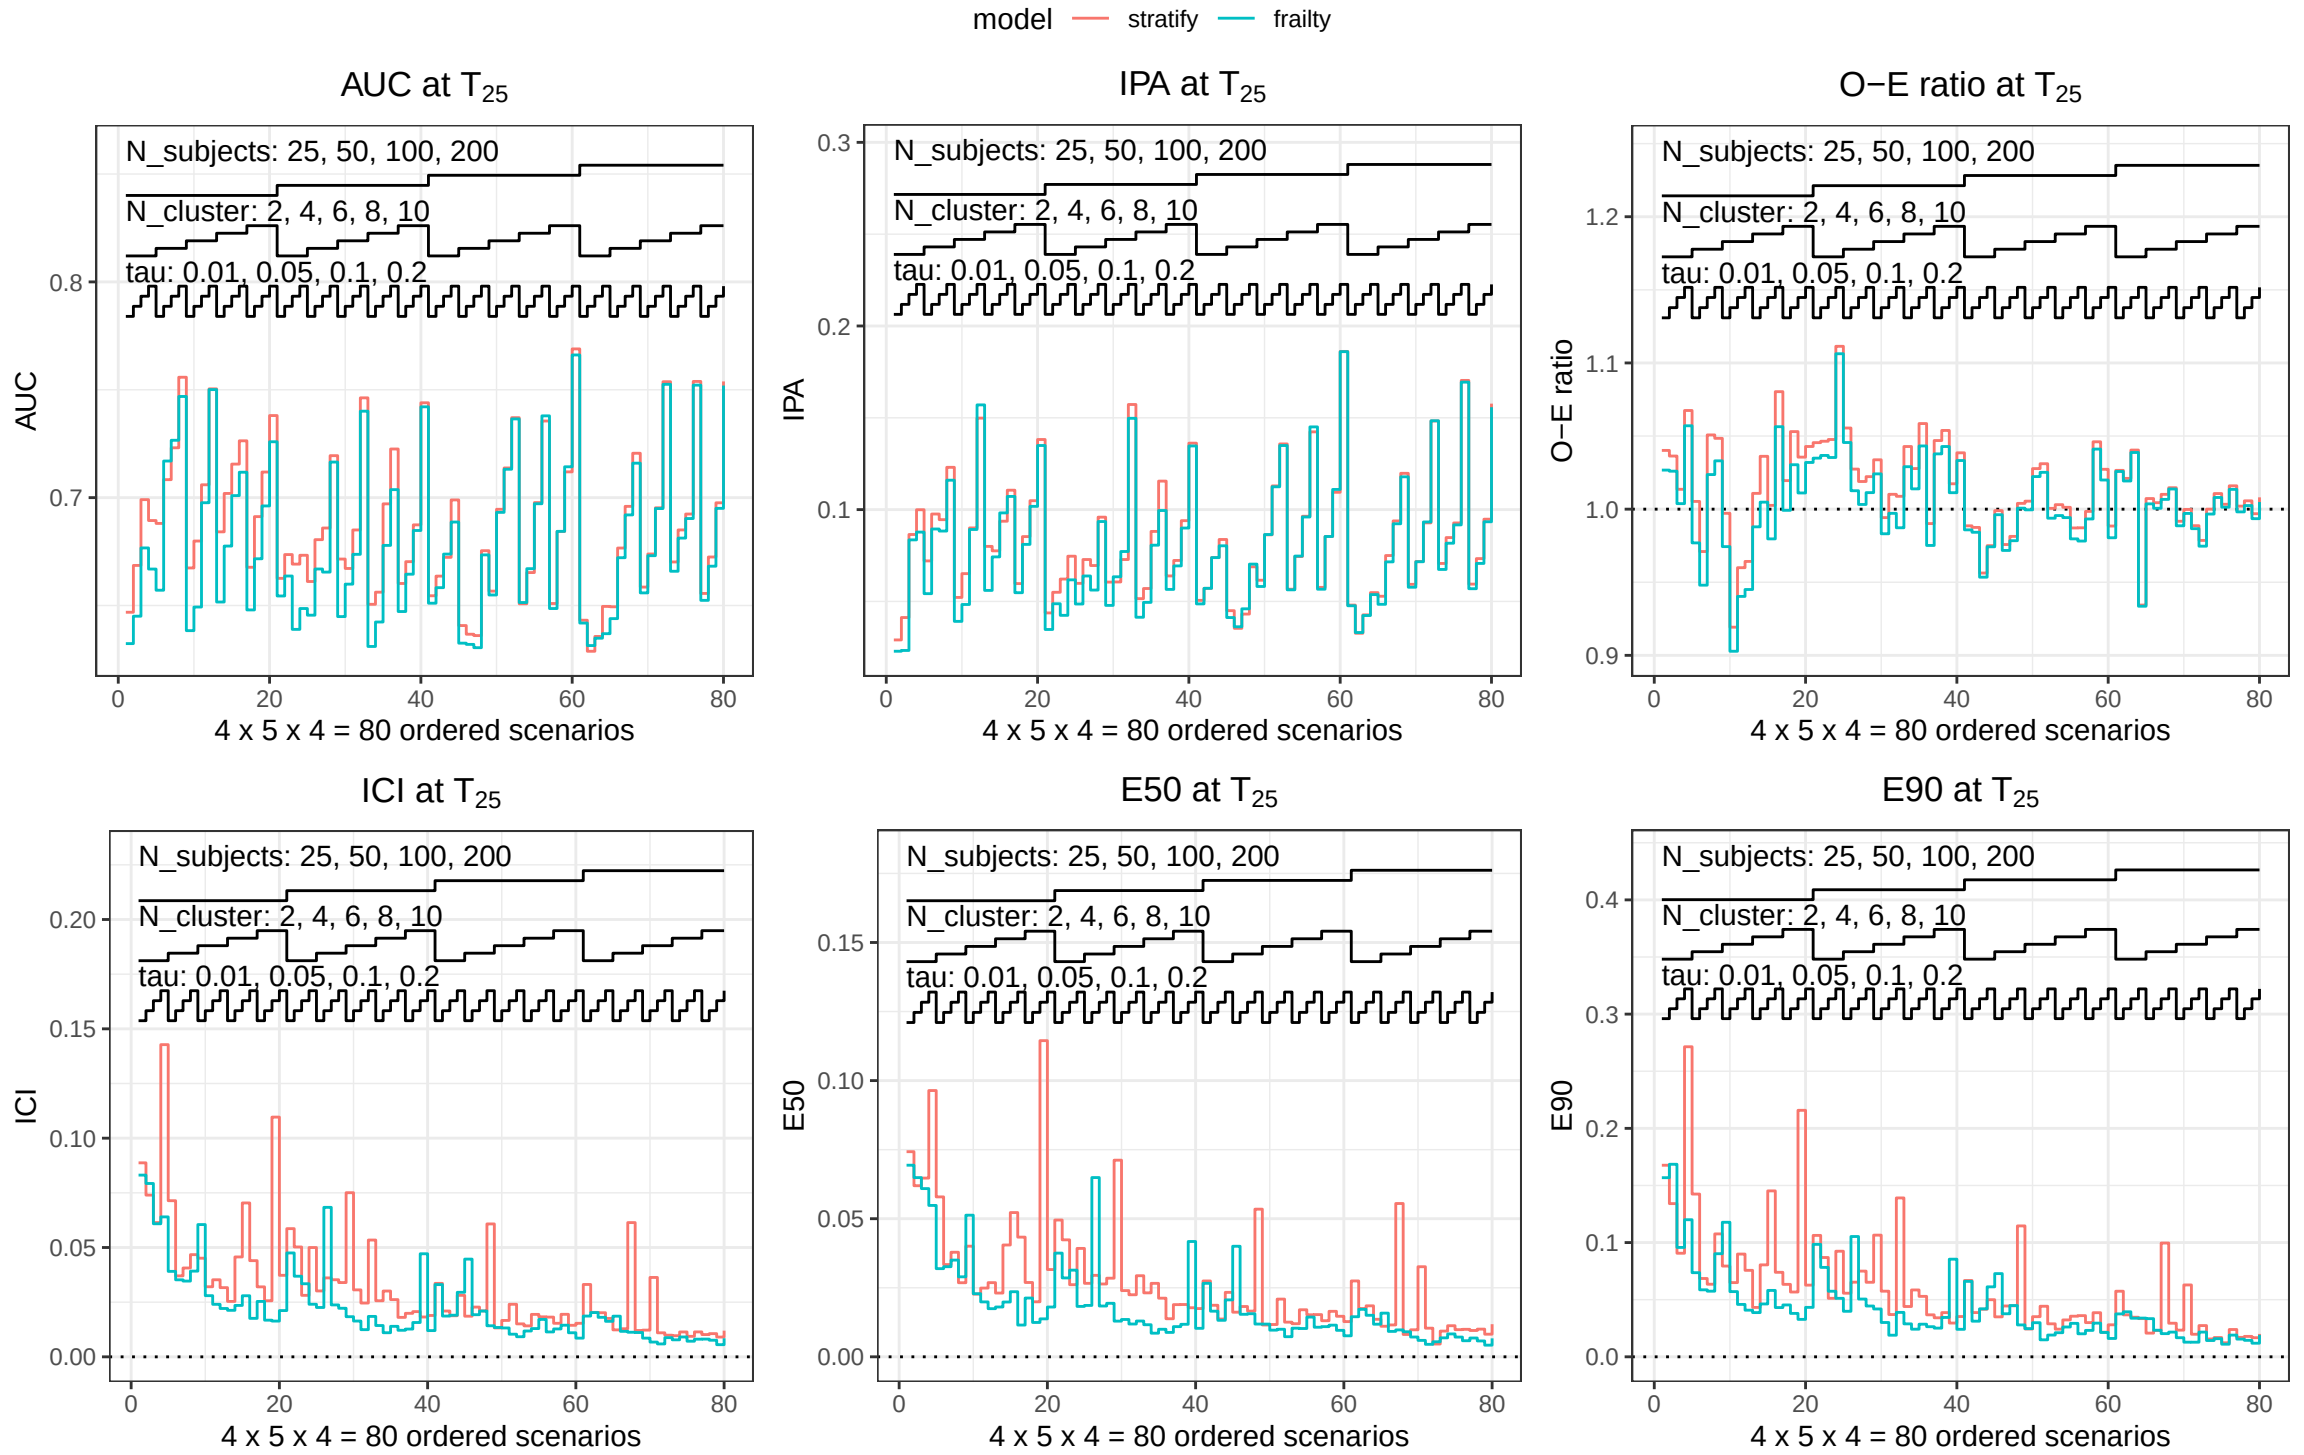

Figure 2. Prediction