

Conformal Survival Bands for Risk Screening under Right-Censoring

Matteo Sesia

University of Southern California, Los Angeles, California, USA

Vladimir Svetnik

Merck & Co., Inc., Rahway, New Jersey, USA

SESIA@MARSHALL.USC.EDU

VLADIMIR_SVETNIK@MERCK.COM

Abstract

We propose a method to quantify uncertainty around individual survival distribution estimates using right-censored data, compatible with any survival model. Unlike classical confidence intervals, the survival bands produced by this method offer predictive rather than population-level inference, making them useful for personalized risk screening. For example, in a low-risk screening scenario, they can be applied to flag patients whose survival band at 12 months lies entirely above 50%, while ensuring that at least half of flagged individuals will survive past that time on average. Our approach builds on recent advances in conformal inference and integrates ideas from inverse probability of censoring weighting and multiple testing with false discovery rate control. We provide asymptotic guarantees and show promising performance in finite samples with both simulated and real data.

Keywords: Censored Data, Conformal Inference, False Discovery Rate, Predictive Calibration, Survival Analysis, Uncertainty Estimation.

1. Introduction

1.1. Background and Motivation

Survival analysis focuses on data involving time-to-event outcomes, such as the time until death or disease relapse in medicine, or mechanical failure in engineering. Its defining challenge is *censoring*, which occurs when the event of interest is not observed for all individuals. In the common case of right-censoring, we only know that the event has not occurred up to a certain time, beyond which the individual is no longer followed. For example, a cancer patient may still be alive at their last clinical follow-up, five years after diagnosis, but their true time of death is unknown because no data are available after that point. In this case, the patient's censoring time is five years and their survival time is unobserved.

In many applications, the goal is to use fitted survival models to generate personalized inferences in the form of *individual survival curves*. These curves estimate, for an individual with specific features, the probability of remaining event-free beyond any future time point. For example, based on a patient's medical history, a model might predict a 90% chance of surviving past one year, 75% past three years, and 40% past five years. These predictions can be intuitively visualized as a decreasing curve over time. Such personalized survival curves are widely used to guide decisions, such as identifying high-risk patients for early intervention or recognizing low-risk individuals who may safely avoid aggressive treatment.

Traditional approaches to survival analysis rely on statistical models that support uncertainty quantification through confidence intervals and hypothesis testing. These include

parametric models, semi-parametric models, and nonparametric estimators such as the Kaplan–Meier (KM) curve (Kaplan and Meier, 1958). Parametric models, such as the exponential or Weibull distributions, assume a specific form for how the event probability evolves over time. Semi-parametric models, most notably the Cox proportional hazards model (Cox, 1972), relax this assumption by leaving the baseline distribution unspecified, while still imposing a fixed relationship between covariates and risk. The Kaplan–Meier estimator, by contrast, avoids strong modeling assumptions altogether, but estimates only population-level survival curves and does not incorporate covariate information.

Although classical survival models have seen many successful applications, they can be limiting in modern settings that involve large datasets with rich covariate information. When the relationship between patient features and survival outcomes is complex or poorly understood, strong modeling assumptions—such as proportional hazards or specific parametric distributions—may be difficult to justify. In these situations, it is often more practical to treat classical models as black-box predictors: tools that generate individual survival curves, but whose internal assumptions are not relied on for inference. This perspective also motivates the growing use of more flexible, data-driven approaches—particularly machine learning methods such as random survival forests (Ishwaran et al., 2008), deep neural networks (Katzman et al., 2018), and gradient boosting (Barnwal et al., 2022)—which can model complex relationships with high-dimensional covariates (Spooner et al., 2020).

While black-box models can produce accurate and personalized survival predictions, they typically lack principled methods for uncertainty quantification. Classical statistical models offer confidence intervals and hypothesis tests, but—as discussed above—their guarantees rely on strong distributional assumptions that may not hold in practice. Machine learning models, by contrast, often omit uncertainty estimates altogether; and when such measures are available, they are typically heuristic and lack formal statistical justification. This limitation is particularly concerning in high-stakes applications such as clinical triage or treatment planning, where decisions must rely on calibrated, trustworthy risk estimates. There is therefore a need for widely applicable statistical tools that can provide rigorous, distribution-free uncertainty quantification for black-box survival models.

1.2. Preview of Our Contributions

We develop a method to construct principled and interpretable “uncertainty bands” around individual survival curves estimated by black-box models, using right-censored data. These bands, which we call *conformal survival bands*, are designed to support clinical screening tasks within a *multiple testing* framework and are rigorously calibrated in a useful predictive sense, as explained below.

Here, “multiple testing” refers to the fact that our method constructs uncertainty bands for each of several individuals in a test set by processing the entire test set jointly. This is analogous to how the Benjamini–Hochberg procedure processes all p -values for multiple hypotheses at once to control the false discovery rate (FDR) (Benjamini and Hochberg, 1995). The multiple testing nature of our approach will become clearer once we present its details and the theoretical guarantees it provides.

A preview of our method is provided in Figure 1, which displays conformal survival bands for four test patients in a simulated dataset. At first glance, these bands may resemble

personalized versions of classical confidence intervals—such as those drawn around Kaplan–Meier survival curves—but their interpretation is different. Classical confidence intervals aim to estimate population-level quantities like the true survival probability at a given time, which is typically only possible under strong parametric assumptions or after binning the covariates into a small number of population subgroups. Our bands, by contrast, are fully nonparametric and support individualized inferences without requiring covariate binning.

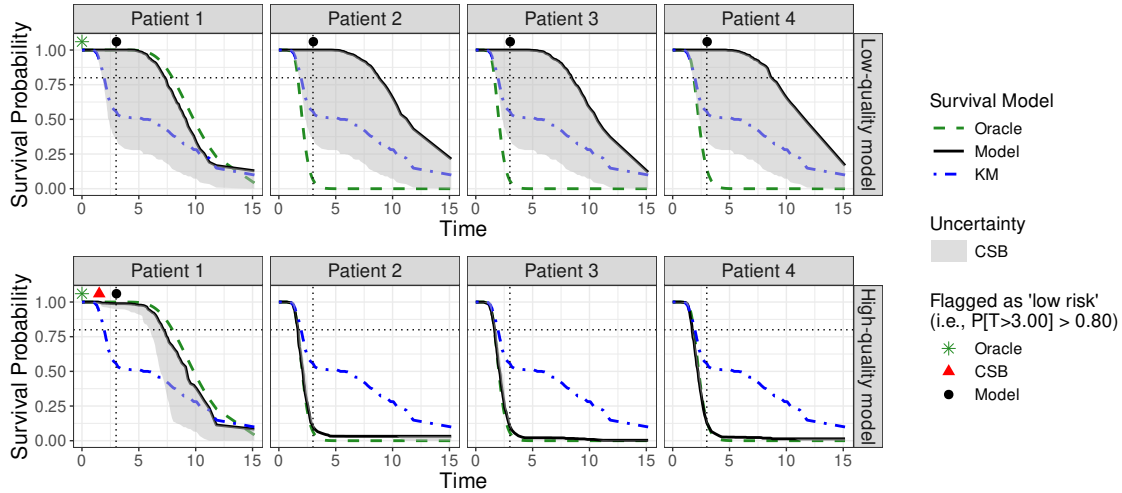


Figure 1: Illustration of the use of conformal survival bands (shaded regions) for screening test patients in simulated censored data. Solid black curves show survival estimates from either an inaccurate (top) or accurate (bottom) survival forest model. Dashed green curves represent the true survival probabilities. The goal is to identify low-risk patients—those with more than 80% probability of surviving beyond 3 years (vertical line). A patient is flagged by our method if their entire conformal band lies above the 80% threshold (horizontal line) at this time, ensuring at least 80% of flagged individuals survive longer. Flagging decisions are indicated by colored markers: red triangles (our method), green asterisks (oracle), and black circles (estimated model). Kaplan-Meier makes no selections here. Patients 2–4 are mistakenly flagged as low risk by the inaccurate model. With the accurate model, our method can identify truly low-risk patients (e.g., Patient 1), unlike the KM curve.

Rather than attempting to directly estimate population-level parameters, our bands are calibrated for distribution-free predictive screening. They allow practitioners to identify individuals whose estimated survival probability lies below or above a clinically meaningful threshold, while guaranteeing that among all such flagged individuals, the false discovery rate is approximately controlled in a precise predictive sense. This is why our method is intrinsically a multiple testing approach that processes the entire test set jointly.

For example, in Figure 1, we illustrate a scenario where one wishes to identify “low-risk” patients—defined as those with more than an $p = 80\%$ chance of surviving beyond time $t = 3$ years. A patient is flagged if their conformal band lies entirely above the 80% line at $t = 3$. Our method *asymptotically* guarantees that, among all flagged individuals, on average at least 80% will survive beyond this time. This guarantee holds *pointwise* over all

fixed choices of the probability threshold p and time horizon t , and for “low-risk” as well as “high-risk” screening, making the method flexible and broadly applicable.

This predictive notion of uncertainty within a multiple testing framework is well-aligned with how survival curves are commonly used in real-world settings: to guide actionable decisions such as prioritizing patients for further testing or early intervention. However, black-box survival models typically provide only point estimates, without any reliable measure of uncertainty or formal calibration guarantees. As in the example shown, screening based on point predictions alone, or even KM estimates, may lead to too many patients incorrectly labeled as low risk—highlighting the practical benefits of our calibrated bands, which offer rigorous control over such errors.

Our approach integrates three key concepts: conformal inference (Vovk et al., 2005; Fontana et al., 2023), inverse probability of censoring weighting (IPCW) (Robins and Rotnitzky, 1992; Robins, 1993), and FDR control (Benjamini and Hochberg, 1995). Concretely, we build on tests for *random* null hypotheses stating that $T \geq t$ or $T \leq t$, where T is the (unobserved) survival time of the test individual and $t > 0$ is a user-specified time threshold. We will compute *conformal p-values* to test these hypotheses and show how to use them to construct the conformal survival bands previewed in Figure 1. Although each of the statistical concepts upon which we build already exists on its own, the way we integrate them is original and supported by novel theoretical results.

1.3. Related Work

Conformal inference for survival analysis is a very active area of research, with several recent works addressing different inferential goals. Some methods focus on prediction of survival times (Boström et al., 2019, 2023; Sun and Wang, 2024) or classification into survival/event categories (Boström et al., 2017), while others aim to refine the point estimates of survival probabilities produced black-box models (Qi et al., 2024). The works most closely related to our own concern the construction of predictive lower bounds for survival times at a fixed confidence level $\alpha \in (0, 1)$, using censored data. These methods seek to deliver distribution-free prediction intervals while correcting for the non-exchangeability (Barber et al., 2023) introduced by censoring.

This direction was initiated by Candès et al. (2023), whose method provides valid lower bounds but is often overly conservative in practice. Gui et al. (2024) introduced refinements that improve coverage tightness, yet even their method tends to produce wide intervals at moderate coverage levels (e.g., $\alpha = 0.5$), limiting its utility in settings that require a finer characterization of predictive uncertainty. In addition, both of these methods assume type-I censoring—where censoring times are fully observed for all individuals—an assumption that is quite limiting in practice. More recent contributions by Davidov et al. (2024), Sesia and Svetnik (2025), and Farina et al. (2025) relax this constraint to accommodate general right-censoring. However, these approaches remain restricted to one-sided predictive inference and do not support threshold-based screening or two-sided uncertainty quantification.

Among the few existing methods that attempt to construct two-sided predictive intervals under censoring are those of Qi et al. (2024), Qin et al. (2024), and Yi et al. (2025). The approach of Qi et al. (2024) approaches the problem by ‘de-censoring’ the data: it estimates latent survival times in the calibration data set using the fitted survival model and

then applies standard conformal techniques as if those predictions were observed. While convenient, this method assumes that the survival model used for de-censoring is reliable. As argued by [Sesia and Svetnik \(2025\)](#), this assumption undermines the primary appeal of conformal inference—namely, its robustness to model misspecification. Indeed, empirical evidence suggests that the method of [Qi et al. \(2024\)](#) may fail to provide valid coverage when the survival model is inaccurate. By contrast, [Qin et al. \(2024\)](#) take a more model-agnostic approach, based on resampling under censoring, but their work focuses on predictive intervals for survival times rather than on calibrating survival screening decisions. Our work is also related to recent, independent research by [Yi et al. \(2025\)](#), who construct two-sided predictive intervals for survival times using a method that, like ours, utilizes conformal p-values with IPC weighting. However, their focus is on survival time prediction rather than decision-based screening or the calibration of individual survival curves.

In this paper, we address the problem of estimating uncertainty around individual survival curves to support threshold-based screening tasks. Our approach is based on computing weighted conformal p-values for testing hypotheses of the form $H_{lt}(t; j) : T_{n+j} \geq t$ and $H_{rt}(t; j) : T_{n+j} \leq t$, where T_{n+j} denotes the (unobserved) survival time of a new individual and $t > 0$ is a fixed time threshold. While it may be possible in principle to obtain such p-values by inverting predictive intervals for survival times obtained using existing methods, this is practically infeasible for several reasons: (i) most methods only provide lower bounds, (ii) even the best available bounds remain conservative, especially at moderate levels of confidence, and (iii) the inversion procedure would be computationally expensive, as it requires solving nested optimization problems for each time threshold t .

Our method integrates and extends two previously distinct lines of work. First, we build on the screening-based conformal inference framework of [Jin and Candès \(2023b\)](#), which develops p-values for testing whether unobserved outcomes exceed user-specified thresholds. [Jin and Candès \(2023a\)](#) later generalized this approach to allow for importance weighting under covariate shift ([Tibshirani et al., 2019](#)), but their theory and assumptions are not suited to censoring. Second, we draw on [Farina et al. \(2025\)](#), who introduced IPCW ([Robins and Rotnitzky, 1992](#); [Robins, 1993](#)) into conformal inference to produce predictive intervals under right-censoring. While their method is tailored to one-sided prediction bounds, we adapt IPCW techniques to develop conformal p-values for our survival threshold hypotheses.

2. Methods

2.1. Problem Setup and Assumptions

We consider a right-censored survival setting based on a sample of n individuals indexed by $[n] := \{1, \dots, n\}$, drawn i.i.d. from an unknown population. For each individual $i \in [n]$, we observe a vector of covariates $X_i \in \mathcal{X} \subseteq \mathbb{R}^d$, along with right-censored survival data: the event indicator $E_i := \mathbb{I}(T_i < C_i)$ and the observed time $\tilde{T}_i := \min(T_i, C_i)$, where $T_i > 0$ is the true survival time and $C_i > 0$ is the censoring time. These n observations form the calibration dataset, denoted by $\mathcal{D}_{\text{cal}} := \{(X_i, \tilde{T}_i, E_i)\}_{i=1}^n$, which will be used to quantify uncertainty—i.e., to calibrate the predictions of a black-box survival model.

We assume access to two black-box models trained on an independent dataset, which may itself be censored and need not follow the same distribution as the calibration or test sets. The only assumption is that the training data are independent of all other

samples, allowing us to treat both models as fixed throughout the analysis. The first model is the *survival model*, $\hat{\mathcal{M}}_T$. This model produces an estimated individual survival function $\hat{S}_T(t | x)$, which is intended to approximate the true conditional survival probability $S_T(t | x) := \mathbb{P}(T > t | X = x)$. The second is an auxiliary *censoring model*, $\hat{\mathcal{M}}_C$, which estimates the conditional survival function of the censoring distribution, $\hat{S}_C(t | x)$, an approximation to $S_C(t | x) := \mathbb{P}(C > t | X = x)$. The role of the censoring model is to reweight the calibration data to correct for the missing information due to censoring, enabling valid uncertainty quantification for the survival model.

In addition to the calibration set, we consider a disjoint test set consisting of m individuals, indexed by $\{n + 1, \dots, n + m\}$, also drawn independently from the same population. For each test individual j , we observe only the covariates $X_j \in \mathcal{X}$. Our goal is to use the survival model $\hat{\mathcal{M}}_T$ to estimate the survival curve $\hat{S}_T(t | X_j)$ for each X_j , and to construct a well-calibrated *conformal survival band* around this curve that reflects uncertainty.

Rather than aiming for our conformal survival bands to provide valid confidence intervals for the true survival function $S_T(t | x)$, which would not be a feasible goal without additional assumptions (Barber, 2020), we focus on a practically useful objective that is more naturally attainable within a conformal inference framework: producing principled and interpretable uncertainty estimates around $\hat{S}_T(t | X_j)$ that can support confident predictive screening decisions. For example, given a survival probability threshold $q \in (0, 1)$ and a clinically meaningful time point $t > 0$, our goal may be to identify *high-risk* test individuals whose predicted survival probability $\hat{S}_T(t | X_j)$ falls *significantly* below q , while guaranteeing that, among all such flagged individuals, the expected proportion who survive beyond t remains controlled below level q . Although inherently predictive in nature, this type of calibration guarantee is intuitive and aligns closely with how physicians and practitioners often interpret the output of survival models in practice: as actionable, patient-specific risk estimates that can support confident decision-making.

Because our method is built on the idea of *conformal p-values*, we first review the key concepts underlying this approach.

2.2. Preliminaries: Review of Conformal Inference without Censoring

To build intuition for our method, we begin by reviewing how existing conformal inference techniques for regression can be applied to survival analysis in a simplified setting without censoring. In this case, imagine observing complete data $\{(X_i, T_i)\}_{i=1}^n$, where each (X_i, T_i) is drawn i.i.d. from an unknown distribution over $\mathcal{X} \times \mathbb{R}_+$. Given a new test point with covariates X_{n+j} , consider the task of testing the null hypothesis

$$H_{\text{lt}}(t; j) : T_{n+j} \geq t, \tag{1}$$

for a fixed time $t > 0$. This hypothesis asserts the individual will experience the event after time t ; rejecting it provides evidence that the individual is unlikely to survive beyond t .