performances at T_{50}

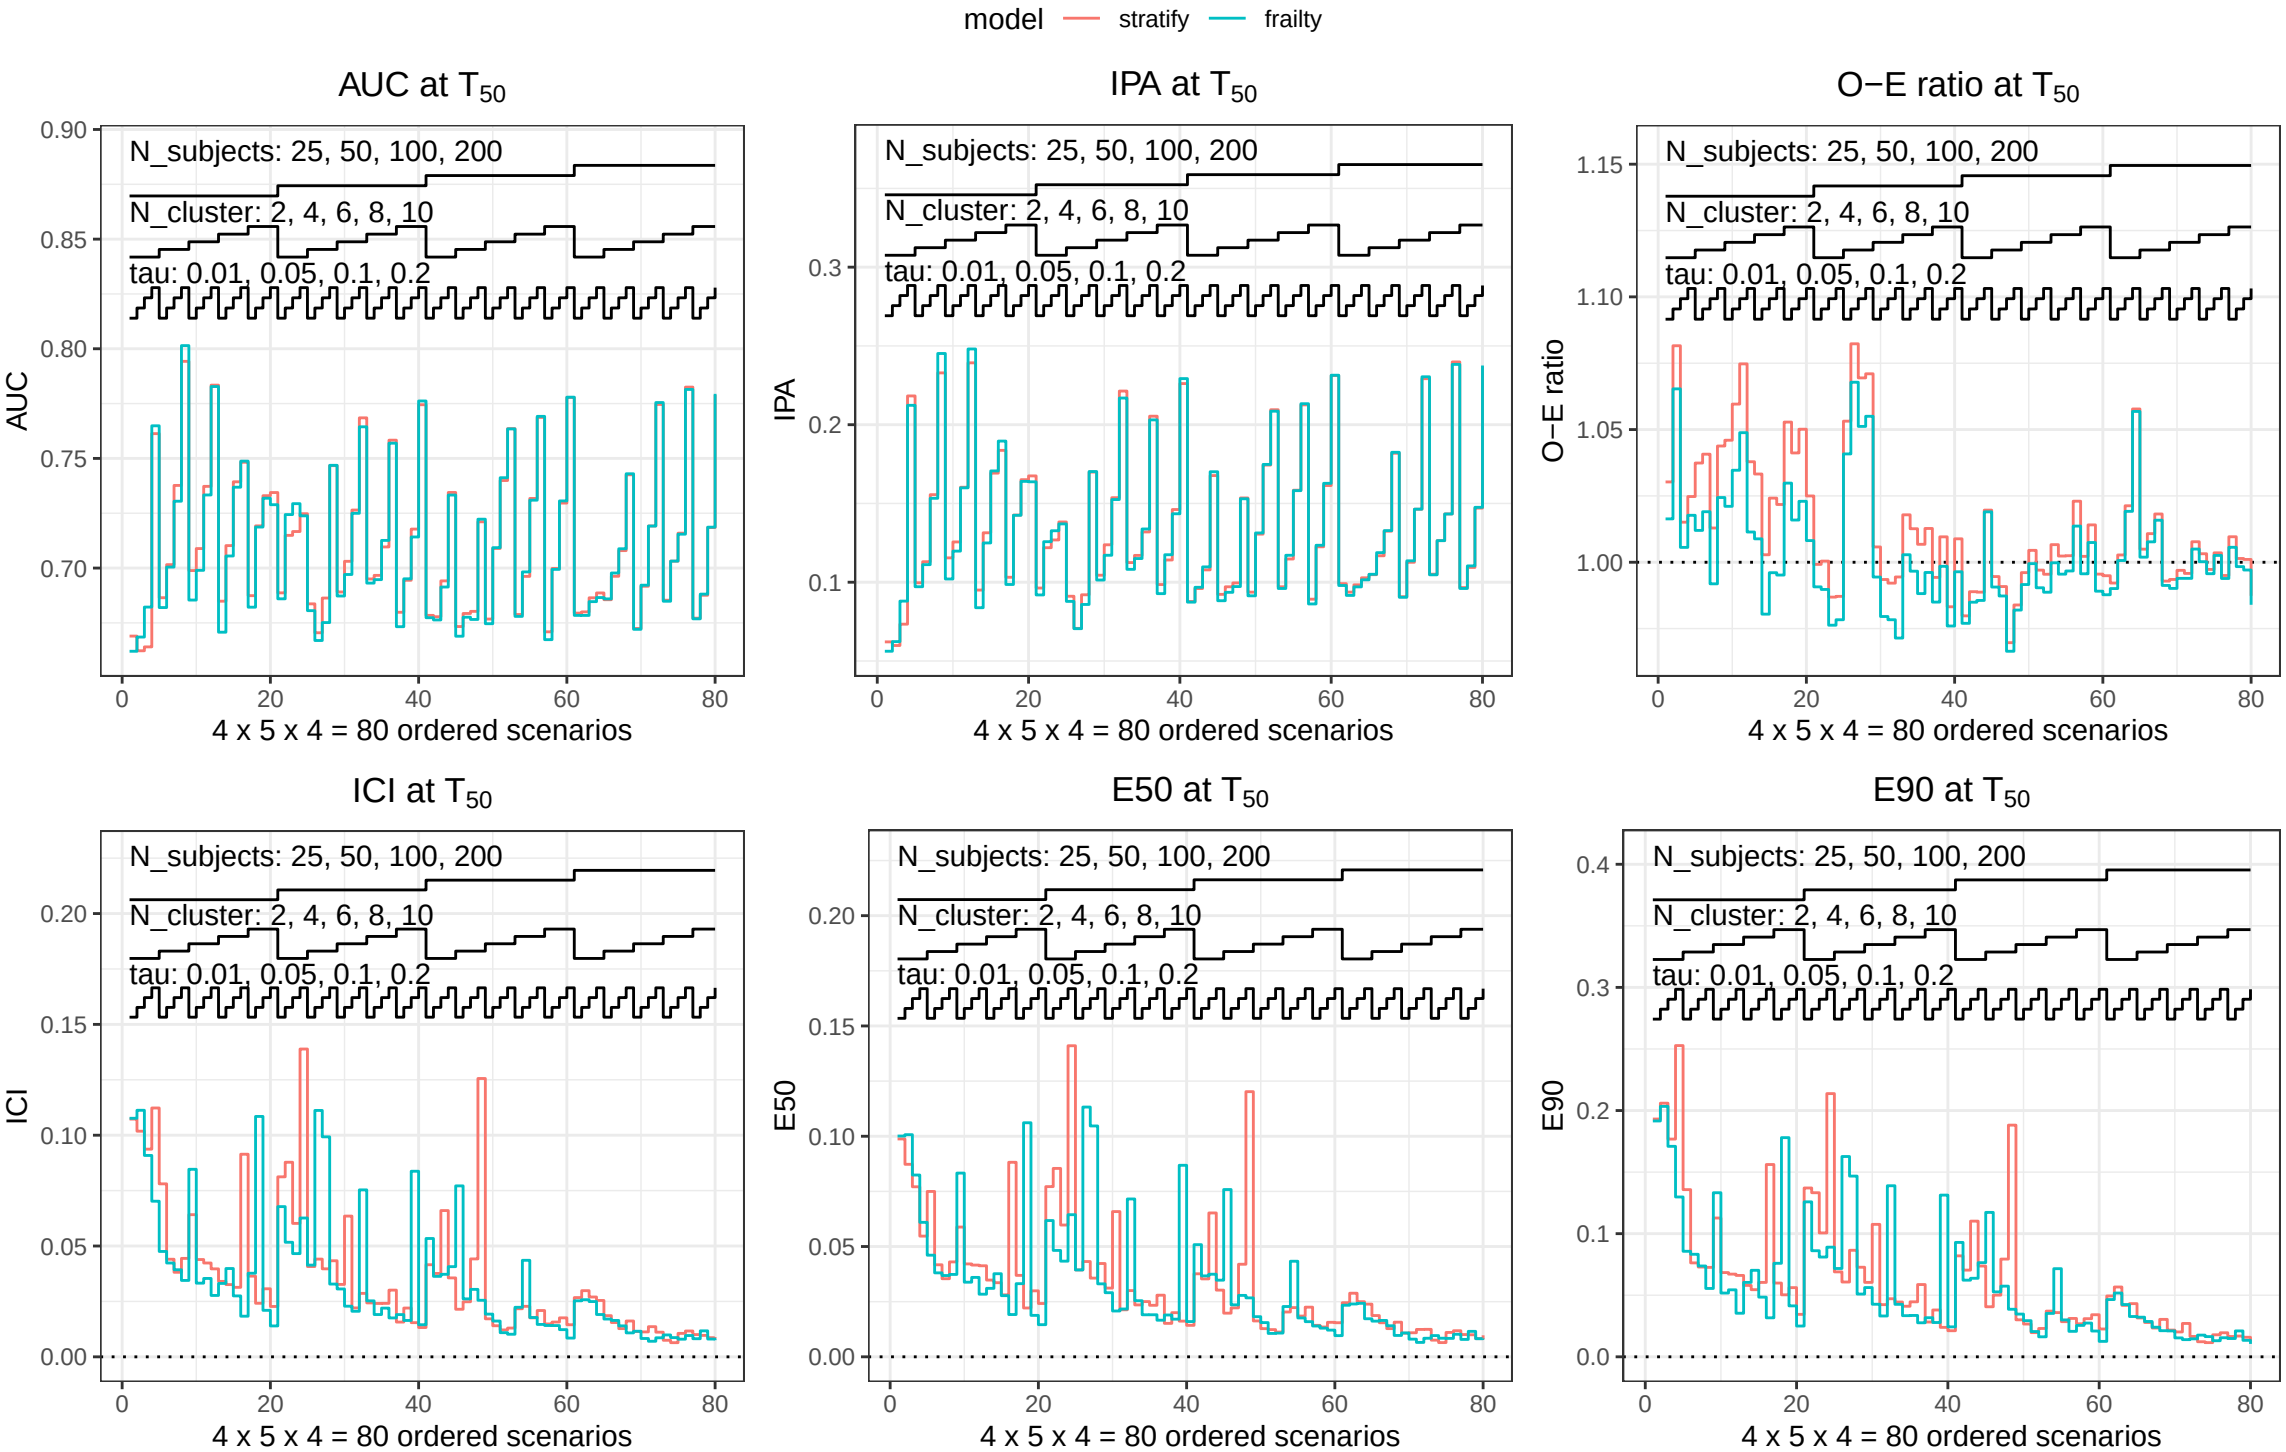