If complete (uncensored) data are available, this hypothesis can be tested as follows. Let $\hat{s}_{\text{lt}}(t'; x, t)$ be any *monotone increasing* function of t' , which may depend on $x \in \mathcal{X}$ and the fixed time $t \in \mathbb{R}_+$ in the definition of the null hypothesis (1). We refer to \hat{s}_{lt} as the *left-tail* nonconformity scoring function.

An intuitive choice of scoring function satisfying the required monotonicity property is

$$\hat{s}_{\text{lt}}(t'; x, t) := \left(1 - \hat{S}_T(t' \mid x)\right) \mathbb{I}[t' \geq t], \quad (2)$$

where $\hat{S}_T(t' \mid x) \in (0, 1)$ denotes the conditional survival probability at time t' for an individual with covariates x , estimated by the model $\hat{\mathcal{M}}_T$. We have found that this scoring function works well in practice, although it is not the only possible option.

Using any suitable monotone increasing scoring function \hat{s}_{lt} , we compute a non-conformity score $\hat{s}_{\text{lt}}(T_i, X_i, t)$ for each calibration point $i \in [n]$ as well as for the test point X_{n+j} , in the latter case using the placeholder value t instead of the unknown event time T_{n+j} —i.e., defining the test score as $\hat{s}_{\text{lt}}(t, X_{n+j}, t)$. Then, following the approach of [Jin and Candès \(2023b\)](#), we define the left-tail conformal p-value as:

$$\tilde{\phi}_{\text{lt}}(t; X_{n+j}) := \frac{1 + \sum_{i=1}^n \mathbb{I}\{\hat{s}_{\text{lt}}(T_i; X_i, t) \geq \hat{s}_{\text{lt}}(t; X_{n+j}, t)\}}{n + 1}. \quad (3)$$

For example, in the special case where the scoring function takes the form defined in (2), under the relatively mild assumption that $\hat{S}_T(t \mid X_{n+j}) < 1$ almost surely, then it is easy to verify that the left-tail conformal p-value in (3) can be written as:

$$\tilde{\phi}_{\text{lt}}(t; X_{n+j}) := \frac{1 + \sum_{i=1}^n \mathbb{I}[T_i \geq t] \mathbb{I}\left\{\hat{S}_T(T_i \mid X_i) \leq \hat{S}_T(t \mid X_{n+j})\right\}}{n + 1}.$$

The following result, due to [Jin and Candès \(2023b\)](#), highlights the statistical validity of the left-tail p-value defined in (3) for testing the null hypothesis (1). A formal proof is included in [Appendix A](#) for completeness. [Appendix A](#) also contains all other proofs.

Proposition 1 ([Jin and Candès \(2023b\)](#)) *Assume the calibration data and test point (X_{n+j}, T_{n+j}) are exchangeable and, for a fixed $t > 0$, the scoring function $\hat{s}_{\text{lt}}(t'; x, t)$ is monotone increasing in t' for all x . Then, for any $\alpha \in (0, 1)$, the conformal p-value $\tilde{\phi}_{\text{lt}}(t; X_{n+j})$ defined in (3), computed using the full (uncensored) data, satisfies:*

$$\mathbb{P}\left[\tilde{\phi}_{\text{lt}}(t; X_{n+j}) \leq \alpha, T_{n+j} \geq t\right] \leq \alpha.$$

This implies $\tilde{\phi}_{\text{lt}}(t; X_{n+j})$ can be interpreted as a p-value for testing the *random* hypothesis (1), in the sense that small values provide evidence against the null. Notably, the guarantee in [Proposition 1](#) does not condition on the null event $\{T_{n+j} \geq t\}$; instead, it controls the *joint* probability that both the null is true and the p-value is small. This marginal formulation differs from the classical frequentist setup, where hypotheses are non-random, but it is sufficient for our purposes. As we will see later, this form of validity is precisely what enables us to obtain valid calibration guarantees for personalized risk screening.

Before turning to censored data, it is helpful to note that the above ideas also apply for testing the complementary *right-tail* null hypothesis

$$H_{\text{rt}}(t; j) : T_{n+j} \leq t, \quad (4)$$

which is useful for identifying low-risk individuals. In this case, we consider a right-tail scoring function assumed to be monotone *decreasing* in t' ; e.g.,

$$\hat{s}_{\text{rt}}(t'; x, t) := \hat{S}_T(t' | x) \mathbb{I}[t' \leq t], \quad (5)$$

interpreted as the predicted probability of experiencing the event after time t' for an individual with covariates x . The corresponding right-tail conformal p-value is

$$\tilde{\phi}_{\text{rt}}(t; X_{n+j}) := \frac{1 + \sum_{i=1}^n \mathbb{I}\{\hat{s}_{\text{rt}}(T_i; X_i, t) \geq \hat{s}_{\text{rt}}(t; X_{n+j}, t)\}}{n + 1},$$

which is super-uniform under the same assumptions as in Proposition 1:

$$\mathbb{P}\left[\tilde{\phi}_{\text{rt}}(t; X_{n+j}) \leq \alpha, T_{n+j} \leq t\right] \leq \alpha.$$

This completes the review of standard conformal p-values in the uncensored setting, for both left- and right-tail survival hypotheses. In the next section, we extend this framework to account for censoring, enabling application to real-world survival data.

2.3. Conformal Inference for Survival Analysis under Right-Censoring

We now return to the right-censored survival setting described in Section 2.1, where the calibration data consist of i.i.d. samples $\{(X_i, \tilde{T}_i, E_i)\}_{i=1}^n$. As in Section 2.2, our goal is to test the left-tail and right-tail hypotheses defined in (1) and (4) for new test points X_{n+j} .

The main challenge compared to the idealized setting of Section 2.2 is that we do not observe the true event times T_i for all calibration individuals. Consequently, we cannot compute the full-data conformal p-values $\hat{\phi}_{\text{lt}}(t; X_{n+j})$ or $\hat{\phi}_{\text{rt}}(t; X_{n+j})$ as defined earlier.

To address this, we adapt the conformal p-value construction using IPCW (Robins and Rotnitzky, 1992; Robins, 1993). This approach uses censored calibration samples but carefully reweights them to correct for censoring-induced selection bias. Specifically, we rely on a fitted censoring model \hat{M}_C that estimates the conditional censoring survival function $\hat{S}_C(t | x) \approx \mathbb{P}(C > t | X = x)$. Based on this model, we define the weight function

$$\hat{w}(t, x) := \frac{1}{\hat{S}_C(t | x)}, \quad (6)$$

approximating the inverse probability of being uncensored at time t given covariates x .

Bringing this into the conformal framework, we define the IPCW version of the left-tail conformal p-value as:

$$\hat{\phi}_{\text{lt}}(t; X_{n+j}) := \frac{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i) \cdot \mathbb{I}\{\hat{s}_{\text{lt}}(T_i; X_i, t) \geq \hat{s}_{\text{lt}}(t; X_{n+j}, t)\}}{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i)}, \quad (7)$$

where the numerator includes only data points whose true event times are observed. For example, in the special case where the scoring function takes the form defined in (2), under the relatively mild assumption that $\hat{S}_T(t | X_{n+j}) < 1$ almost surely, the left-tail conformal p-value in (7) can be written as:

$$\hat{\phi}_{\text{lt}}(t; X_{n+j}) := \frac{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i) \cdot \mathbb{I}[T_i \geq t] \cdot \mathbb{I}\{\hat{S}_T(T_i | X_i) \leq \hat{S}_T(t | X_{n+j})\}}{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i)}.$$

Analogously, we define the IPCW right-tail p-value as:

$$\hat{\phi}_{\text{rt}}(t; X_{n+j}) := \frac{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i) \cdot \mathbb{I}\{\hat{s}_{\text{rt}}(T_i; X_i, t) \geq \hat{s}_{\text{rt}}(t; X_{n+j}, t)\}}{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i)}. \quad (8)$$

In the special case where the scoring function takes the form defined in (5), under the relatively mild assumption that $\hat{S}_T(t | X_{n+j}) > 0$ almost surely, this can be written as:

$$\hat{\phi}_{\text{rt}}(t; X_{n+j}) := \frac{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i) \cdot \mathbb{I}[T_i \leq t] \cdot \mathbb{I}\left\{\hat{S}_T(T_i | X_i) \geq \hat{S}_T(t | X_{n+j})\right\}}{1 + \sum_{i=1}^n E_i \cdot \hat{w}(T_i, X_i)}.$$

These IPCW p-values are fully computable from the observed censored data. As we now show, they yield asymptotically valid inference under the assumption that the censoring model $\hat{\mathcal{M}}_C$ is consistent, along with mild regularity conditions on the data distribution.

We define the censoring weight estimation error as:

$$\Delta_N := \left(\mathbb{E} \left[\left(\frac{1}{\hat{w}(T; X)} - \frac{1}{w^*(T; X)} \right)^2 \right] \right)^{1/2},$$

where $w^*(t, x) := 1/S_C(t | x)$ denotes the true (unknown) censoring weight function, and N represents the number of training data points.

We now list the assumptions under which asymptotic validity holds:

(A1) The data are split into three independent parts:

- a training set of cardinality N used to estimate $\hat{w}(t; x)$, $\hat{s}_{\text{lt}}(t'; x, t)$, and $\hat{s}_{\text{rt}}(t'; x, t)$;
- an i.i.d. calibration set (X_i, T_i, C_i) for $i = 1, \dots, n$, which is censored;
- and an i.i.d. test point $(X_{n+j}, T_{n+j}, C_{n+j})$, of which we only see the covariates.

(A2) The scoring function $\hat{s}_{\text{lt}}(t'; x, t)$ is almost surely monotone increasing in t' for all x, t , while $\hat{s}_{\text{rt}}(t'; x, t)$ is monotone decreasing. Moreover, $\hat{s}_{\text{lt}}(t; x, t) \in (0, 1)$ and $\hat{s}_{\text{rt}}(t; x, t) \in (0, 1)$ for all x, t , almost surely.

(A3) The probability of observing an event is bounded away from zero: $\pi := \mathbb{P}(T \leq C) > 0$.

(A4) The estimated weights are bounded below: $\hat{w}(T; X) \geq \omega_{\min} > 0$ almost surely.

(A5) The weight estimation error vanishes asymptotically: $\Delta_N \rightarrow 0$ as $N \rightarrow \infty$.

Theorem 2 *Under Assumptions (A1)–(A5), for any fixed $t > 0$ and $\alpha \in (0, 1)$, the IPCW conformal p-values defined in (7) satisfy:*

$$\begin{aligned} \limsup_{N, n \rightarrow \infty} \mathbb{P} \left[\hat{\phi}_{\text{lt}}(t; X_{n+1}) \leq \alpha, T_{n+1} \geq t \right] &\leq \alpha, \\ \limsup_{N, n \rightarrow \infty} \mathbb{P} \left[\hat{\phi}_{\text{rt}}(t; X_{n+1}) \leq \alpha, T_{n+1} \leq t \right] &\leq \alpha. \end{aligned}$$

In Section 2.4, we will explain how to use IPCW conformal p-values to construct *conformal survival bands*.

2.4. Conformal Survival Bands for Screening High- or Low-Risk Individuals

Consider m test patients with feature vectors $X_{n+1}, X_{n+2}, \dots, X_{n+m}$. For each $j \in [m] := \{1, \dots, m\}$ and a fixed time $t > 0$, define

$$\hat{U}(t; X_{n+j}) := \hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+j}), \quad (9)$$

where $\hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+j})$ denotes the Benjamini-Hochberg (BH) adjusted left-tail p-value.

More precisely, let $\hat{\phi}_{(1)} \leq \hat{\phi}_{(2)} \leq \dots \leq \hat{\phi}_{(m)}$ denote the ordered values of the unadjusted left-tail p-values $\{\hat{\phi}_{\text{lt}}(t; X_{n+j})\}_{j=1}^m$, and let $\pi(j)$ be the index such that $\hat{\phi}_{(j)} = \hat{\phi}_{\text{lt}}(t; X_{n+\pi(j)})$. Then the BH-adjusted p-values are defined as

$$\hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+\pi(j)}) := \min_{k \geq j} \left\{ \frac{m}{k} \cdot \hat{\phi}_{(k)} \right\} \wedge 1, \quad \text{for } j = 1, \dots, m. \quad (10)$$

Each value $\hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+j})$ represents the smallest FDR level at which the null hypothesis $H_{\text{lt}}(t; j)$ can be rejected by the BH procedure (Benjamini and Hochberg, 1995). In practice, these adjusted p-values can be very easily computed using the `p.adjust` function in R with the argument `method = "BH"`.

The quantity $\hat{U}(t; X_{n+j})$ can then be interpreted as a calibrated upper bound on the survival probability at time t for individual X_{n+j} . If $\hat{U}(t; X_{n+j}) \leq \alpha$ for some threshold $\alpha \in (0, 1)$, then the individual may be confidently flagged as *high-risk*, with the guarantee that the expected proportion of false discoveries—individuals who survive past t despite being flagged as “high-risk”—remains below α in the large-sample limit.

Symmetrically, we define

$$\hat{L}(t; X_{n+j}) := 1 - \hat{\phi}_{\text{rt}}^{\text{BH}}(t; X_{n+j}), \quad (11)$$

where $\hat{\phi}_{\text{rt}}^{\text{BH}}(t; X_{n+j})$ is the BH-adjusted p-value computed from the collection of right-tail IPCW conformal p-values $\{\hat{\phi}_{\text{rt}}(t; X_{n+\ell})\}_{\ell=1}^m$. If $\hat{L}(t; X_{n+j}) \geq 1 - \alpha$, the patient can be confidently flagged as *low-risk*, with the guarantee that the expected proportion of individuals who fail before t —despite being classified as low-risk—approximately remains below α .

This construction defines our calibrated individual survival band

$$[\hat{L}(t; X_{n+j}), \hat{U}(t; X_{n+j})] \quad (12)$$

which could be interpreted as a calibrated range for likely survival probabilities at time t —offering statistically principled decision support for confident high- and low-risk screening. The full procedure described above is summarized in Algorithm 1.

This algorithm is straightforward to implement and computationally efficient. Model training is performed once on the training dataset, and nonconformity scores for the uncensored calibration points are computed a single time, independently of the test set and the evaluation time grid. The only components that depend on the chosen time points are the nonconformity scores for the test samples, the IPCW conformal p-values, and their BH-adjusted counterparts. While the BH adjustment requires sorting across test samples, it is very fast in practice. All other steps can be parallelized over both time points and test samples. As a result, the overall computational cost of conformal calibration is modest, and typically negligible compared to the cost of training the survival or censoring models.

Algorithm 1 Construction of Conformal Survival Bands

input Censored training dataset $\mathcal{D}_{\text{train}}$; censored calibration dataset $\mathcal{D}_{\text{cal}} = \{(X_i, \tilde{T}_i, E_i)\}_{i=1}^n$; test covariates X_{n+1}, \dots, X_{n+m} ; time grid $\mathcal{T} = \{t_1, \dots, t_K\}$; trainable survival model $\hat{\mathcal{M}}_T$ and censoring model $\hat{\mathcal{M}}_C$.

- 1: Define the active calibration subset: $\mathcal{D}'_{\text{cal}} := \{(X_i, T_i) : E_i = 1, i \in [n]\}$.
- 2: Train $\hat{\mathcal{M}}_T$ and $\hat{\mathcal{M}}_C$ on $\mathcal{D}_{\text{train}}$.
- 3: Estimate censoring weights $\hat{w}(t, x)$ for the data in $\mathcal{D}'_{\text{cal}}$ using $\hat{\mathcal{M}}_C$ as in Eq. (6).
- 4: **for** each time $t \in \mathcal{T}$ **do**
- 5: **for** each calibration point $(X_i, T_i) \in \mathcal{D}'_{\text{cal}}$ **do**
- 6: Compute $\hat{s}_{\text{lt}}(T_i; X_i, t)$ and $\hat{s}_{\text{rt}}(T_i; X_i, t)$ using Eqs. (2) and (5).
- 7: **end for**
- 8: **for** each test point X_{n+j} , for $j = 1, \dots, m$ **do**
- 9: Compute scores $\hat{s}_{\text{lt}}(t; X_{n+j}, t)$ and $\hat{s}_{\text{rt}}(t; X_{n+j}, t)$ using Eqs. (2) and (5).
- 10: Compute IPCW p-values $\hat{\phi}_{\text{lt}}(t; X_{n+j})$ and $\hat{\phi}_{\text{rt}}(t; X_{n+j})$ using Eqs. (7) and (8).
- 11: **end for**
- 12: Apply BH procedure to p-values across test set to obtain adjusted values $\hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+j})$ and $\hat{\phi}_{\text{rt}}^{\text{BH}}(t; X_{n+j})$ for all $j \in [m]$ as in Eq. (10).
- 13: **for** each test point $j \in [m]$ **do**
- 14: Compute endpoints $\hat{U}(t; X_{n+j})$ and $\hat{L}(t; X_{n+j})$ using Eqs. (9) and (11).
- 15: **end for**
- 16: **end for**

output Personalized survival bands $[\hat{L}(t; X_{n+j}), \hat{U}(t; X_{n+j})]$ for each $j \in [m]$ and all $t \in \mathcal{T}$.