Figure 3. Prediction performances at T_{75}

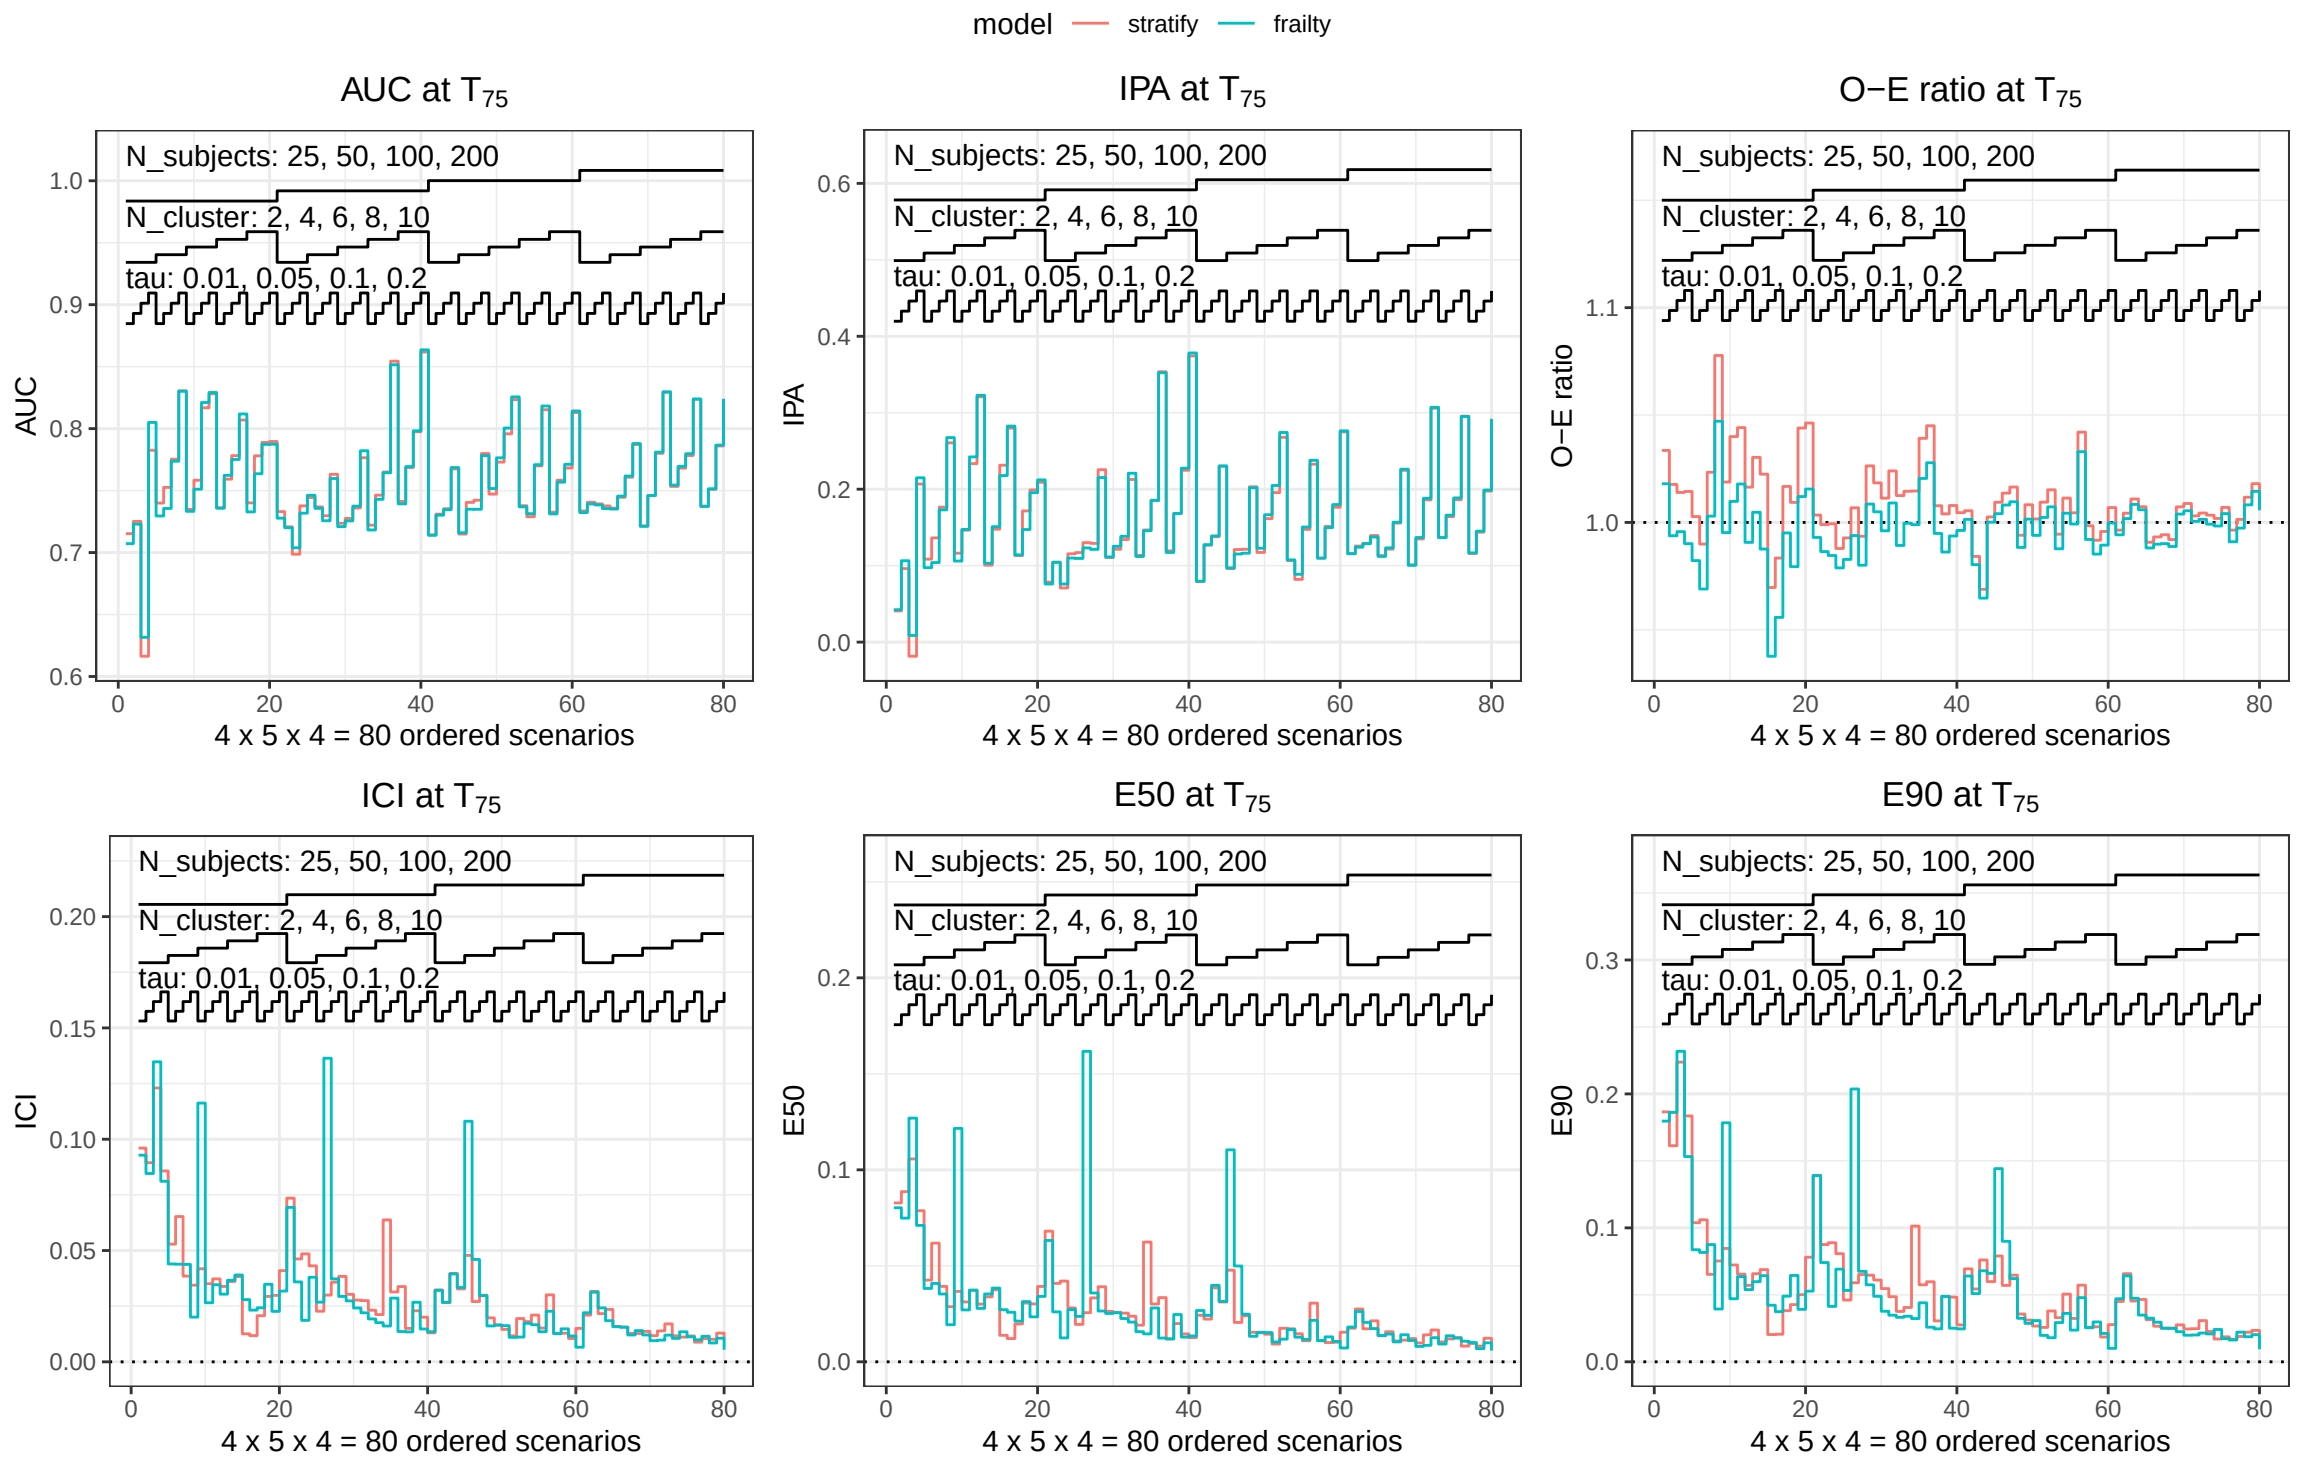

Figure S1. SE of the prediction