To implement Algorithm 1 in an even more data-efficient way, one may choose to re-allocate all censored observations from the calibration set \mathcal{D}_{cal} into the training set. This yields a larger training dataset $\mathcal{D}'_{\text{train}} := \mathcal{D}_{\text{train}} \cup \{(X_i, \tilde{T}_i, E_i) : E_i = 0, i \in [n]\}$. This does not affect the statistical validity of our method, since the conformal p -values $\hat{\phi}_{\text{lt}}(t; X_{n+j})$ and $\hat{\phi}_{\text{rt}}(t; X_{n+j})$ are computed exclusively using uncensored calibration points. While this approach deviates from the standard sample-splitting paradigm commonly used in conformal inference, it remains valid within our framework due to the use of inverse probability of censoring weighting (IPCW), which automatically corrects for the potential selection bias introduced by reallocation during the calibration phase. However, it is not guaranteed that re-allocating the censored observations to the training set will improve model estimation, as doing so introduces its own form of sampling bias in the training data. Exploring the trade-offs involved in this reallocation strategy—particularly in combination with the possible use of importance weighting during training—remains an open direction for future work.

2.5. Theoretical Guarantees for Calibrated Survival Bands

We will now formally state the asymptotic statistical guarantee provided by our method, which requires the following additional assumption:

(A5) For all $t' > 0$, the score functions $\hat{s}_{\text{lt}}(t'; X, t)$ admit a Lebesgue density $f_{t'}$ satisfying

$$0 < f_{\min} \leq f_{t'}(u) \leq f_{\max} < \infty \quad \text{for all } u > 0,$$

for some constants $f_{\min}, f_{\max} \in (0, \infty)$.

Theorem 3 Under Assumptions (A1)–(A5), fix a target level $\alpha \in (0, 1)$ and survival cutoff time $t > 0$. Let $p_j := \hat{\phi}_{\text{lt}}(t; X_{n+j})$ denote the IPCW p -value used to test the null hypothesis $H_{\text{lt}}(t; j) : T_{n+j} \geq t$, for $j = 1, \dots, m$. Define the false discovery rate (FDR) as

$$\text{FDR}_{m,n} := \mathbb{E} \left[\frac{|\mathcal{R} \cap \mathcal{H}_0|}{|\mathcal{R}| \vee 1} \right],$$

where $\mathcal{R} := \{j : p_j \leq \hat{\tau}\}$ is the rejection set obtained by applying the BH procedure at level α to (p_1, \dots, p_m) and $\mathcal{H}_0 := \{j : T_{n+j} \geq t\}$ is the (random) set of true null hypotheses. Then,

$$\limsup_{N, n \rightarrow \infty} \text{FDR}_{m,n} \leq \alpha \quad \text{for any fixed } m.$$

If in addition $m = m_n \rightarrow \infty$ and $m_n^2 \epsilon_n \rightarrow 0$, then

$$\limsup_{N, m, n \rightarrow \infty} \text{FDR}_{m,n} \leq \alpha.$$

An analogous result holds for the right-tail p -values $\hat{\phi}_{\text{rt}}(t; X_{n+j})$, used to test $H_{\text{rt}}(t; j)$.

A corollary of Theorem 3 is that high- and low-risk screening rules based on the survival bands produced by Algorithm 1 described above are asymptotically well-calibrated. For any $t > 0$ and $\alpha \in (0, 1)$, define the set of test individuals flagged as high-risk at time t as

$$\mathcal{F}_{\text{hi}}(t; \alpha) := \left\{ j \in [m] : \hat{U}(t; X_{n+j}) \leq \alpha \right\} = \left\{ j \in [m] : \hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+j}) \leq \alpha \right\},$$

and similarly, define the set of individuals flagged as low-risk as

$$\mathcal{F}_{\text{lo}}(t; \alpha) := \left\{ j \in [m] : \hat{L}(t; X_{n+j}) \geq 1 - \alpha \right\} = \left\{ j \in [m] : \hat{\phi}_{\text{rt}}^{\text{BH}}(t; X_{n+j}) \leq \alpha \right\},$$

Then, Theorem 3 implies that, in the large-sample limit, the expected proportion of individuals in $\mathcal{F}_{\text{hi}}(t; \alpha)$ who actually survive past t , and the expected proportion of individuals in $\mathcal{F}_{\text{lo}}(t; \alpha)$ who fail before t , are both asymptotically controlled below level α .

2.6. Extension: Making the Survival Bands Doubly Robust

The validity guarantee provided by Theorem 3 relies on the assumption that the censoring model used to estimate the inverse probability of censoring (IPC) weights is asymptotically consistent. Under this condition, our method enables approximately well-calibrated screening rules without requiring consistency of the black-box survival model.

We now describe a simple extension that makes our method *more conservative*, aiming to improve robustness to misspecification of the censoring model. Specifically, this variant targets *double robustness*: it is designed to deliver approximately calibrated screening rules as long as *either* the censoring model *or* the survival model is consistent—though not necessarily both. We do not prove a formal double robustness guarantee, as the paper is already quite long. Nonetheless, we believe that a formal double robustness result for this extension could be established using techniques similar to those in [Sesia and Svetnik \(2025\)](#).

For each test point $j \in [m]$, we define the adjusted conformal survival band:

$$[\hat{L}^{\text{DR}}(t; X_{n+j}), \hat{U}^{\text{DR}}(t; X_{n+j})], \quad (13)$$

where

$$\hat{U}^{\text{DR}}(t; X_{n+j}) := \min \left\{ \hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+j}), \hat{S}_T(t | X_{n+j}) \right\}, \quad (14)$$

and

$$\hat{L}^{\text{DR}}(t; X_{n+j}) := \max \left\{ 1 - \hat{\phi}_{\text{rt}}^{\text{BH}}(t; X_{n+j}), \hat{S}_T(t | X_{n+j}) \right\}. \quad (15)$$

Above, $\hat{S}_T(t | X_{n+j})$ is the point estimate of the survival probability produced by $\hat{\mathcal{M}}_T$, and $\hat{\phi}_{\text{lt}}^{\text{BH}}(t; X_{n+j})$ and $\hat{\phi}_{\text{rt}}^{\text{BH}}(t; X_{n+j})$ are the upper and lower bounds produced by Algorithm 1. This adjusted construction was used in the illustrative example shown in Figure 1.

A useful side benefit of this adjusted band is enhanced interpretability: Equation (13) guarantees that the predicted survival probability $\hat{S}_T(t | X_{n+j})$ always lies within the band itself. This aligns with common practitioner expectations, simplifying communication.

3. Numerical Experiments with Synthetic Data

3.1. Setup

Synthetic data. We consider four data-generating distributions, summarized in Table 3 (Appendix B), which span a range of interesting settings, inspired by previous works. In each setting, $p = 100$ covariates $X = (X_1, \dots, X_p)$ are generated independently, while T and C are sampled independently conditional on X , from either a log-normal distribution— $\log T | X \sim \mathcal{N}(\mu(X), \sigma(X))$ —or an exponential distribution— $C | X \sim \text{Exp}(\lambda(X))$.

The first three settings are borrowed from [Sesia and Svetnik \(2025\)](#) and are ordered by decreasing modeling difficulty. The first two simulate challenging scenarios where the true survival distribution is highly nonlinear or complex, making accurate model fitting difficult and highlighting the need for robust uncertainty estimation via conformal inference. The third setting, originally appearing in [Candès et al. \(2023\)](#), corresponds to a simpler survival distribution that is easier to learn with standard survival models, reducing the gap between model-based and conformal prediction. Finally, the fourth setting, which is new, introduces a covariate shift between the training and calibration/test distributions, leading to potential model misspecification even if the survival model fits the training distribution well, and further emphasizing the value of conformal inference.

Design, Methods, and Performance Metrics. We generate independent training and calibration datasets with varying sample sizes, along with a test dataset containing 1000 samples. Right-censoring is introduced by replacing the true event time T and censoring time C with the observed time $\tilde{T} = \min(T, C)$ and event indicator $E = \mathbb{I}(T \leq C)$. The censored training data are used to fit survival and censoring models, as described below.

Our goal is to evaluate the proposed conformal survival band (CSB) method and compare it to natural baseline approaches for screening test patients according to rules of the form $P[T > t] > p$ (for *low-risk* selection) or $P[T > t] < p$ (for *high-risk* selection), evaluated at various time thresholds t and probability levels p . We consider three methods for

screening patients based on estimated conditional survival probabilities: (i) *model*, using the point estimates of conditional survival probabilities computed by a black-box survival model $\hat{\mathcal{M}}^{\text{surv}}$ fitted on all available data (training and calibration); (ii) *KM*, using a Kaplan–Meier estimator fitted on the calibration data; and (iii) *CSB*, using the conformal survival band produced by Algorithm 1, implemented with the doubly robust extension described in Section 2.6.

We focus on the doubly robust implementation of CSB in our experiments because it is more interpretable, ensuring that the predicted survival probability always lies within the band, as a practitioner would likely expect. Although formally studying the double robustness property of this extension is beyond the scope of this paper, this approach is already supported by the statistical guarantees established in Section 2.5, since it is always more conservative than the baseline method described in Section 2.4.

In addition, we compare against an ideal *oracle* that uses the true conditional survival probabilities from the data-generating distribution; although not practical, the oracle provides valuable insight into the achievable performance limits.

Performance is evaluated on the independent test set using several metrics. The *screened proportion* measures the fraction of test patients selected by the screening rule; this quantity should ideally be large, reflecting high power, provided that selections are accurate. The *survival rate* measures the proportion of selected patients who survive beyond time t , and should be larger than p for low-risk screening and smaller than p for high-risk screening. Finally, *precision* and *recall* are computed relative to the oracle screening decisions, treating screening as a binary classification task, and both should ideally be high.

All experiments are independently repeated 100 times, and results are reported as means across repetitions, with error bars representing two standard errors (SE).

Models. We consider four model families for fitting the censoring and survival models, denoted respectively by $\hat{\mathcal{M}}^{\text{cens}}$ and $\hat{\mathcal{M}}^{\text{surv}}$, to ensure consistent comparisons across different calibration methods. The model families are: (1) *grf*, a generalized random forest (R package `grf`); (2) *survreg*, an accelerated failure time (AFT) model with a log-normal distribution (R package `survival`); (3) *rf*, a random survival forest (R package `randomForestSRC`); and (4) *cox*, the Cox proportional hazards model (R package `survival`). The survival model $\hat{\mathcal{M}}^{\text{surv}}$ is used both by the *model* screening method and as input to the CSB method. In contrast, the censoring model $\hat{\mathcal{M}}^{\text{cens}}$ is only used within the CSB method, to compute the weights for constructing conformal survival bands.

Leveraging Prior Knowledge on $P_{C|X}$. To examine the effect of incorporating prior knowledge about the censoring distribution, we fit $\hat{\mathcal{M}}^{\text{cens}}$ using only the first $p_1 \leq p$ covariates, knowing that C is independent of (X_{p_1+1}, \dots, X_p) given (X_1, \dots, X_{p_1}) within the data-generating distributions used in these experiments (Table 3). Therefore, for $10 \leq p_1 \leq p = 100$, this prior knowledge helps improve the censoring model by excluding irrelevant predictors and mitigating overfitting. We start with $p_1 = 10$ and later evaluate the impact of larger p_1 , representing weaker prior knowledge. The case $p_1 = p$ corresponds to no prior knowledge, where all covariates are used to fit the censoring model.

3.2. Results

Effect of the Training Sample Size. Figure 2 compares the performance of the four methods in Setting 1 (the harder case), using the *grf* models. The training sample size varies between 100 and 10,000, while the calibration sample size is fixed at 500. We evaluate two screening rules: selecting low-risk patients with $P(T > 6.00) > 0.80$ and selecting high-risk patients with $P(T > 12.00) < 0.80$.

Low-risk screening results (top panel). The *model* selects too many patients, resulting in an excess of false positives; the average survival rate among its selected patients falls below 60%, despite the target threshold of 80%. The *KM* method performs even worse, selecting nearly all patients, with a survival rate close to 50%. In contrast, *CSB* achieves the desired survival rate among selected patients, provided that the training sample size is not too small, and its performance improves steadily as the sample size increases, approaching that of the oracle. Even when the training size is small (e.g., 100), and the fitted survival and censoring models are relatively inaccurate, *CSB* is more robust than the *model*. As expected, the oracle achieves a survival rate above 80%—in fact closer to 100%—while selecting approximately 50% of the test patients. It is important to note that the survival rate of selected patients does not need to match the threshold p exactly, even for the oracle, because many patients have true survival probabilities much higher (or lower) than p .

High-risk screening results (bottom panel). In this setting, all methods lead to subsets of selected patients whose survival rates are below 80%, as desired. However, the *KM* method again selects too many patients, including many whose true survival probabilities are actually higher than 80%, resulting in lower precision. This behavior stems from the non-personalized nature of *KM* estimates: *KM* cannot flexibly select subsets of patients based on individual characteristics and must either select nearly all or none. In contrast, the *model* and *CSB* methods achieve higher precision and recall, and both approach the oracle performance as the training sample size increases.

Figures 5 and 6 in Appendix B present similar results for Settings 2 and 3, respectively. Across these experiments, the *CSB* method consistently achieves survival rates on the correct side of the target threshold and maintains relatively high precision and recall compared to the other methods, more closely approximating the performance of the ideal oracle as the training sample size increases.

Effect of the Calibration Sample Size. Figure 3 examines the impact of the calibration sample size on the performance of *CSB* in the same challenging synthetic data setting as before, with the training sample size fixed at 5000. In this regime, where the survival and censoring models are already accurate thanks to a relatively large training set, increasing the calibration sample size improves screening performance: the survival rates among selected patients consistently falls on the correct side of the target threshold, and precision and recall tend to approach the behavior of the ideal oracle. In general, however, if the training sample size is small and the fitted models are inaccurate, it may be more beneficial to prioritize training over calibration. Overall, a moderate calibration sample size (on the order of a few hundred) is typically sufficient to obtain reliable conformal inferences. Figures 7 and 8 in Appendix B present qualitatively similar results from analogous experiments conducted under synthetic data under Settings 2 and 3.

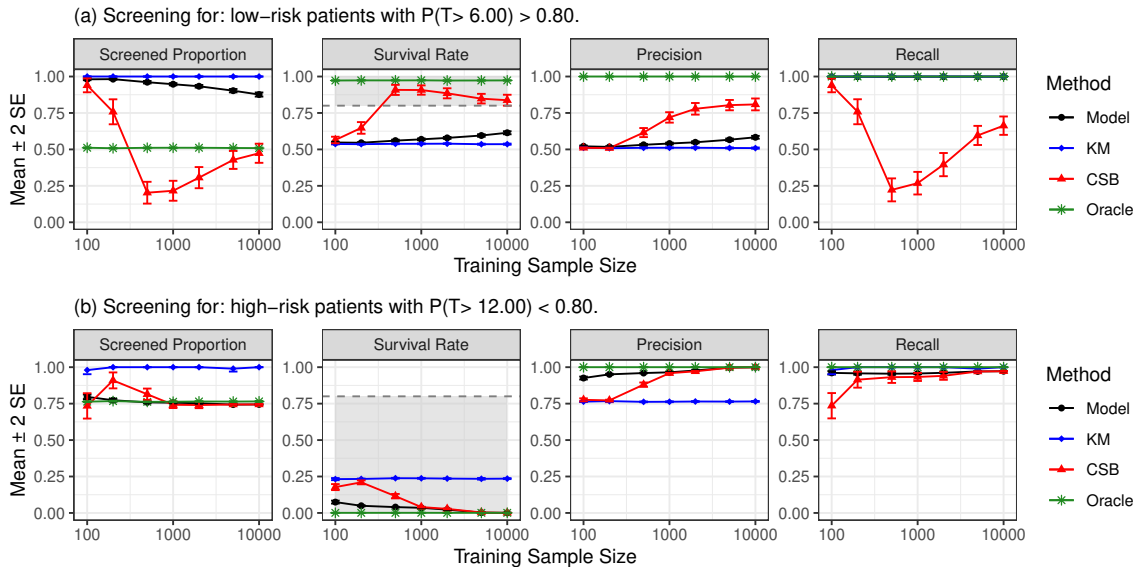


Figure 2: Effect of training sample size on patient screening performance of different methods in a challenging synthetic data scenario with complex survival and censoring distributions, using *grf* models. Top: results for low-risk screening at time $t = 6$; bottom: high-risk screening at time $t = 12$. The calibration sample size is fixed at 500. The conformal survival band (CSB) method successfully achieves survival rates above the target $p = 0.8$ for low-risk screening and below $p = 0.8$ for high-risk screening, while making personalized selections that approach the oracle performance as the training sample size increases.

Effect of the Training Sample Size for the Censoring Model. Figure 4 presents results from experiments similar to those in Figure 2, but here the censoring model is fit using only a subset from a total of 5000 training samples. The goal is to study how the quality of the censoring model affects the performance of the CSB method. In the top panel (low-risk screening), we observe that when the censoring model is trained on too few samples, CSB may fail to provide valid screening selections. As the training size increases and the censoring model improves, the survival rate among selected patients eventually falls on the correct side of the target threshold, consistent with our asymptotic theoretical results. In the bottom panel (high-risk screening), CSB maintains validity even with a small censoring training set, but its power (i.e., ability to identify appropriate patients) improves as the censoring model quality increases. Figures 9 and 10 in Appendix B present qualitatively similar results for the relatively easier Settings 2 and 3.

Additional Experiments. Additional experiments studying the effect of the number of features used to fit the censoring model are presented in Figures 11 and 12 in Appendix B. These results show that when too many features are included—especially with limited sample sizes—performance may degrade due to challenges in accurately estimating the censoring model, consistent with the theory. Tables 4–9 in Appendix B report on additional experiments using different survival and censoring models, leading to qualitatively similar results.

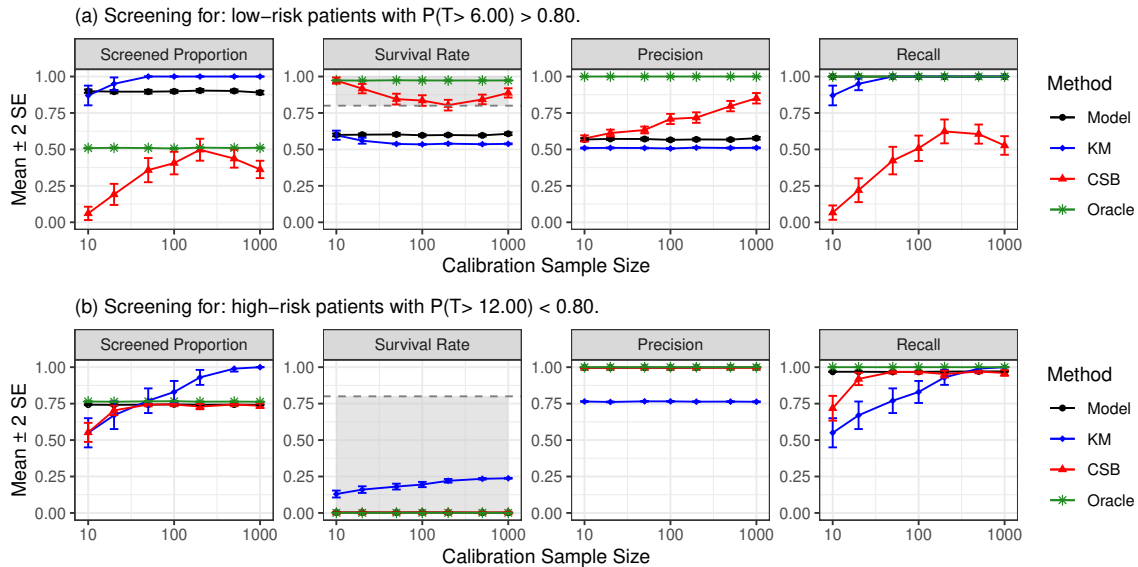


Figure 3: Effect of calibration sample size on patient screening performance with conformal survival bands (CSB) in a challenging synthetic data scenario with complex survival and censoring distributions, as in Figure 2. The training sample size is fixed at 5000. The CSB method successfully achieves survival rates above the target $p = 0.8$ for low-risk screening and below $p = 0.8$ for high-risk screening, while making selections that more closely resemble those of an ideal oracle as the calibration sample size increases.

Effect of Distribution Shift. Finally, we evaluate robustness under a distribution shift between the training and calibration/test distributions, corresponding to Setting 4 described in Table 3. Specifically, we now allow the survival distribution to change between training and calibration/test phases, which may lead to model misspecification even if the survival model fits the training data well. In these experiments, the training set consists of 5000 samples and the calibration set consists of 500 samples.

Table 1 summarizes the performance of different methods for low-risk screening under this setting. The “low-quality model” refers to a survival model trained on the shifted training distribution, resulting in inaccurate predictions on the calibration and test distributions. In this case, the *model* baseline selects all patients as “low risk” at time $t = 3$, but the observed survival rate is only about 54%, far below the target threshold of $p = 0.80$. By contrast, the *KM* and *CSB* methods conservatively select no patients, thereby maintaining validity at the cost of reduced power. With a “high-quality model” trained directly on the calibration/test distribution, the *CSB* method successfully screens approximately half of the patients while maintaining valid survival rates, closely approaching oracle performance. Meanwhile, the *KM* method remains overly conservative due to its lack of adaptivity, continuing to select no patients even when safe screening is feasible. Figure 1, introduced in Section 1, illustrates survival curves, conformal survival bands, and screening decisions corresponding to these experiments for four example patients in the low-risk setting.

Additional results for high-risk screening under distribution shift are presented in Table 10 in Appendix B. In this case, model misspecification primarily reduces the power of all

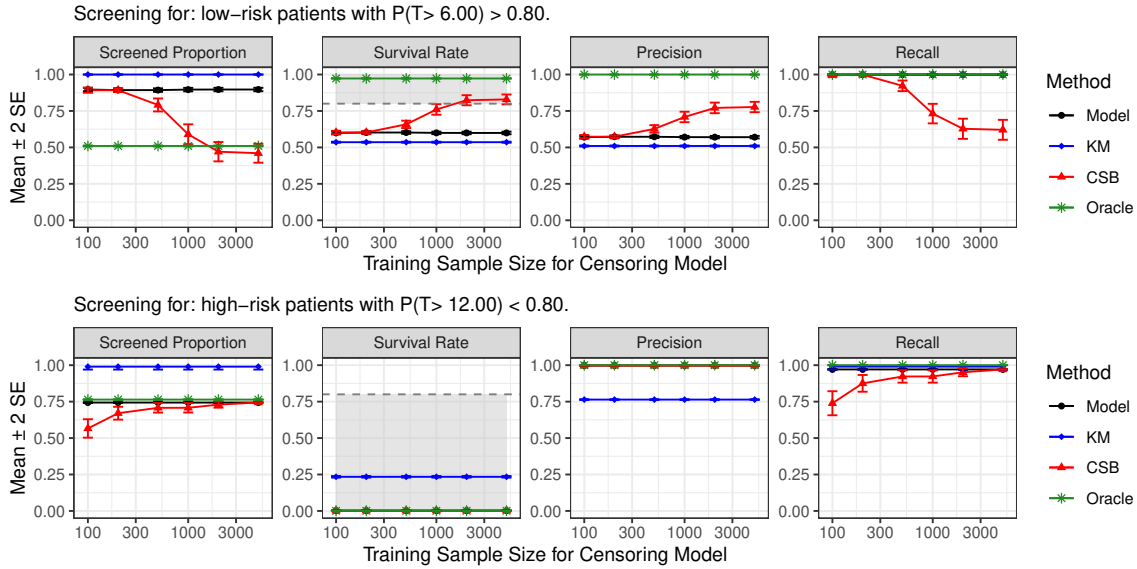


Figure 4: Effect of the training sample size used for fitting the censoring model on patient screening performance with conformal survival bands (CSB) in a challenging synthetic data scenario with complex survival and censoring distributions, as in Figure 2. The overall training sample size is fixed at 5000, but only a subset is used to fit the censoring model. CSB tends to produce higher-quality selections as the censoring model improves with more training data, achieving survival rates on the correct side of the target threshold (top), and improving power toward oracle performance (bottom).

Table 1: Performance of different methods for screening low-risk patients with $P(T > 3) > 0.80$ under distribution shift (Setting 4 from Table 3). Results are shown separately for a low-quality survival model trained on the shifted training distribution and a high-quality model trained directly on the calibration/test distribution. The survival rate highlighted in red illustrates how a misspecified *grf* model can lead to invalid screening rules, resulting in substantially lower survival rates than expected among patients labeled as “low-risk.”

Method	Screened	Survival	Precision	Recall
Low-Quality Model				
Model	1.000 ± 0.000	0.538 ± 0.003	0.499 ± 0.003	1.000 ± 0.000
KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
Oracle	0.499 ± 0.003	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
High-Quality Model				
Model	0.494 ± 0.004	1.000 ± 0.000	1.000 ± 0.000	0.994 ± 0.001
KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
CSB	0.492 ± 0.004	1.000 ± 0.000	1.000 ± 0.000	0.989 ± 0.001
Oracle	0.497 ± 0.004	1.000 ± 0.000	1.000 ± 0.000	1.000 ± 0.000

methods, leading to fewer selections. A similar figure illustrating survival curves, conformal bands, and screening decisions for the high-risk case is provided in Appendix B as Figure 13.

4. Application to Real Data

4.1. Setup

We apply our method to the same seven publicly available datasets analyzed by Sesia and Svetnik (2025): COLON, GBSG, HEART, METABRIC, PBC, RETINOPATHY, and VALCT. Details regarding the number of observations, covariates, and data sources are provided in Table 11 in Appendix C. Additional information about the standard preprocessing procedures applied to these datasets—designed to ensure compatibility with various survival and censoring models—is also available in Appendix C.

We follow the procedure from Section 3, comparing *CSB* to the *model* and *KM* benchmarks. Since the true data-generating distribution is unknown, the oracle method is not applicable. All methods are applied separately using each of the same four model types from Section 3 to estimate the survival distribution—generalized random forests (*grf*), parametric survival regression (*survreg*), random survival forests (*rf*), and Cox proportional hazards models (*cox*)—with the censoring distribution always estimated using *grf*. Each dataset is split into 60% training, 20% calibration, and 20% testing sets, and all experiments are repeated 100 times with independent random splits.

As in the synthetic experiments, we screen test patients using rules of the form $P[T > t] > p$ (low-risk) or $P[T > t] < p$ (high-risk), evaluated at fixed time thresholds t and probability levels p . We consider four tasks: low- and high-risk selection with $p = 0.8$ at time t_1 , and with $p = 0.25$ at t_2 . To facilitate comparisons while accounting for heterogeneity in the distribution of survival times across data sets, the time thresholds t_1 and t_2 are set to the 0.1 and 0.9 quantiles of all observed times, respectively.

Performance is evaluated on the test set using two metrics. The *screened proportion* measures the fraction of test patients selected for each screening task, and should ideally be large under accurate selection. The *survival rate* is the proportion of selected patients who survive beyond time t . Since this cannot be computed exactly due to censoring, we report deterministic bounds: the lower bound treats censored patients as failures at their censoring times, and the upper bound assumes they survive indefinitely. Together, these provide a conservative interval for the true survival rate among selected patients.

4.2. Results

Table 2 summarizes the results of this data analysis, aggregating the performance of each screening method across tasks, datasets, and repetitions, separately for each survival model. For each combination of survival model and screening method, we report the average proportion of test patients selected (*screened proportion*) and the empirical distribution of verification outcomes: *valid*, *dubious*, and *invalid*. Verification outcomes are determined for each task by assessing whether the survival rate bounds, averaged over 100 random repetitions, fall entirely on the correct side of the threshold p (valid), straddle it (dubious), or lie entirely on the wrong side (invalid), accounting for uncertainty via two standard errors.

Table 2: Summary of screening performance for high- and low-risk patient selection methods based on different survival models, aggregated across datasets, screening tasks, and repetitions. The *screened proportion* denotes the average fraction of test patients selected. Although false positives cannot be directly verified due to censoring, approximate verification is possible and classified as *valid*, *dubious*, or *invalid*. Note that the aggregated screened proportion is always 0.5 for the *Model* and *KM* benchmarks, which do not account for uncertainty and always assign each patient to either the high- or low-risk group for each pair (p, t) of screening parameters. In contrast, CSB may leave some patients unclassified—neither high- nor low-risk—when the (p, t) point falls within the “gray area” uncertainty band (e.g., see Figure 1), resulting in lower and more variable aggregated screening rates.

Method	Survival Model	Screened Proportion	Verification Outcome		
			Valid	Dubious	Invalid
Model	grf	0.500	0.643	0.357	0.000
	survreg	0.500	0.714	0.250	0.036
	rf	0.500	0.536	0.464	0.000
	Cox	0.500	0.571	0.393	0.036
KM	grf	0.500	0.786	0.214	0.000
	survreg	0.500	0.786	0.214	0.000
	rf	0.500	0.786	0.214	0.000
	Cox	0.500	0.786	0.214	0.000
CSB	grf	0.293	0.929	0.071	0.000
	survreg	0.245	0.964	0.036	0.000
	rf	0.301	0.857	0.143	0.000
	Cox	0.299	0.857	0.143	0.000

CSB yields the highest proportion of valid selections across all survival models, with few or no dubious or invalid outcomes, despite screening fewer patients. In contrast, the *Model* and *KM* benchmarks screen more patients but exhibit lower verification quality. These results illustrate a fundamental trade-off between power and reliability, unsurprisingly.

Compared to the synthetic data experiments from Section 3, the advantage of CSB is somewhat less obvious here. This may be due in part to imperfect verification, stemming from censoring in the test data and the absence of oracle knowledge. Moreover, survival models may estimate the conditional survival distributions more accurately here than in the synthetic settings, reducing the relative gains from conformal inference. As noted by [Sesia and Svetnik \(2025\)](#), these datasets do not appear to be especially difficult to model. Nevertheless, in general, it remains difficult to assess the reliability of screening results produced by black-box methods. Our approach can provide more confidence in the calibration of such screening results, which may be particularly valuable in high-stakes applications where the cost of false positives is substantial, even if it entails some reduction in power.

Additional breakdowns of these results for specific screening tasks are provided in Appendix C, with low-risk and high-risk summaries shown in Tables 12 and 13, respectively.

5. Discussion

This paper introduced a conformal inference method for constructing uncertainty bands around individual survival curves under right-censoring, enabling statistically principled personalized risk screening under minimal assumptions on the data-generating process.

Experiments with synthetic data showed that screening based on uncalibrated black-box models can be unreliable, particularly in hard-to-model settings, whereas our method provides greater robustness, albeit with some conservativeness. In our real data applications, standard survival models seemed to produce reasonable screening performance, likely because the fitted models were able to approximate the true survival distribution quite well. Nonetheless, our method remains appealing in more complex or high-stakes scenarios, where model misspecification is a concern and formal uncertainty quantification is essential.

Two limitations of the current approach are its reliance on asymptotic FDR control and its focus on pointwise inference, which assumes that the screening thresholds t and p are fixed in advance. While it may be possible to obtain finite-sample and simultaneous inference guarantees, doing so would likely require a significantly more conservative method. Understanding this trade-off between statistical rigor and screening power in greater depth presents an intriguing direction for future research.

An additional limitation is that our method can be conservative in some settings, resulting in fewer selections than the ideal oracle or even fewer than the black-box survival model when that model appears well calibrated. A possible way to mitigate this over-conservativeness in future work would be to incorporate an adaptive estimator of the proportion of null hypotheses (Gao, 2024), similar to how Storey’s method extends the BH procedure (Storey et al., 2004).

Another promising direction for improvement concerns the treatment of censored calibration points. In Section 2.4, we proposed reallocating these samples to the training set but did not do this in practice due to concerns about sampling bias potentially degrading model quality. Recent work by Farina et al. (2025) introduces a more complex strategy for incorporating censored observations directly into the calibration phase, which may be possible to adapt to our setting. A potentially simpler alternative is to reallocate censored calibration samples to the training set and apply importance weighting during training to mitigate bias, while retaining the calibration procedure described in this paper, which uses only uncensored observations. Both strategies deserve further investigation.

Additional extensions include de-randomizing our inferences via model aggregation across different data splits (Carlsson et al., 2014; Linusson et al., 2017), potentially using e-values (Vovk and Wang, 2023; Wang and Ramdas, 2022; Bashari et al., 2023), and adapting the method to handle noisy or corrupted data (Sesia et al., 2024; Clarkson et al., 2024).

Software Availability

Open-source software implementing our methods and experiments is available at https://github.com/msesia/conformal_survival_screening.

Acknowledgements

The authors thank two anonymous referees for detailed feedback and valuable suggestions on an earlier version of this manuscript. M. S. is partly supported by NSF grant DMS 2210637, a Capital One CREDIF Research Award, and a Google Research Scholar Award.

References

- Rina Foygel Barber. Is distribution-free inference possible for binary regression? *Electron. J. Statist.*, 14(2):3487–3524, 2020.
- Rina Foygel Barber, Emmanuel J Candès, Aaditya Ramdas, and Ryan J Tibshirani. Conformal prediction beyond exchangeability. *Ann. Stat.*, 51(2):816–845, 2023.
- Avinash Barnwal, Hyunsu Cho, and Toby Hocking. Survival regression with accelerated failure time model in xgboost. *J. Comput. Graph. Stat.*, 31(4):1292–1302, 2022.
- Meshi Bashari, Amir Epstein, Yaniv Romano, and Matteo Sesia. Derandomized novelty detection with fdr control via conformal e-values. In *Advances in Neural Information Processing Systems*, volume 36, pages 65585–65596, 2023.
- Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *J. R. Stat. Soc. Ser. B Methodol.*, 57(1):289–300, 1995.
- AL Blair, DR Hadden, JA Weaver, DB Archer, PB Johnston, and CJ Maguire. The 5-year prognosis for vision in diabetes. *The Ulster medical journal*, 49(2):139, 1980.
- Henrik Boström, Lars Asker, Ram Gurung, Isak Karlsson, Tony Lindgren, and Panagiotis Papapetrou. Conformal prediction using random survival forests. In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 812–817. IEEE, 2017.
- Henrik Boström, Ulf Johansson, and Anders Vesterberg. Predicting with confidence from survival data. In *Conformal and Probabilistic Prediction and Applications*, pages 123–141. PMLR, 2019.
- Henrik Boström, Henrik Linusson, and Anders Vesterberg. Mondrian predictive systems for censored data. In *Proceedings of the Twelfth Symposium on Conformal and Probabilistic Prediction with Applications*, volume 204 of *Proceedings of Machine Learning Research*, pages 399–412. PMLR, 13–15 Sep 2023.
- Emmanuel Candès, Lihua Lei, and Zhimei Ren. Conformalized survival analysis. *J. R. Stat. Soc. Ser. B Methodol.*, 85(1):24–45, 2023.
- Lars Carlsson, Martin Eklund, and Ulf Norinder. Aggregated conformal prediction. In *Artificial Intelligence Applications and Innovations: AIAI 2014 Workshops: CoPA, MHDW, IIVC, and MT4BD, Rhodes, Greece, September 19-21, 2014. Proceedings 10*, pages 231–240. Springer, 2014.