performances at T_{25}

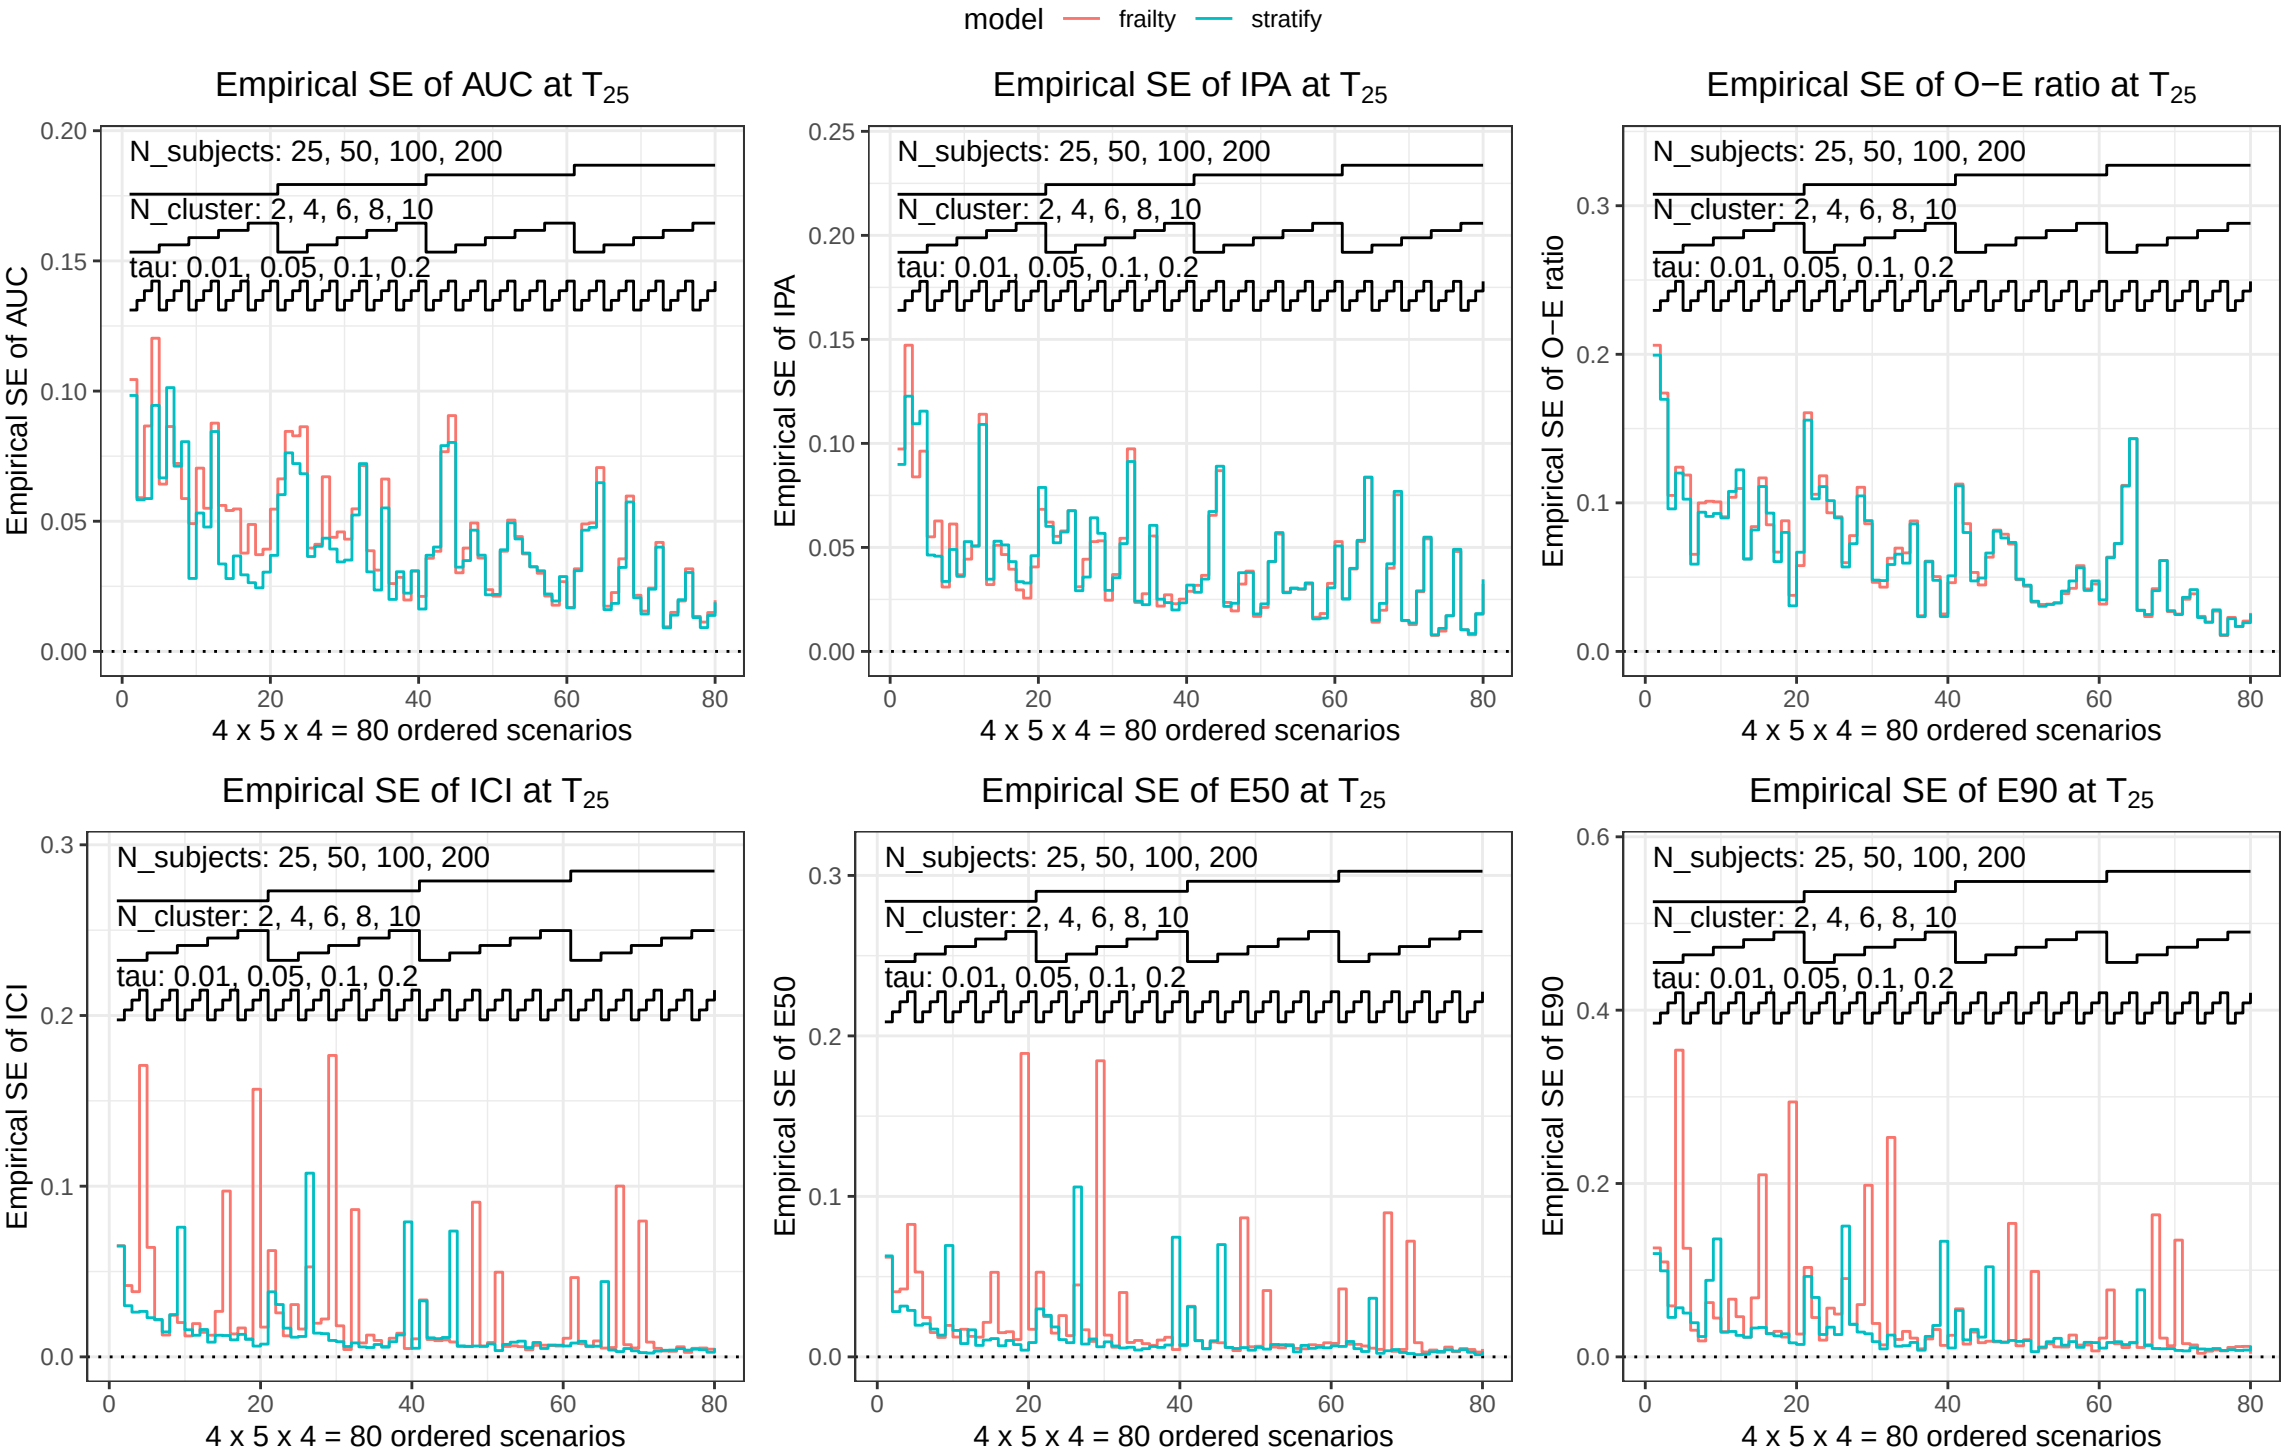