- Jason Clarkson, Wenkai Xu, Mihai Cucuringu, and Gesine Reinert. Split conformal prediction under data contamination. In *Proceedings of the Thirteenth Symposium on Conformal and Probabilistic Prediction with Applications*, volume 230 of *Proceedings of Machine Learning Research*, pages 5–27. PMLR, 09–11 Sep 2024.
- David R Cox. Regression models and life-tables. *J. R. Stat. Soc. Ser. B Methodol.*, 34(2): 187–202, 1972.
- John Crowley and Marie Hu. Covariance analysis of heart transplant survival data. *J. Am. Stat. Assoc.*, 72(357):27–36, 1977.
- Christina Curtis, Sohrab P Shah, Suet-Feung Chin, Gulisa Turashvili, Oscar M Rueda, Mark J Dunning, Doug Speed, Andy G Lynch, Shamith Samarajiwa, Yinyin Yuan, et al. The genomic and transcriptomic architecture of 2,000 breast tumours reveals novel subgroups. *Nature*, 486(7403):346–352, 2012.
- Hen Davidov, Shai Feldman, Gil Shamai, Ron Kimmel, and Yaniv Romano. Conformalized survival analysis for general right-censored data. In *The Thirteenth International Conference on Learning Representations*, 2024.
- Rebecca Farina, Eric J Tchetgen Tchetgen, and Arun Kumar Kuchibhotla. Doubly robust and efficient calibration of prediction sets for censored time-to-event outcomes. *arXiv preprint arXiv:2501.04615*, 2025.
- Matteo Fontana, Gianluca Zeni, and Simone Vantini. Conformal prediction: a unified review of theory and new challenges. *Bernoulli*, 29(1):1–23, 2023.
- Zijun Gao. An adaptive null proportion estimator for false discovery rate control. *Biometrika*, 112(1):asae051, 10 2024. ISSN 1464-3510.
- Yu Gui, Rohan Hore, Zhimei Ren, and Rina Foygel Barber. Conformalized survival analysis with adaptive cut-offs. *Biometrika*, 111(2):459–477, 2024.
- Hemant Ishwaran, Udaya B. Kogalur, Eugene H. Blackstone, and Michael S. Lauer. Random survival forests. *Ann. Appl. Stat.*, 2(3):841 – 860, 2008.
- Ying Jin and Emmanuel J Candès. Model-free selective inference under covariate shift via weighted conformal p-values. *arXiv preprint arXiv:2307.09291*, 2023a.
- Ying Jin and Emmanuel J Candès. Selection by prediction with conformal p-values. *J. Mach. Learn. Res.*, 24(244):1–41, 2023b.
- John D Kalbfleisch and Ross L Prentice. *The statistical analysis of failure time data*. John Wiley & Sons, 2002.
- Edward L Kaplan and Paul Meier. Nonparametric estimation from incomplete observations. *J. Am. Stat. Assoc.*, 53(282):457–481, 1958.
- Jared L Katzman, Uri Shaham, Alexander Cloninger, Jonathan Bates, Tingting Jiang, and Yuval Kluger. Deepsurv: personalized treatment recommender system using a Cox proportional hazards deep neural network. *BMC Med. Res. Methodol.*, 18:1–12, 2018.

- Henrik Linusson, Ulf Norinder, Henrik Boström, Ulf Johansson, and Tuve Löfström. On the calibration of aggregated conformal predictors. In *Conformal and probabilistic prediction and applications*, pages 154–173. PMLR, 2017.
- Charles G Moertel, Thomas R Fleming, John S Macdonald, Daniel G Haller, John A Laurie, Phyllis J Goodman, James S Ungerleider, William A Emerson, Douglas C Tormey, John H Glick, et al. Levamisole and fluorouracil for adjuvant therapy of resected colon carcinoma. *New England Journal of Medicine*, 322(6):352–358, 1990.
- Shi-ang Qi, Yakun Yu, and Russell Greiner. Conformalized survival distributions: A generic post-process to increase calibration. *arXiv preprint arXiv:2405.07374*, 2024.
- Jing Qin, Jin Piao, Jing Ning, and Yu Shen. Conformal predictive intervals in survival analysis: a re-sampling approach. *arXiv preprint arXiv:2408.06539*, 2024.
- James M Robins. Information recovery and bias adjustment in proportional hazards regression analysis of randomized trials using surrogate markers. In *Proc. Biopharm. Sect. Am. Stat. Assoc.*, volume 24, pages 24–33. San Francisco CA, 1993.
- James M Robins and Andrea Rotnitzky. Recovery of information and adjustment for dependent censoring using surrogate markers. In *AIDS epidemiology: methodological issues*, pages 297–331. Springer, 1992.
- Matteo Sesia and Vladimir Svetnik. Doubly robust conformalized survival analysis with right-censored data. In *International Conference on Machine Learning*. PMLR, 2025.
- Matteo Sesia, YX Rachel Wang, and Xin Tong. Adaptive conformal classification with noisy labels. *J. R. Stat. Soc. Ser. B Methodol.*, page qkae114, 2024.
- Annette Spooner, Emily Chen, Arcot Sowmya, Perminder Sachdev, Nicole A Kochan, Julian Trollor, and Henry Brodaty. A comparison of machine learning methods for survival analysis of high-dimensional clinical data for dementia prediction. *Scientific reports*, 10(1):20410, 2020.
- John D Storey, Jonathan E Taylor, and David Siegmund. Strong control, conservative point estimation and simultaneous conservative consistency of false discovery rates: a unified approach. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 66(1):187–205, 2004.
- Xiaolin Sun and Yanhua Wang. Conformal prediction with censored data using kaplan-meier method. In *Journal of Physics: Conference Series*, volume 2898, page 012030. IOP Publishing, 2024.
- Terry M Therneau, Patricia M Grambsch, Terry M Therneau, and Patricia M Grambsch. *The Cox model*. Springer, 2000.
- Ryan J Tibshirani, Rina Foygel Barber, Emmanuel Candès, and Aaditya Ramdas. Conformal prediction under covariate shift. *Advances in neural information processing systems*, 32, 2019.

- Vladimir Vovk and Ruodu Wang. Confidence and discoveries with e-values. *Statistical Science*, 38(2):329–354, 2023.
- Vladimir Vovk, Alexander Gammerman, and Glenn Shafer. *Algorithmic learning in a random world*, volume 29. Springer, 2005.
- Ruodu Wang and Aaditya Ramdas. False discovery rate control with e-values. *J. R. Stat. Soc. Ser. B Methodol.*, 84(3):822–852, 2022.
- Menghan Yi, Ze Xiao, Huixia Judy Wang, and Yanlin Tang. Survival conformal prediction under random censoring. *Stat*, 14(2):e70052, 2025.

Appendix A. Mathematical Proofs

A.1. Proof of Proposition 1

Proof [of Proposition 1]

Under the null hypothesis that $T_{n+j} \geq t$, because $\hat{s}_{\text{lt}}(t'; x, t)$ is monotone increasing in t' for all x , we have that, almost surely

$$\hat{s}_{\text{lt}}(t; X_{n+j}, t) \leq \hat{s}_{\text{lt}}(T_{n+j}; X_{n+j}, t),$$

and thus

$$\begin{aligned} \tilde{\phi}_{\text{lt}}(t; X_{n+j}) &:= \frac{1 + \sum_{i=1}^n \mathbb{I}\{\hat{s}_{\text{lt}}(T_i; X_i, t) \geq \hat{s}_{\text{lt}}(t; X_{n+j}, t)\}}{n+1} \\ &\geq \frac{1 + \sum_{i=1}^n \mathbb{I}\{\hat{s}_{\text{lt}}(T_i; X_i, t) \geq \hat{s}_{\text{lt}}(T_{n+j}; X_{n+j}, t)\}}{n+1}. \end{aligned}$$

Hence

$$\begin{aligned} \mathbb{P}\left[\tilde{\phi}_{\text{lt}}(t; X_{n+j}) \leq \alpha, T_{n+j} \leq t\right] &\leq \mathbb{P}\left[\frac{1 + \sum_{i=1}^n \mathbb{I}\{\hat{s}_{\text{lt}}(T_i; X_i, t) \geq \hat{s}_{\text{lt}}(T_{n+j}; X_{n+j}, t)\}}{n+1} \leq \alpha\right] \\ &\leq \alpha, \end{aligned}$$

where the last inequality follows immediately through standard conformal inference arguments because $(\hat{s}_{\text{lt}}(T_1; X_1, t), \dots, \hat{s}_{\text{lt}}(T_n; X_n, t), \hat{s}_{\text{lt}}(T_{n+j}; X_{n+j}, t))$ are exchangeable. ■

A.2. Proofs of Theorems 2 and 3

Proof [of Theorem 2: Asymptotic validity of IPCW p-values] We apply the finite-sample bound from Theorem 5, which states that for any $t > 0$ and any $n \geq 1$,

$$\mathbb{P}\left[\hat{\phi}_{\text{lt}}(t; X_{n+1}) \leq \alpha, T_{n+1} \geq t\right] \leq \alpha + \frac{1}{\omega_{\min}\pi} \left[\frac{2(2\omega_{\min} + 1)}{n} + 8\sqrt{\frac{2\log(2n)}{n}} \right] + 2\Delta_N.$$

By assumption:

- $\hat{w}(T; X) \geq \omega_{\min} > 0$ almost surely,
- $\Delta_N := \left(\mathbb{E} \left[\left(\frac{1}{\hat{w}(T; X)} - \frac{1}{w^*(T; X)} \right)^2 \right] \right)^{1/2} \rightarrow 0$ as $N \rightarrow \infty$.

Since the remaining terms in the bound all vanish as $N, n \rightarrow \infty$, it follows that

$$\limsup_{N, n \rightarrow \infty} \mathbb{P}\left[\hat{\phi}_{\text{lt}}(t; X_{n+1}) \leq \alpha, T_{n+1} \geq t\right] \leq \alpha.$$

An identical argument applies to the limit involving $\hat{\phi}_{\text{rt}}(t; X_{n+1})$, completing the proof. ■

Proof [of Theorem 3: Asymptotic FDR control] We prove the result for the left-tail p-values $\hat{\phi}_{\text{lt}}(t; X_{n+j})$; the argument for right-tail p-values is analogous.

The proof proceeds in three steps:

- (1) We compare each empirical p-value to an oracle counterpart and establish uniform closeness.
- (2) We show that the BH rejection sets based on these two sets of p-values are close with high probability.
- (3) We conclude FDR control for the empirical procedure via stability and the oracle FDR guarantee.

Step 1: Empirical p-values are close to oracle p-values. For each test point $j \in \{1, \dots, m\}$, define the empirical and oracle p-values as:

$$p_j := \hat{\phi}_{\text{lt}}(t; X_{n+j}), \quad p_j^* := \phi_{\text{lt}}^*(t; X_{n+j}) := \mathbb{P}(\hat{s}_{\text{lt}}(T; X, t) \geq \hat{s}_{\text{lt}}(t; X_{n+j}, t)),$$

where the probability in p_j^* is taken over an independent copy $(T, X) \sim P$.

By Corollary 6, for any $\delta > 0$, with probability at least $1 - \delta$ over the calibration data,

$$\max_{1 \leq j \leq m} |p_j - p_j^*| \leq \epsilon_n,$$

where

$$\epsilon_n = \frac{1}{\omega_{\min} \pi} \left[\frac{2(\omega_{\min} + 1)}{n} + 8 \sqrt{\frac{\log(2n/\delta)}{n}} \right] + 2\Delta_N.$$

Step 2: Stability of BH rejection sets. Let $\mathcal{R} := \{j : p_j \leq \hat{\tau}\}$ and $\mathcal{R}^* := \{j : p_j^* \leq \hat{\tau}^*\}$ be the rejection sets obtained by applying the BH procedure at level α to the empirical and oracle p-values, respectively.

Before we can apply Lemma 8 (a general stability bound for the BH procedure), we need to verify that the oracle p-values p_j^* have a density bounded above. Since

$$p_j^* = \phi_{\text{lt}}^*(t; X_{n+j}) = 1 - F(\hat{s}_{\text{lt}}(t; X_{n+j}, t)),$$

where F is the CDF of the score variable $\hat{s}_{\text{lt}}(T; X, t)$, with corresponding density $f(u) = F'(u)$ uniformly bounded above and below across all $u > 0$. Then, the change-of-variables formula gives that the density f^* of p_j^* is:

$$f^*(u) = \frac{f_t(F^{-1}(1-u))}{f(F^{-1}(1-u))},$$

which, thanks to Assumptions (A2) and (A5), we know to be bounded from both above and below by:

$$\frac{f_{\min}}{f_{\max}} \leq f^*(u) \leq \frac{f_{\max}}{f_{\min}} \quad \text{for all } u \in [0, 1].$$

Now we can apply Lemma 8 with $p = (p_1, \dots, p_m)$ and $q = (p_1^*, \dots, p_m^*)$, obtaining:

$$\mathbb{P}(\mathcal{R} \neq \mathcal{R}^*) \leq \delta + 2\epsilon_n m^2 \cdot \frac{f_{\max}}{f_{\min}}.$$

Step 3: Transferring FDR control from the oracle to empirical procedure. By Lemma 4, the oracle p-values p_j^* satisfy the super-uniformity property:

$$\mathbb{P}(p_j^* \leq \alpha, T_{n+j} \geq t) \leq \alpha, \quad \text{for all } \alpha \in (0, 1),$$

and the p-values p_j^* are mutually independent (conditional on the calibration data). Therefore, by Theorem 9, BH applied to (p_1^*, \dots, p_m^*) controls the FDR below α :

$$\text{FDR}^* := \mathbb{E} \left[\frac{|\mathcal{R}^* \cap \mathcal{H}_0|}{|\mathcal{R}^*| \vee 1} \right] \leq \alpha,$$

where $\mathcal{H}_0 := \{j : T_{n+j} \geq t\}$ is the (random) set of true nulls.

Conclusion. By definition of the FDR and triangle inequality,

$$|\text{FDR} - \text{FDR}^*| \leq \mathbb{P}(\mathcal{R} \neq \mathcal{R}^*).$$

Combining this with the oracle guarantee $\text{FDR}^* \leq \alpha$ and the stability bound from Lemma 8, we obtain the finite-sample bound:

$$\text{FDR} \leq \alpha + \delta + 2\epsilon_n m^2 \cdot \frac{f_{\max}}{f_{\min}}.$$

Letting $\delta \rightarrow 0$ and noting that $\epsilon_n \rightarrow 0$ as $n \rightarrow \infty$, we conclude:

$$\limsup_{N, n \rightarrow \infty} \text{FDR}_n \leq \alpha \quad \text{for fixed } m.$$

Moreover, if $m = m_n \rightarrow \infty$ and $m_n^2 \epsilon_n \rightarrow 0$, then the finite-sample error term vanishes:

$$\delta + 2\epsilon_n m^2 \cdot \frac{f_{\max}}{f_{\min}} \rightarrow 0,$$

and we conclude the joint limit:

$$\limsup_{N, m, n \rightarrow \infty} \text{FDR}_{m, n} \leq \alpha.$$

■

A.3. Auxiliary Results

A.3.1. VALIDITY OF ORACLE P-VALUES

Lemma 4 *Let (X, T) and (X', T') be independent copies of a pair of random variables taking values in $\mathcal{X} \times \mathbb{R}$, and let $s : \mathbb{R} \times \mathcal{X} \rightarrow \mathbb{R}$ be a fixed function that is monotone increasing in its first argument. Define the oracle p-value function*

$$\phi^*(t, x) := \mathbb{P}(s(T, X) \geq s(t, x)),$$

where the probability is taken over (X, T) . Then, for any $\alpha \in [0, 1]$,

$$\mathbb{P}(\phi^*(t, X') \leq \alpha, T' \geq t) \leq \alpha.$$

Proof [of Lemma 4] Let $S := -s(T, X)$ and $S' := -s(T', X')$, and define the cumulative distribution function $F_S(u) := \mathbb{P}(S \leq u)$. By definition of ϕ^* , we can write:

$$\phi^*(t, X') = \mathbb{P}(S \leq -s(t, X')) = F_S(-s(t, X')).$$

Now observe that, under the condition $T' \geq t$ and monotonicity of s in t , we have:

$$s(t, X') \leq s(T', X') \quad \Rightarrow \quad -s(t, X') \geq -s(T', X'),$$

and since F_S is non-decreasing we have the almost-surely inequality:

$$\phi^*(t, X') = F_S(-s(t, X')) \geq F_S(-s(T', X')) = F_S(S').$$

Hence:

$$\mathbb{P}(\phi^*(t, X') \leq \alpha, T' \geq t) \leq \mathbb{P}(F_S(S') \leq \alpha, T' \geq t) \leq \mathbb{P}(F_S(S') \leq \alpha) = \alpha,$$

where the final equality uses the probability integral transform: since S' has the same distribution as S and F_S is its CDF, the random variable $F_S(S')$ is uniform on $[0, 1]$. ■

A.3.2. FINITE-SAMPLE VERSION OF THEOREM 2

Theorem 5 *Under the assumptions of Theorem 2,*

$$\mathbb{P}\left[\hat{\phi}_{\text{lt}}(t; X_{n+1}) \leq \alpha, T_{n+1} \geq t\right] \leq \alpha + \frac{1}{\omega_{\min}\pi} \left[\frac{2(2\omega_{\min} + 1)}{n} + 8\sqrt{\frac{2\log(2n\sqrt{n})}{n}} \right] + 2\Delta_N,$$

for any $t > 0$, and

$$\mathbb{P}\left[\hat{\phi}_{\text{rt}}(t; X_{n+1}) \leq \alpha, T_{n+1} \leq t\right] \leq \alpha + \frac{1}{\omega_{\min}\pi} \left[\frac{2(2\omega_{\min} + 1)}{n} + 8\sqrt{\frac{2\log(2n\sqrt{n})}{n}} \right] + 2\Delta_N.$$

Proof [of Theorem 5] We prove the result for $\hat{\phi}_{\text{lt}}$; the argument for $\hat{\phi}_{\text{rt}}$ is identical and omitted.

For $\delta \in (0, 1)$, define the event

$$\mathcal{E}_{t,\delta} := \left\{ \sup_{x \in \mathbb{R}^d} \left| \hat{\phi}_{\text{lt}}(t; x) - \phi_{\text{lt}}^*(t; x) \right| \leq \epsilon_n(\delta) \right\},$$

$$\epsilon_n(\delta) := \frac{1}{\omega_{\min}\pi} \left[\frac{2(\omega_{\min} + 1)}{n} + 8\sqrt{\frac{\log(2n/\delta)}{n}} \right] + 2\Delta_N.$$

Corollary 6 implies:

$$\begin{aligned} & \mathbb{P}\left[\hat{\phi}_{\text{lt}}(t; x) \leq \alpha, T_{n+1} \geq t \mid X_{n+1} = x\right] \\ &= \mathbb{P}\left[\hat{\phi}_{\text{lt}}(t; x) \leq \alpha, \mathcal{E}_{t,\delta}, T_{n+1} \geq t \mid X_{n+1} = x\right] + \mathbb{P}\left[\hat{\phi}_{\text{lt}}(t; x) \leq \alpha, \mathcal{E}_{t,\delta}^c, T_{n+1} \geq t \mid X_{n+1} = x\right] \\ &\leq \mathbb{P}\left[\phi_{\text{lt}}^*(t; x) \leq \alpha + \epsilon_n(\delta), T_{n+1} \geq t \mid X_{n+1} = x\right] + \delta. \end{aligned}$$

Therefore,

$$\begin{aligned} \mathbb{P} \left[\hat{\phi}_{\text{lt}}(t; x) \leq \alpha, T_{n+1} \geq t \right] &\leq \mathbb{P} [\phi_{\text{lt}}^*(t; x) \leq \alpha + \epsilon_n(\delta), T_{n+1} \geq t] + \delta \\ &\leq \alpha + \epsilon_n(\delta) + \delta, \end{aligned}$$

where the second inequality above follows directly from Lemma 4. Finally, setting $\delta = 1/n$:

$$\mathbb{P} \left[\hat{\phi}_{\text{lt}}(t; X_{n+1}) \leq \alpha, T_{n+1} \geq t \right] \leq \alpha + \frac{1}{\omega_{\min} \pi} \left[\frac{2(2\omega_{\min} + 1)}{n} + 8\sqrt{\frac{2\log(2n)}{n}} \right] + 2\Delta_N.$$

■

Corollary 6 (Specialized Instance of Theorem 7) *Under the assumptions of Theorem 2, for any $x \in \mathbb{R}^d$ and $t > 0$, define:*

$$\begin{aligned} \phi_{\text{lt}}^*(t; x) &:= \mathbb{P}(\hat{s}_{\text{lt}}(T; X, t) \geq \hat{s}_{\text{lt}}(t; x, t)), \\ \phi_{\text{rt}}^*(t; x) &:= \mathbb{P}(\hat{s}_{\text{rt}}(T; X, t) \geq \hat{s}_{\text{rt}}(t; x, t)), \end{aligned}$$

where the probabilities are taken with respect to the randomness in (X, T) .

Then, for any $t > 0$ and $\delta \in (0, 1)$, with probability at least $1 - \delta$ over the randomness in the calibration data $\mathcal{D} = \{(X_i, T_i, C_i)\}_{i=1}^n$:

$$\sup_{x \in \mathbb{R}^d} \left| \hat{\phi}_{\text{lt}}(t; x) - \phi_{\text{lt}}^*(t; x) \right| \leq \frac{1}{\omega_{\min} \pi} \left[\frac{2(\omega_{\min} + 1)}{n} + 8\sqrt{\frac{\log(2n/\delta)}{n}} \right] + 2\Delta_N.$$

Similarly, for any $t > 0$ and $\delta \in (0, 1)$, with probability at least $1 - \delta$:

$$\sup_{x \in \mathbb{R}^d} \left| \hat{\phi}_{\text{rt}}(t; x) - \phi_{\text{rt}}^*(t; x) \right| \leq \frac{1}{\omega_{\min} \pi} \left[\frac{2(\omega_{\min} + 1)}{n} + 8\sqrt{\frac{\log(2n/\delta)}{n}} \right] + 2\Delta_N.$$

Proof [of Corollary 6] We prove the bound for $\hat{\phi}_{\text{lt}}$; the argument for $\hat{\phi}_{\text{rt}}$ is identical and omitted. Define the following variables:

$$\begin{aligned} Z_i &= (X_i, T_i), \\ D_i &= \mathbf{1}(T_i \leq C_i), \\ E_i &= \hat{w}(T_i; X_i), \\ e(Z_i) &= w^*(T_i; X_i), \\ \psi(Z_i) &= \hat{s}_{\text{lt}}(T_i; X_i, t). \end{aligned}$$

Recall the definitions of $R_n(u)$ and $R(u)$ from Theorem 7:

$$\begin{aligned} R_n(u) &:= \frac{1 + \sum_{i=1}^n (D_i/E_i) \cdot \mathbf{1}(\psi(Z_i) \geq u)}{1 + \sum_{i=1}^n (D_i/E_i)}, \\ R(u) &:= \mathbb{E} \left[\frac{D}{e(Z)} \cdot \mathbf{1}(\psi(Z) \geq u) \right]. \end{aligned}$$

For any $x \in \mathbb{R}^d$ and $t > 0$, define

$$\hat{\phi}_{\text{lt}}(t; x) := R_n(\hat{s}_{\text{lt}}(t; x, t)), \quad \phi_{\text{lt}}^*(t; x) := R(\hat{s}_{\text{lt}}(t; x, t)).$$

Therefore,

$$\begin{aligned} \sup_{x \in \mathbb{R}^d} \left| \hat{\phi}_{\text{lt}}(t; x) - \phi_{\text{lt}}^*(t; x) \right| &= \sup_{x \in \mathbb{R}^d} |R_n(\hat{s}_{\text{lt}}(t; x, t)) - R(\hat{s}_{\text{lt}}(t; x, t))| \\ &\leq \sup_{u \in \mathbb{R}} |R_n(u) - R(u)|, \end{aligned}$$

and the result follows directly from Theorem 7, noting that $E_i \geq \omega_{\min}$ and $\mathbb{E}[D] = \pi$. \blacksquare

A.3.3. UNIFORM CONCENTRATION FOR GENERAL REWEIGHTED SURVIVAL ESTIMATOR

Theorem 7 *Let $(Z_i, D_i, E_i)_{i=1}^n$ be i.i.d. samples from some unknown distribution \mathcal{P} , where $Z_i \in \mathbb{R}^d$, $D_i \in \{0, 1\}$, and $E_i \in [\underline{e}, 1]$ almost surely for some constant $\underline{e} > 0$. For any $z \in \mathbb{R}^d$, define $e(z) := \mathbb{P}(D_i = 1 \mid Z_i = z)$, which may be interpreted as a propensity score. For any fixed function $\psi : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ and any $u \in \mathbb{R}$, define*

$$R_n(u) := \frac{1 + \sum_{i=1}^n (D_i/E_i) \cdot \mathbf{1}(\psi(Z_i) \geq u)}{1 + \sum_{i=1}^n (D_i/E_i)}, \quad R(u) := \mathbb{E} \left[\frac{D}{e(X)} \cdot \mathbf{1}(\psi(Z) \geq u) \right],$$

where (Z, D, E) is a generic random sample from \mathcal{P} . Define also

$$\Delta_N := \left(\mathbb{E} \left[\left(\frac{1}{E} - \frac{1}{e(X)} \right)^2 \right] \right)^{1/2}, \quad \pi := \mathbb{E}[D].$$

Then for any $\delta \in (0, 1)$, with probability at least $1 - \delta$,

$$\sup_{u \in \mathbb{R}} |R_n(u) - R(u)| \leq \frac{1}{\underline{e}\pi} \left[\frac{2(\underline{e} + 1)}{n} + 8\sqrt{\frac{\log(2n/\delta)}{n}} \right] + 2\Delta_N. \quad (16)$$

Proof [of Theorem 7] The proof splits the error into two parts: a stochastic deviation term (comparing the estimator $R_n(u)$ to a population analogue $\tilde{R}(u)$ based on the approximate weights E_i), and a bias term due to discrepancy between E_i and the propensity score $e(X_i)$.

Define

$$\tilde{R}(u) := \frac{\mathbb{E}[(D/E) \cdot \mathbf{1}(\psi(Z) \geq u)]}{\mathbb{E}[D/E]},$$

where the expectations are over the joint distribution of (X, D, Z, E) . Then decompose the error:

$$|R_n(u) - R(u)| \leq |R_n(u) - \tilde{R}(u)| + |\tilde{R}(u) - R(u)|.$$

Step 1 (Bounding the stochastic deviation). We begin by computing a uniform bound for $|R_n(u) - \tilde{R}(u)|$. Let $W_i := D_i/E_i \in [0, 1/\underline{e}]$ and $W := D/E \in [0, 1/\underline{e}]$. Define

$$\begin{aligned} N_n(u) &:= \frac{1 + \sum_{i=1}^n W_i \cdot \mathbf{1}(\psi(Z_i) \geq u)}{1 + n}, \\ A_n &:= \frac{1 + \sum_{i=1}^n W_i}{1 + n}, \\ \tilde{N}(u) &:= \mathbb{E}[W \cdot \mathbf{1}(\psi(Z) \geq u)], \\ \tilde{A} &:= \mathbb{E}[W], \end{aligned}$$

so that

$$R_n(u) = \frac{N_n(u)}{A_n}, \quad \tilde{R}(u) = \frac{\tilde{N}(u)}{\tilde{A}}.$$

Using the inequality

$$\left| \frac{a}{b} - \frac{a'}{b'} \right| \leq \frac{|a - a'|}{b'} + |a| \left| \frac{1}{b} - \frac{1}{b'} \right|,$$

with $a = N_n(u)$, $a' = \tilde{N}(u)$, $b = A_n$, $b' = \tilde{A}$, and the fact that $N_n(u) \leq A_n$, we obtain

$$\begin{aligned} |R_n(u) - \tilde{R}(u)| &\leq \frac{|N_n(u) - \tilde{N}(u)|}{\tilde{A}} + A_n \left| \frac{1}{A_n} - \frac{1}{\tilde{A}} \right| \\ &= \frac{|N_n(u) - \tilde{N}(u)| + |A_n - \tilde{A}|}{\tilde{A}} \\ &\leq \frac{|N_n(u) - \tilde{N}(u)| + |A_n - \tilde{A}|}{\pi}, \end{aligned} \tag{17}$$

because $E \leq 1$ and thus $\tilde{A} = \mathbb{E}[W] = \mathbb{E}[D/E] \geq \mathbb{E}[D] = \pi$.

Uniform concentration bound for $|N_n(u) - \tilde{N}(u)|$. To bound the deviation

$$\sup_{u \in \mathbb{R}} \left| \frac{1}{n} \sum_{i=1}^n W_i \mathbf{1}(\psi(Z_i) \geq u) - \mathbb{E}[W \cdot \mathbf{1}(\psi(Z) \geq u)] \right|,$$

note that the class $\{u \mapsto W_i \mathbf{1}(\psi(Z_i) \geq u)\}$ has envelope bounded by $1/\underline{e}$ and VC dimension 1. Applying Theorem 4.10 and Lemma 4.14 in Wainwright (2019), we obtain: with probability at least $1 - \delta$,

$$\sup_{u \in \mathbb{R}} \left| \frac{1}{n} \sum_{i=1}^n W_i \mathbf{1}(\psi(Z_i) \geq u) - \tilde{N}(u) \right| \leq \frac{4}{\underline{e}} \sqrt{\frac{\log(n+1)}{n}} + \sqrt{\frac{2 \log(1/\delta)}{n \underline{e}^2}}.$$

Now recall that

$$N_n(u) = \frac{1}{1+n} + \frac{1}{1+n} \sum_{i=1}^n W_i \mathbf{1}(\psi(Z_i) \geq u),$$

so we may write

$$|N_n(u) - \tilde{N}(u)| \leq \frac{1 + 1/\underline{e}}{1 + n} + \left| \frac{1}{n} \sum_{i=1}^n W_i \mathbf{1}(\psi(Z_i) \geq u) - \tilde{N}(u) \right|.$$

Therefore, with probability at least $1 - \delta$,

$$\sup_{u \in \mathbb{R}} |N_n(u) - \tilde{N}(u)| \leq \frac{1 + 1/\underline{e}}{1 + n} + \frac{1}{\underline{e}} \left(4\sqrt{\frac{\log(n+1)}{n}} + \sqrt{\frac{2\log(1/\delta)}{n}} \right). \quad (18)$$

Hoeffding bound for $|A_n - \tilde{A}|$. Recall that

$$A_n := \frac{1 + \sum_{i=1}^n W_i}{1 + n}, \quad \tilde{A} := \mathbb{E}[W].$$

We decompose the deviation:

$$|A_n - \tilde{A}| \leq \frac{1 + 1/\underline{e}}{1 + n} + \left| \frac{1}{n} \sum_{i=1}^n W_i - \mathbb{E}[W] \right|.$$

Applying Hoeffding's inequality for $W_i \in [0, 1/\underline{e}]$, we obtain that, with probability at least $1 - \delta$,

$$|A_n - \tilde{A}| \leq \frac{1 + 1/\underline{e}}{1 + n} + \frac{1}{\underline{e}} \sqrt{\frac{\log(2/\delta)}{2n}}. \quad (19)$$

Combining (17) with (18) and (19), we obtain that, with probability at least $1 - \delta$,

$$\begin{aligned} \sup_{u \in \mathbb{R}} |R_n(u) - \tilde{R}(u)| &\leq \frac{1}{\underline{e}\pi} \left[\frac{2(\underline{e} + 1)}{1 + n} + 4\sqrt{\frac{\log(n+1)}{n}} + \sqrt{\frac{2\log(1/\delta)}{n}} + \sqrt{\frac{\log(2/\delta)}{2n}} \right] \\ &\leq \frac{1}{\underline{e}\pi} \left[\frac{2(\underline{e} + 1)}{n} + 8\sqrt{\frac{\log(2n/\delta)}{n}} \right]. \end{aligned} \quad (20)$$

Step 2 (Bounding the bias due to approximation). We now bound the difference between $\tilde{R}(u)$ and the ideal target $R(u)$, which uses the true propensity $e(X)$. Using the same ratio inequality as before we obtain:

$$|\tilde{R}(u) - R(u)| = \left| \tilde{N}(u) - R(u) \right| + \left| 1 - \tilde{A} \right|. \quad (21)$$

To bound the first term, note that, using the Cauchy-Schwarz inequality,

$$\begin{aligned} \left| \tilde{N}(u) - R(u) \right| &= \left| \mathbb{E} \left[\left(\frac{1}{E} - \frac{1}{e(X)} \right) D \cdot \mathbf{1}(\psi(Z) \geq u) \right] \right| \\ &\leq \mathbb{E} \left[\left| \frac{1}{E} - \frac{1}{e(X)} \right| \right] \leq \sqrt{\mathbb{E} \left[\left(\frac{1}{E} - \frac{1}{e(X)} \right)^2 \right]} =: \Delta_N. \end{aligned}$$

For the second term, using the same argument;

$$|1 - \tilde{A}| = \left| 1 - \mathbb{E} \left[\frac{D}{E} \right] \right| = \left| \mathbb{E} \left[\left(\frac{1}{e(X)} - \frac{1}{E} \right) D \right] \right| \leq \Delta_N.$$

Therefore, we conclude:

$$\sup_{u \in \mathbb{R}} |\tilde{R}(u) - R(u)| \leq 2\Delta_N. \quad (22)$$

Step 3 (Conclusion). Combining the stochastic deviation bound (20) with the approximation bias bound (22), we conclude that with probability at least $1 - \delta$,

$$\begin{aligned} \sup_{u \in \mathbb{R}} |R_n(u) - R(u)| &\leq \sup_u |R_n(u) - \tilde{R}(u)| + \sup_u |\tilde{R}(u) - R(u)| \\ &\leq \frac{1}{\underline{e}\pi} \left[\frac{2(\underline{e} + 1)}{n} + 8\sqrt{\frac{\log(2n/\delta)}{n}} \right] + 2\Delta_N. \end{aligned}$$

■

A.3.4. STABILITY AND VALIDITY OF THE BH PROCEDURE

Lemma 8 (Stability of BH under small perturbations) *Let $p = (p_1, \dots, p_m)$ and $q = (q_1, \dots, q_m)$ be two vectors of p-values in $[0, 1]^m$. Suppose:*

- (Closeness) *For some $\epsilon > 0$ and $\delta \in (0, 1)$, we have*

$$\mathbb{P} \left(\max_{1 \leq j \leq m} |p_j - q_j| \leq \epsilon \right) \geq 1 - \delta.$$

- (Smoothness) *Each q_j has a density f_j supported on $[0, 1]$ satisfying $f_j(t) \leq f_{\max}$ for all $t \in [0, 1]$.*

Let \mathcal{R}_p and \mathcal{R}_q be the rejection sets from applying the BH procedure at level α to p and q , respectively. Then:

$$\mathbb{P}(\mathcal{R}_p \neq \mathcal{R}_q) \leq \delta + 2\epsilon m^2 f_{\max}.$$

Proof [of Lemma 8] Define the event

$$\mathcal{E} := \left\{ \max_{1 \leq j \leq m} |p_j - q_j| \leq \epsilon \right\},$$

which holds with probability at least $1 - \delta$ by assumption.