Figure S2. SE of the prediction performances at T_{50}

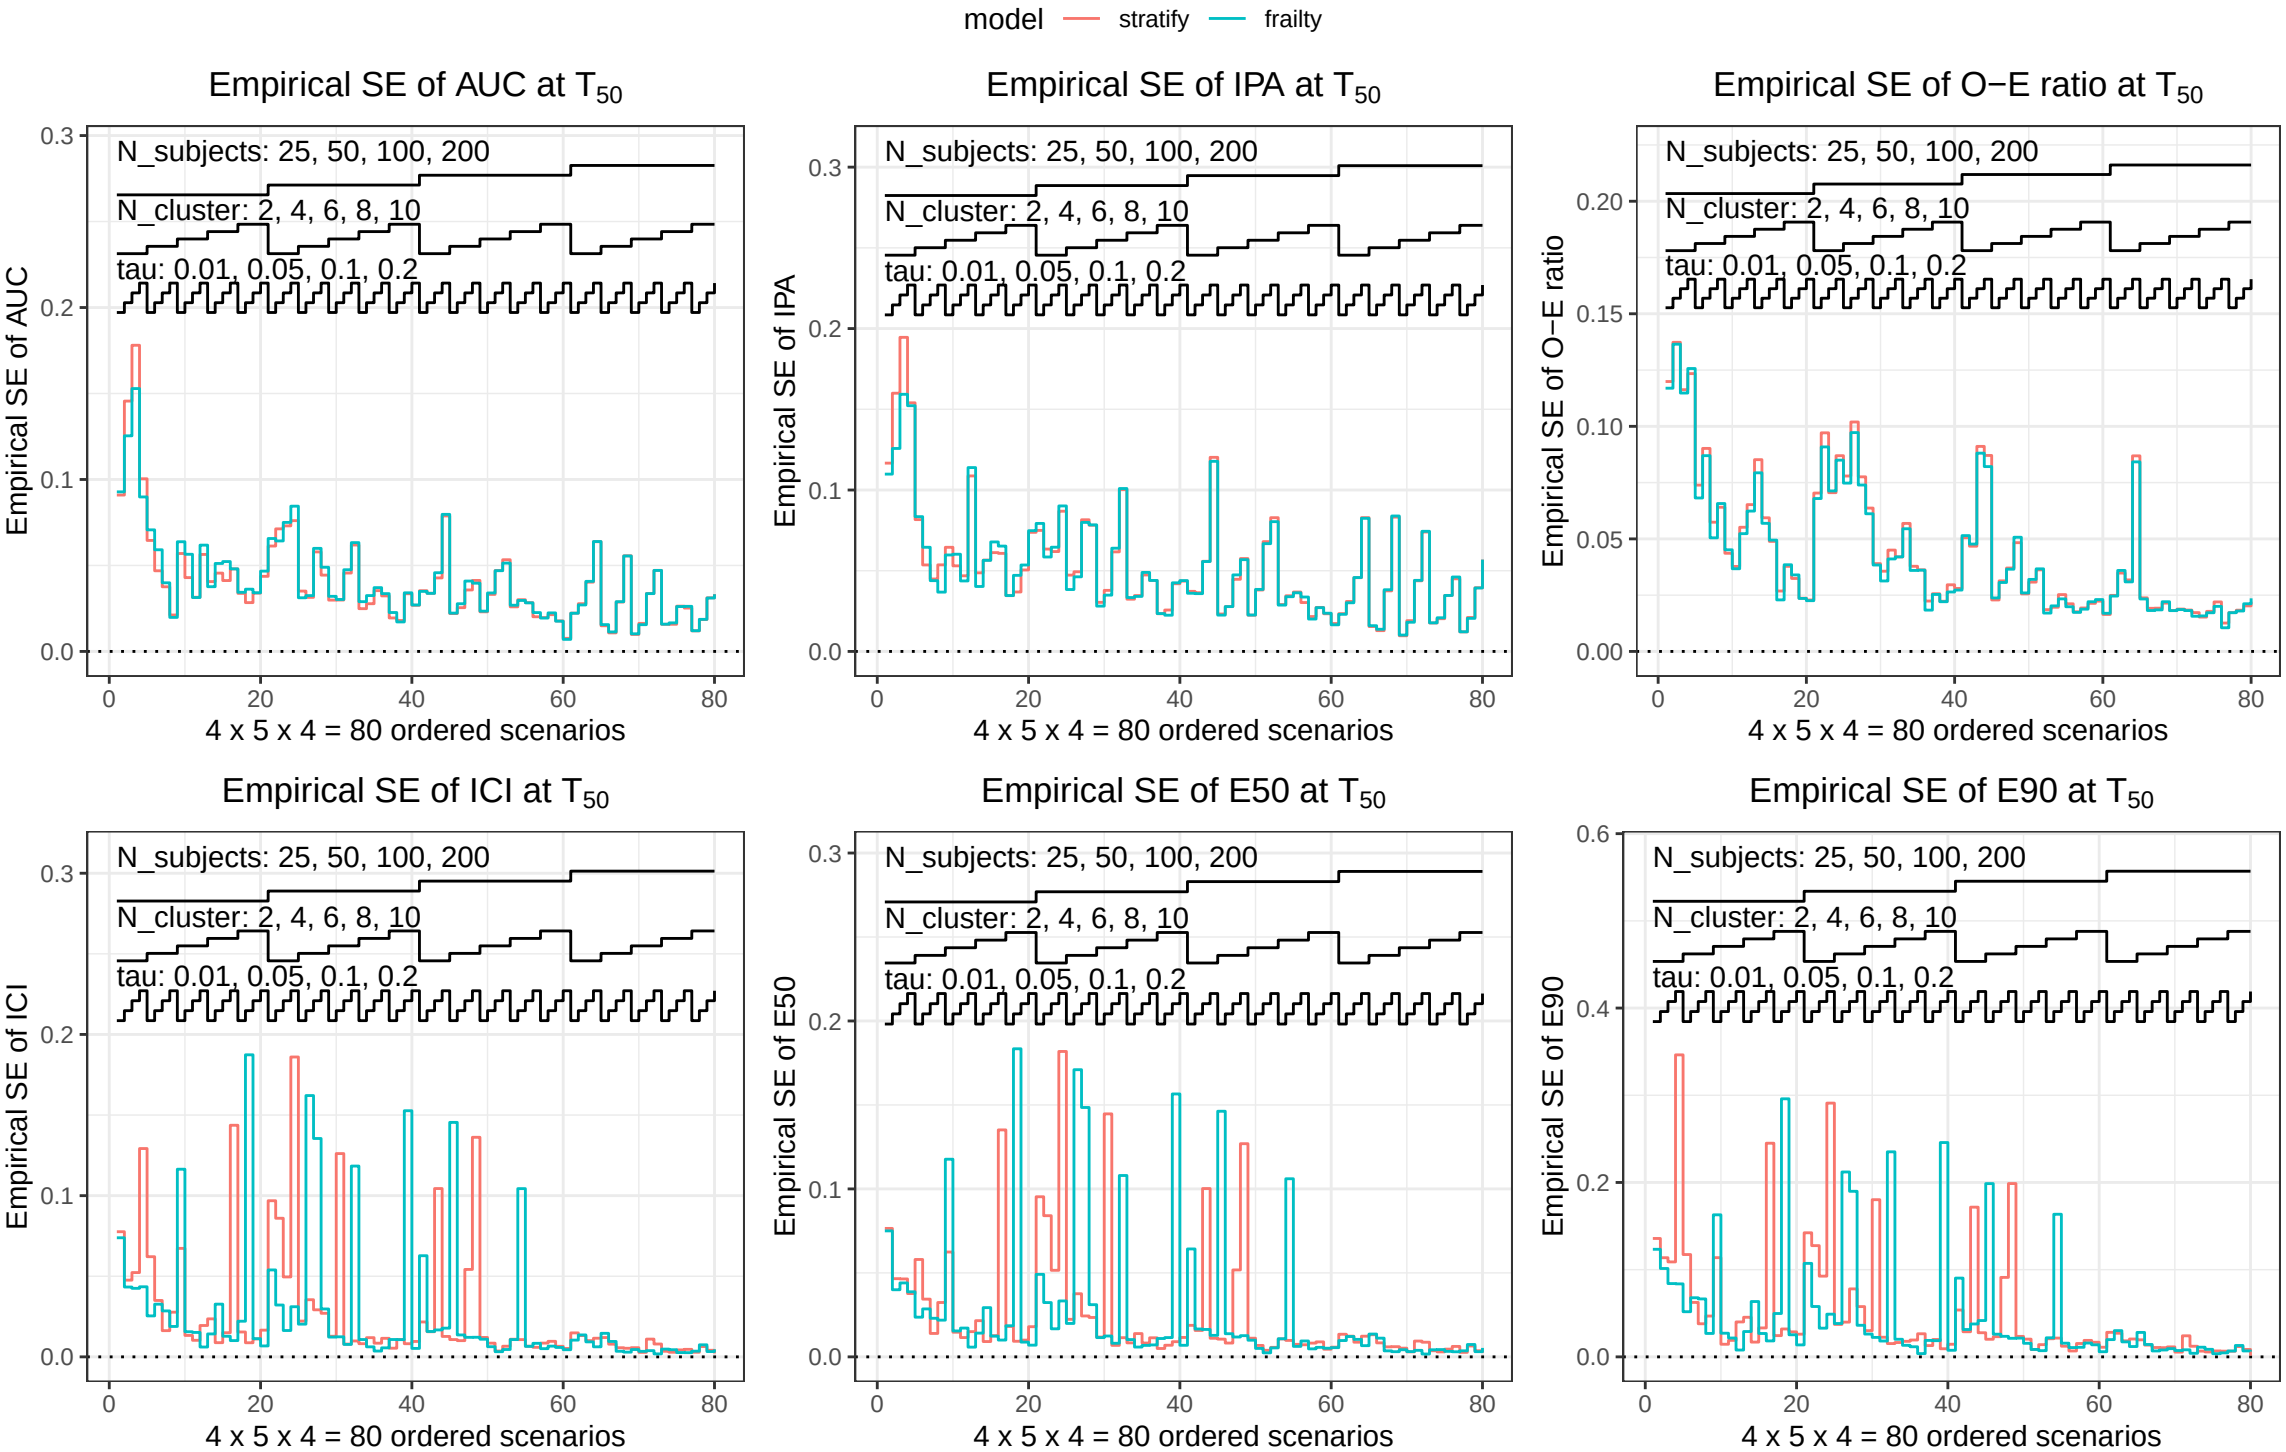
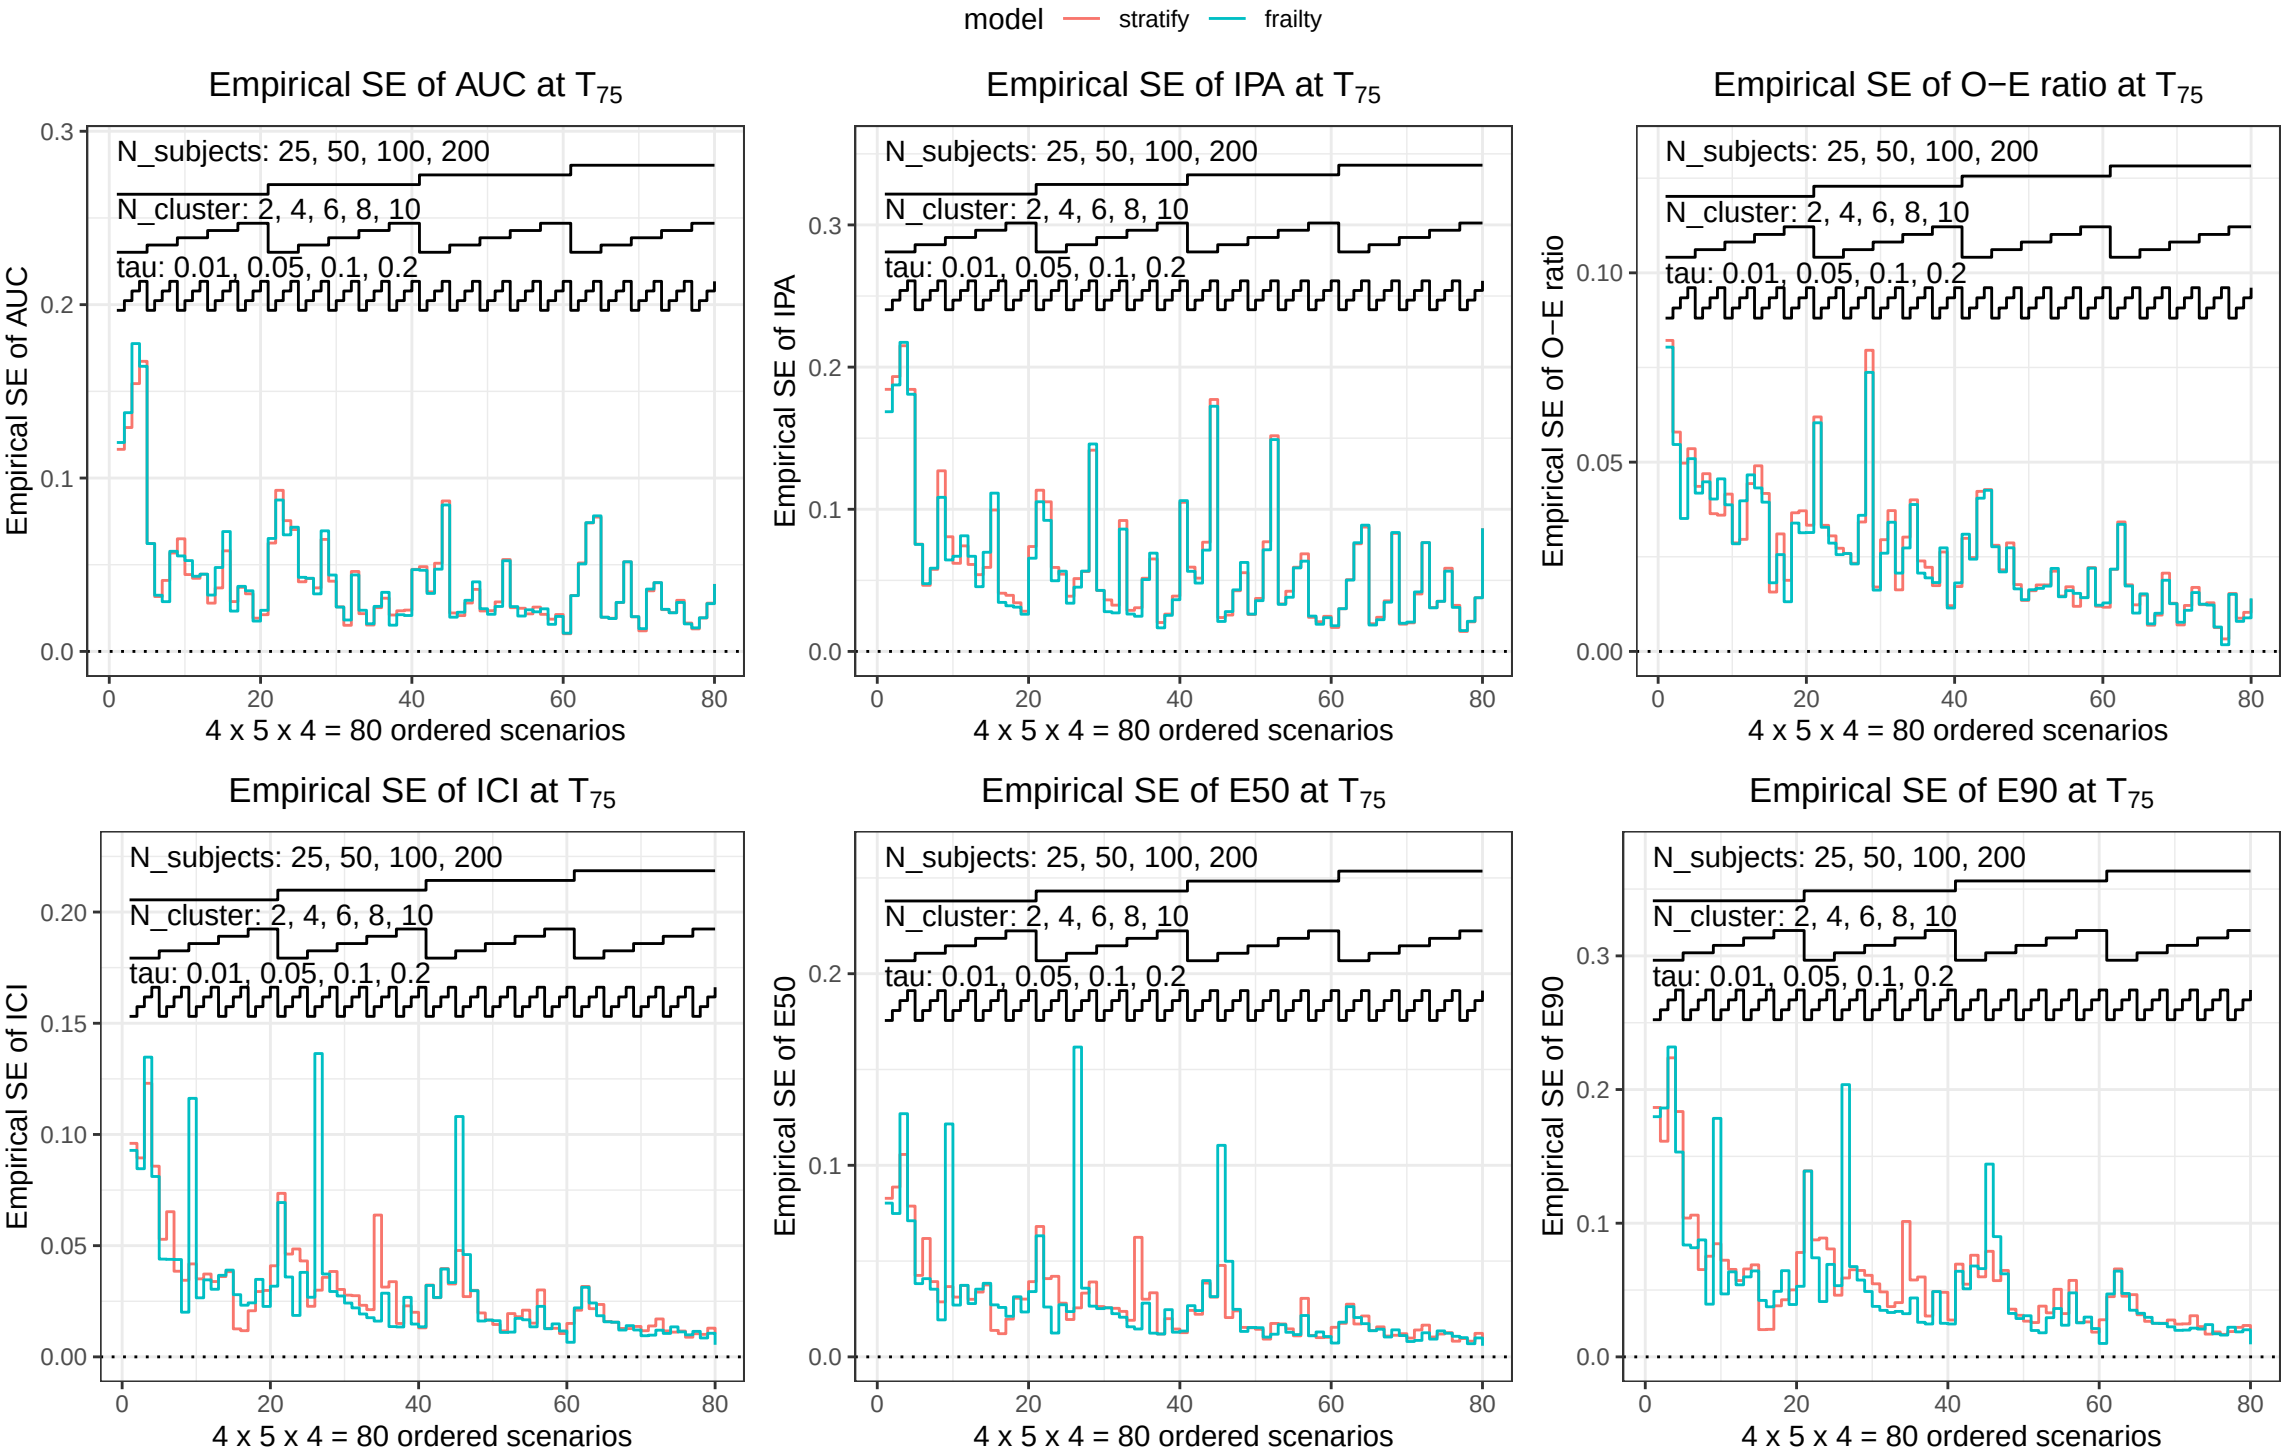
Figure S3. SE of the prediction performances at T_{75}