For each $k \in \{1, \dots, m\}$, the BH procedure compares p-values to candidate thresholds $\tau_k := \frac{\alpha k}{m}$. Define the ‘‘instability region’’:

$$\mathcal{B}_\epsilon := \bigcup_{k=1}^m (\tau_k - \epsilon, \tau_k + \epsilon),$$

which is the union of m intervals of width 2ϵ , and hence has total Lebesgue measure at most $2\epsilon m$.

For any fixed index j , the probability that $q_j \in \mathcal{B}_\epsilon$ is easily bounded using the smoothness assumption: $\mathbb{P}(q_j \in \mathcal{B}_\epsilon) \leq 2\epsilon m f_{\max}$. Applying a union bound over all $j = 1, \dots, m$, we get:

$$\mathbb{P}(\exists j : q_j \in \mathcal{B}_\epsilon) \leq 2\epsilon m^2 f_{\max}.$$

Let A be the “bad event” where either closeness fails or some q_j lies near a threshold:

$$A := \mathcal{E}^c \cup \{\exists j : q_j \in \mathcal{B}_\epsilon\}.$$

Then:

$$\mathbb{P}(A) \leq \delta + 2\epsilon m^2 f_{\max}.$$

On the complement A^c (i.e., when p and q are close and all q_j are away from thresholds), the BH procedure makes the same rejection decisions on p and q because: each possible threshold τ_k is fixed; each p_j and q_j differ by at most ϵ ; no q_j lies within ϵ of any threshold. Hence, for every j , the relation $p_j \leq \tau_k$ matches $q_j \leq \tau_k$ for all k . Therefore, the sorted comparisons and BH cutoffs lead to identical rejection sets: $\mathcal{R}_p = \mathcal{R}_q$ on the event A^c . Thus,

$$\mathbb{P}(\mathcal{R}_p \neq \mathcal{R}_q) \leq \mathbb{P}(A) \leq \delta + 2\epsilon m^2 f_{\max}. \quad \blacksquare$$

Theorem 9 (*Jin and Candès, 2023b, Thm. 3*) *Let $(H_1, \dots, H_m) \in \{0, 1\}^m$ be random variables indicating the status of m hypotheses, where $H_j = 1$ means that the j th hypothesis is a true null. Let $(p_1, \dots, p_m) \in [0, 1]^m$ be random variables, interpreted as p -values for the corresponding hypotheses. Assume:*

- *The p -values p_1, \dots, p_m are mutually independent.*
- *The hypothesis indicators H_1, \dots, H_m are mutually independent.*
- *Each p -value p_j is independent of H_i for all $i \neq j$.*
- *For each $j \in \{1, \dots, m\}$, the p -value p_j satisfies the joint super-uniformity condition:*

$$\mathbb{P}(p_j \leq \alpha, H_j = 1) \leq \alpha \quad \text{for all } \alpha \in [0, 1].$$

Then the Benjamini-Hochberg (BH) procedure at level $\alpha \in (0, 1)$, applied to (p_1, \dots, p_m) , controls the false discovery rate:

$$\text{FDR} := \mathbb{E} \left[\frac{|\mathcal{R} \cap \mathcal{H}_0|}{|\mathcal{R}| \vee 1} \right] \leq \alpha,$$

where:

- $\mathcal{H}_0 := \{j : H_j = 1\}$ *is the (random) set of true null hypotheses,*
- $\mathcal{R} := \{j : p_j \leq \hat{\tau}\}$ *is the BH rejection set,*

- $\hat{\tau} := \max \left\{ \frac{k\alpha}{m} : p_{(k)} \leq \frac{k\alpha}{m} \right\}$ is the BH threshold, with $p_{(1)} \leq \dots \leq p_{(m)}$ the order statistics of (p_1, \dots, p_m) .

Note: This result is a simplified version of the argument in Jin and Candès (2023b), specialized to the case of independent p-values (rather than PRDS).

Proof [of Theorem 9] Let $\mathcal{R} := \{j : p_j \leq \hat{\tau}\}$ be the BH rejection set, and let $R_j := \mathbf{1}\{j \in \mathcal{R}\}$. Then the false discovery rate can be written as:

$$\text{FDR} = \mathbb{E} \left[\frac{|\mathcal{R} \cap \mathcal{H}_0|}{|\mathcal{R}| \vee 1} \right] = \mathbb{E} \left[\sum_{j=1}^m \frac{H_j R_j}{|\mathcal{R}| \vee 1} \right].$$

Using the identity $R_j = \mathbf{1}\{p_j \leq |\mathcal{R}| \alpha/m\}$ and decomposing over possible values of $|\mathcal{R}| = k$, we write:

$$\text{FDR} = \sum_{j=1}^m \sum_{k=1}^m \frac{1}{k} \mathbb{E} \left[\mathbf{1}\{|\mathcal{R}| = k\} \cdot H_j \cdot \mathbf{1}\left\{p_j \leq \frac{k\alpha}{m}\right\} \right].$$

Now fix any index $j \in \{1, \dots, m\}$, and define the modified rejection set $\mathcal{R}(p_j \rightarrow 0)$ to be the BH rejection set obtained by replacing p_j with 0 while keeping all other p-values fixed. Because the BH procedure is monotone in each coordinate, and because reducing p_j does not reduce the size of the rejection set, we have:

$$\mathbf{1}\{|\mathcal{R}| = k\} \cdot \mathbf{1}\left\{p_j \leq \frac{k\alpha}{m}\right\} = \mathbf{1}\{|\mathcal{R}(p_j \rightarrow 0)| = k\} \cdot \mathbf{1}\left\{p_j \leq \frac{k\alpha}{m}\right\}.$$

Let $\mathcal{F}_j := \sigma(p_1, \dots, p_{j-1}, 0, p_{j+1}, \dots, p_m)$ denote the sigma-algebra generated by the other p-values excluding p_j . Then, using independence of H_j and p_j from \mathcal{F}_j , and applying the tower property:

$$\begin{aligned} \text{FDR} &= \sum_{j=1}^m \sum_{k=1}^m \frac{1}{k} \mathbb{E} \left[\mathbf{1}\{|\mathcal{R}(p_j \rightarrow 0)| = k\} \cdot H_j \cdot \mathbf{1}\left\{p_j \leq \frac{k\alpha}{m}\right\} \right] \\ &= \sum_{j=1}^m \sum_{k=1}^m \frac{1}{k} \mathbb{E} \left[\mathbf{1}\{|\mathcal{R}(p_j \rightarrow 0)| = k\} \cdot \mathbb{E} \left[H_j \cdot \mathbf{1}\left\{p_j \leq \frac{k\alpha}{m}\right\} \mid \mathcal{F}_j \right] \right] \\ &= \sum_{j=1}^m \sum_{k=1}^m \frac{1}{k} \mathbb{E} \left[\mathbf{1}\{|\mathcal{R}(p_j \rightarrow 0)| = k\} \cdot \mathbb{P} \left(p_j \leq \frac{k\alpha}{m}, H_j = 1 \right) \right] \\ &\leq \sum_{j=1}^m \sum_{k=1}^m \frac{1}{k} \cdot \mathbb{E} \left[\mathbf{1}\{|\mathcal{R}(p_j \rightarrow 0)| = k\} \cdot \frac{k\alpha}{m} \right] \\ &= \frac{\alpha}{m} \sum_{j=1}^m \sum_{k=1}^m \mathbb{P} (|\mathcal{R}(p_j \rightarrow 0)| = k) \leq \alpha. \end{aligned}$$

■

Appendix B. Additional Results from Experiments with Synthetic Data

B.1. Additional Details on Data Generating Distributions

Table 3: Summary of the four synthetic data generation settings considered for numerical experiments. Settings 1, 2, and 3 are adapted from [Sesia and Svetnik \(2025\)](#), with Setting 3 originally introduced in [Candès et al. \(2023\)](#). Setting 4 is new. In this table, X denotes a p -dimensional vector of covariates for a random individual, and X_j indicates its j -th element for any $j \in [p]$. Not all covariates are used to generate Y ; most serve as noisy variables, making it more challenging to fit accurate survival and censoring models. In some experiments, we simplify the training process by excluding a subset of these noisy variables.

Setting	p	Covariate, Survival, and Censoring Distributions
1	100	X : Unif($[0, 1]^p$) $T \mid X$: LogNormal, $\mu_s(X) = (X_2 > \frac{1}{2}) + (X_3 < \frac{1}{2}) + (1 - X_1)^{0.25}$, $\sigma_s(X) = \frac{1-X_1}{10}$ $C \mid X$: LogNormal, $\mu_c(X) = (X_2 > \frac{1}{2}) + (X_3 < \frac{1}{2}) + (1 - X_1)^4 + \frac{4}{10}$, $\sigma_c(X) = \frac{X_2}{10}$
2	100	X : Unif($[0, 1]^p$) $T \mid X$: LogNormal, $\mu_s(X) = X_1^{0.25}$, $\sigma_s(X) = 0.1$ $C \mid X$: LogNormal, $\mu_c(X) = X_1^4 + \frac{4}{10}$, $\sigma_c(X) = 0.1$
3	100	X : Unif($[-1, 1]^p$) $T \mid X$: LogNormal, $\mu_s(X) = \log(2) + 1 + 0.55(X_1^2 - X_3X_5)$, $\sigma_s(X) = X_{10} + 1$ $C \mid X$: Exponential, $\lambda_c(X) = 0.4$
4	100	X : Unif($[-1, 1]^p$) $T \mid X$: LogNormal, $\mu_s(X) = 0.2(1 + X_1)X_2 + \log(2)\mathbf{1}(X_3 > 0) + \log(10)\mathbf{1}(X_3 \leq 0)$, $\sigma_s(X) = 0.25$ $C \mid X$: LogNormal, $\mu_c(X) = 2 + 0.5X_1$, $\sigma_c(X) = 0.1$ Shifted $T \mid X$: LogNormal, $\mu_s(X) = 0.2(1 + X_1)X_2 + \log(10)$, $\sigma_s(X) = 0.25$

B.2. Effect of the Training Sample Size

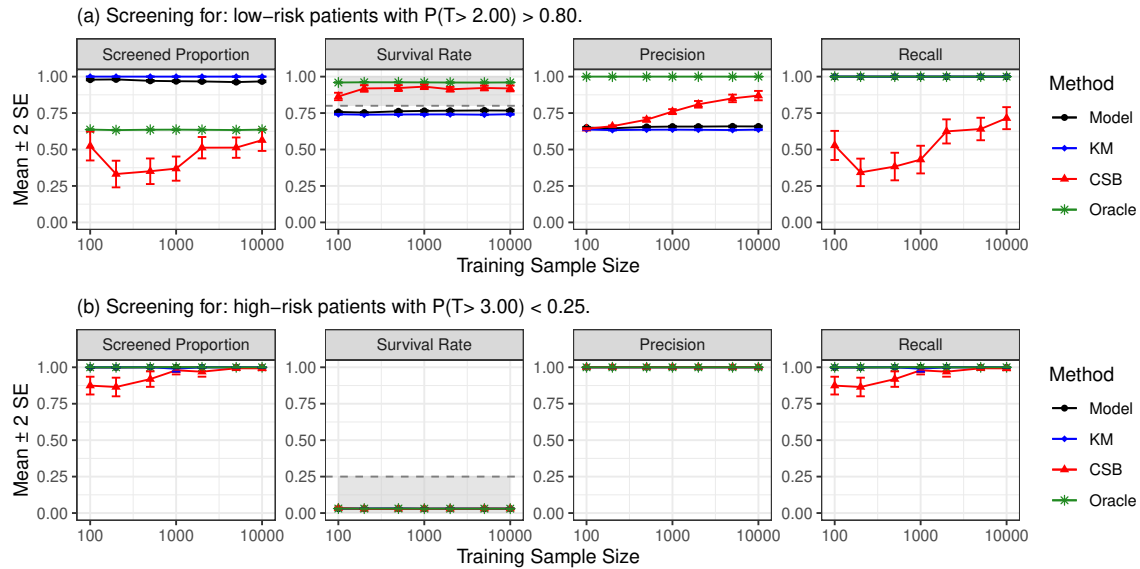


Figure 5: Patient screening performance as a function of training sample size in a moderately challenging synthetic data scenario (Setting 2 from Table 3). Other details are as in Figure 2.

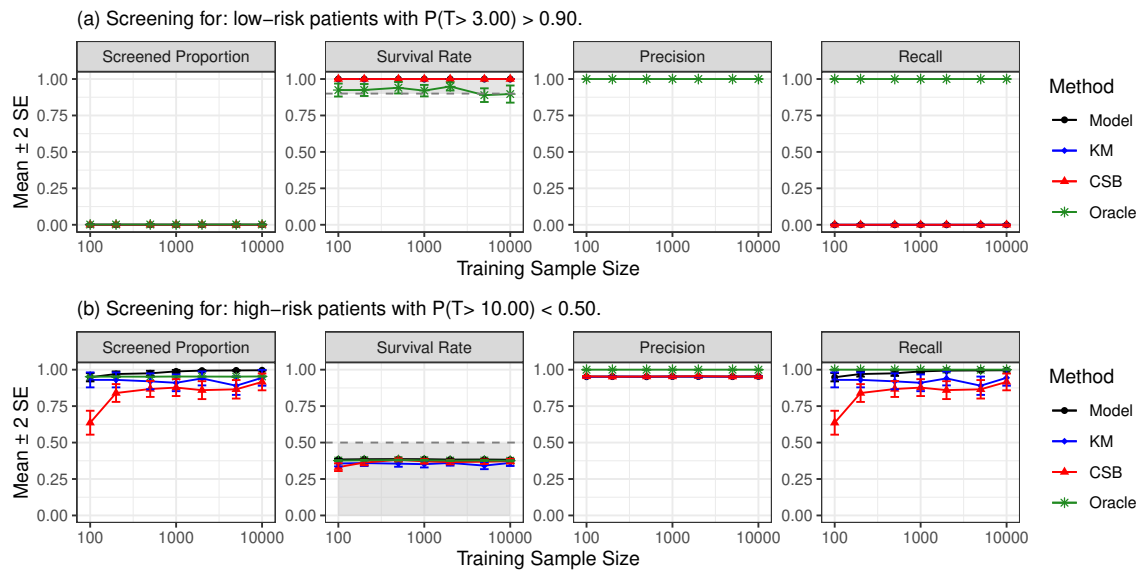


Figure 6: Patient screening performance as a function of training sample size in a relatively easy synthetic data scenario (Setting 3 from Table 3). Other details are as in Figure 2.

B.3. Effect of the Calibration Sample Size

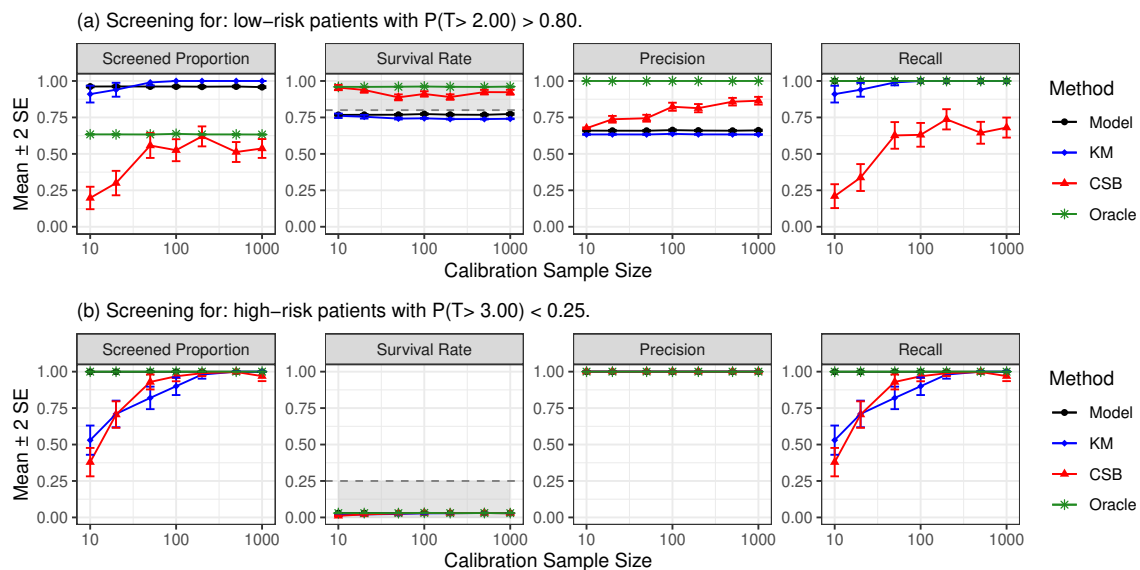


Figure 7: Effect of calibration sample size on patient screening performance with conformal survival bands (CSB) in a moderately challenging synthetic data scenario (Setting 2 from Table 3). Other details are as in Figure 3.

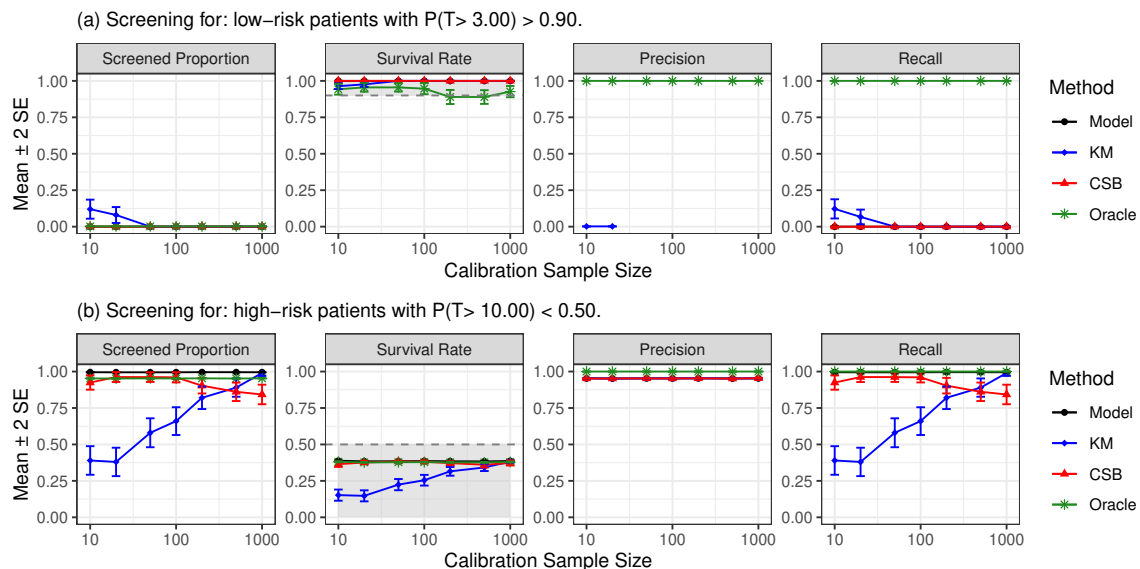


Figure 8: Effect of calibration sample size on patient screening performance with conformal survival bands (CSB) in a relatively easy synthetic data scenario (Setting 3 from Table 3). Other details are as in Figure 3.

B.4. Effect of the Training Sample Size for the Censoring Model

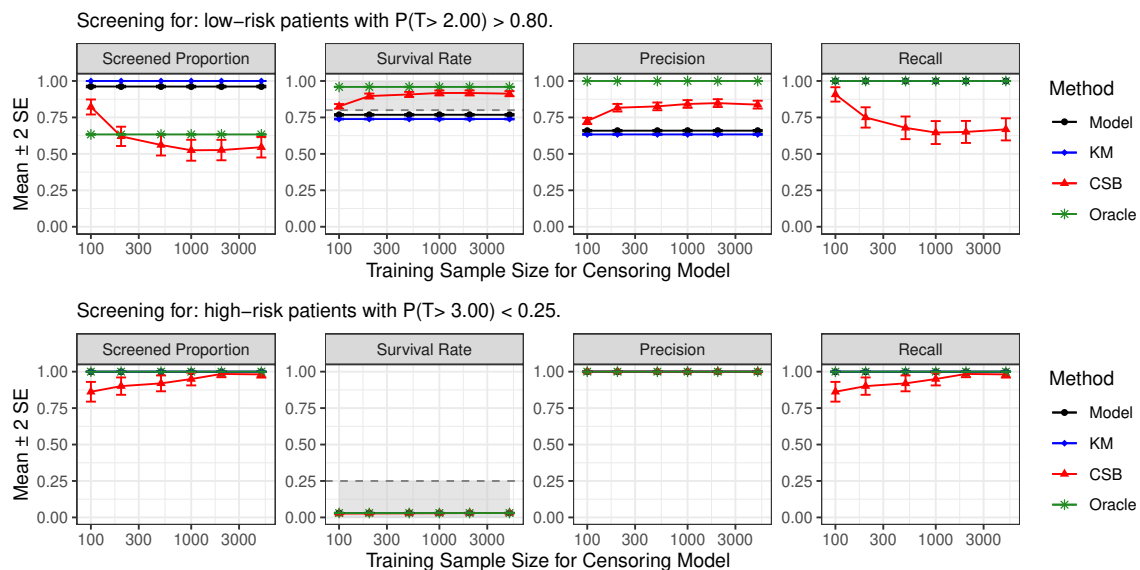


Figure 9: Effect of the training sample size used for fitting the censoring model on patient screening performance with conformal survival bands (CSB) in a moderately challenging synthetic data scenario (Setting 2 from Table 3). Other details are as in Figure 4.

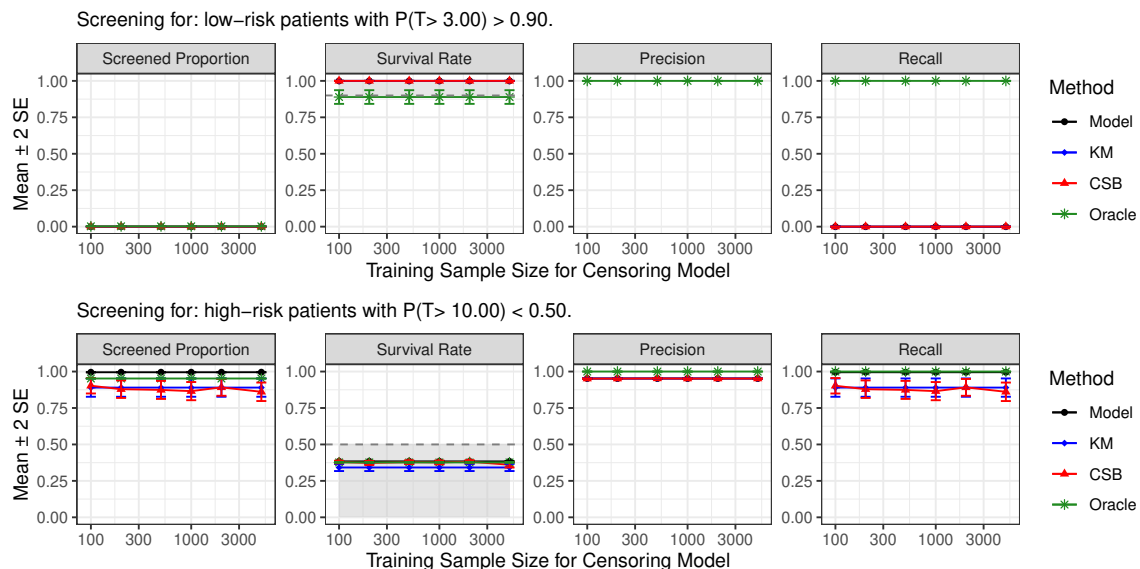


Figure 10: Effect of the training sample size used for fitting the censoring model on patient screening performance with conformal survival bands (CSB) in a relatively easy synthetic data scenario (Setting 3 from Table 3). Other details are as in Figure 4.

B.5. Effect of the Number of Features for the Censoring Model

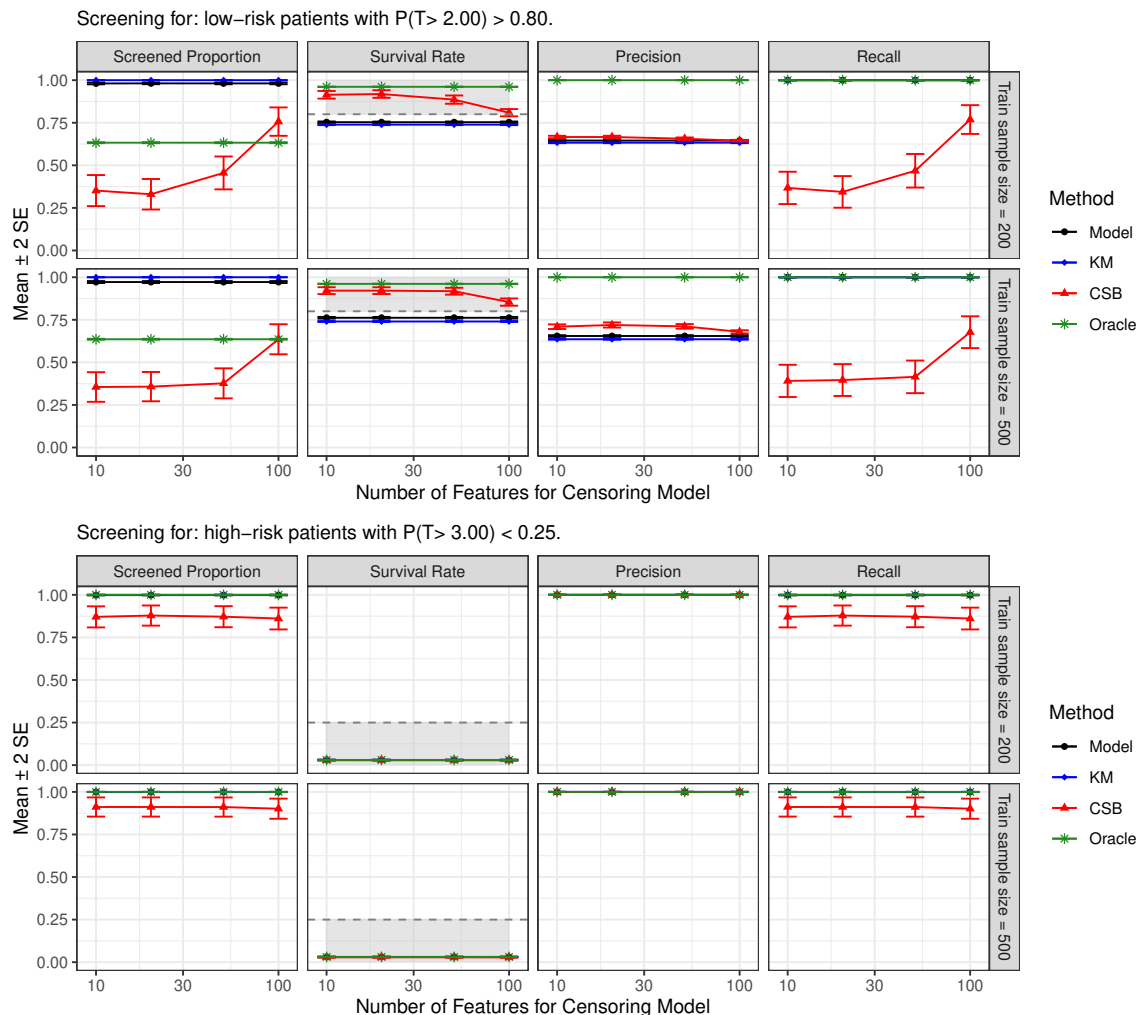


Figure 11: Impact of the number of features used in fitting the censoring model on patient screening performance with conformal survival bands (CSB) in a challenging synthetic data scenario (Setting 1 from Table 3). The total number of features is 100, but only a subset—including the relevant ones—is used to fit the censoring model. Top: Low-risk screening. Screening performance is more sensitive to the number of features when the training sample size is small (200); using too many features degrades performance due to the difficulty of accurately estimating the censoring model. When the sample size is larger (500), performance is less sensitive to the number of features. Bottom: High-risk screening (training sample sizes = 200 and 500). In this case, screening performance shows even less sensitivity to the number of features used to fit the censoring model.

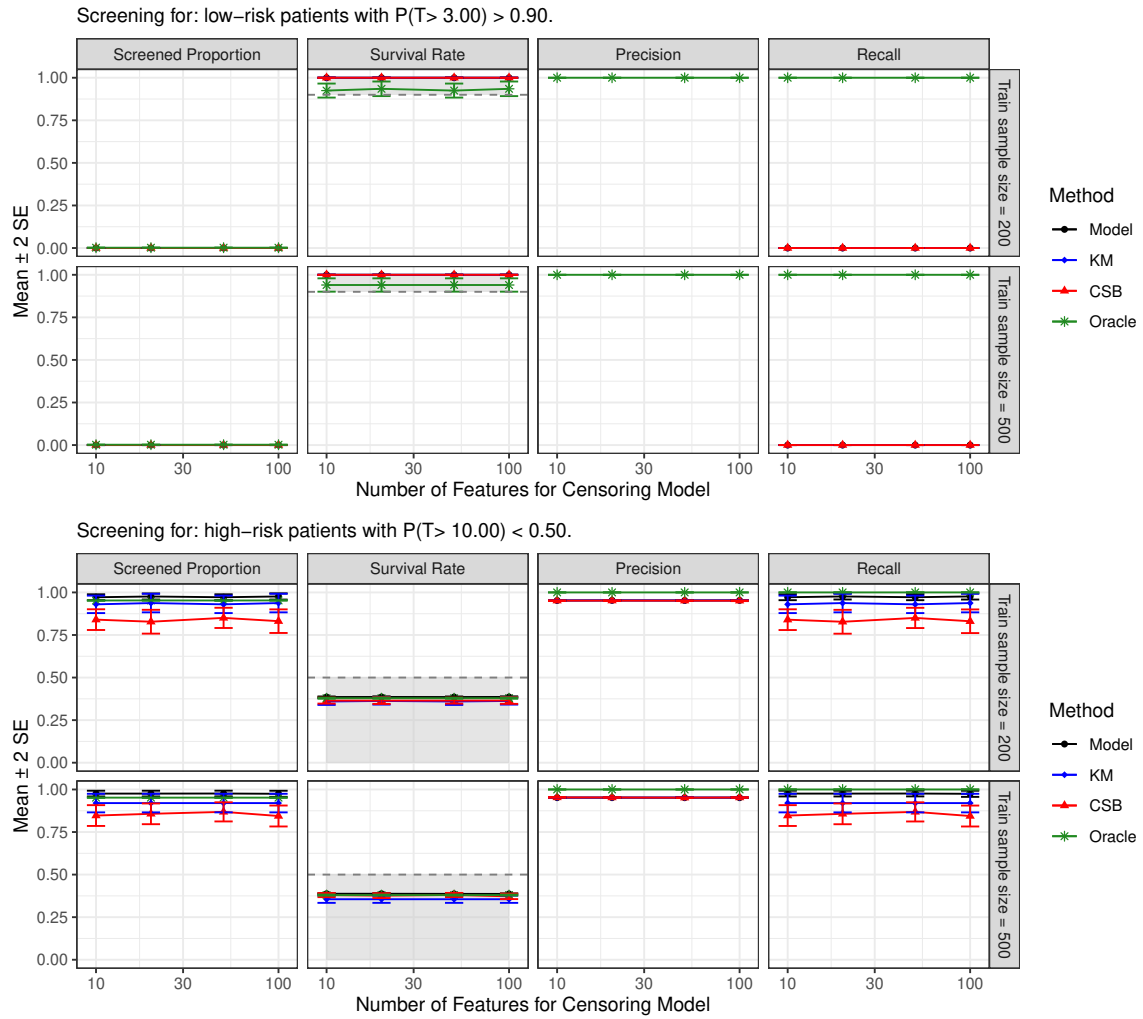


Figure 12: Impact of the number of features used in fitting the censoring model on patient screening performance with conformal survival bands (CSB) in an easier synthetic data scenario (Setting 3 from Table 3). Other details are as in Figure 11.

B.6. Effect of Different Censoring and Survival Models

Table 4: Performance of different methods for screening low-risk patients with $P(T > 6) > 0.80$ in a challenging synthetic scenario (Setting 1 from Table 3), using various censoring and survival models. The training sample size is 1000. Other details are as in Figure 2.

Survival Model	Method	Screened	Survival	Precision	Recall
Censoring Model: grf					
grf	Model	0.946 ± 0.008	0.570 ± 0.006	0.541 ± 0.006	0.999 ± 0.001
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.221 ± 0.071	0.902 ± 0.033	0.685 ± 0.034	0.266 ± 0.079
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.995 ± 0.001	0.541 ± 0.003	0.514 ± 0.003	1.000 ± 0.000
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.049 ± 0.043	0.977 ± 0.020	0.506 ± 0.003	0.050 ± 0.044
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.814 ± 0.007	0.659 ± 0.006	0.626 ± 0.006	0.994 ± 0.001
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.308 ± 0.061	0.857 ± 0.028	0.722 ± 0.019	0.413 ± 0.078
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.764 ± 0.005	0.652 ± 0.004	0.619 ± 0.004	0.925 ± 0.004
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.303 ± 0.047	0.834 ± 0.022	0.794 ± 0.022	0.433 ± 0.057
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Censoring Model: Cox					
grf	Model	0.945 ± 0.008	0.571 ± 0.007	0.542 ± 0.007	0.999 ± 0.001
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.688 ± 0.075	0.692 ± 0.036	0.619 ± 0.030	0.763 ± 0.071
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.995 ± 0.001	0.541 ± 0.003	0.514 ± 0.003	1.000 ± 0.000
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.583 ± 0.098	0.733 ± 0.045	0.519 ± 0.003	0.590 ± 0.099
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.814 ± 0.007	0.659 ± 0.006	0.626 ± 0.006	0.994 ± 0.001
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.700 ± 0.045	0.692 ± 0.018	0.648 ± 0.014	0.864 ± 0.047
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.764 ± 0.005	0.652 ± 0.004	0.619 ± 0.004	0.925 ± 0.004
	KM	1.000 ± 0.000	0.539 ± 0.003	0.511 ± 0.003	1.000 ± 0.000
	CSB	0.661 ± 0.036	0.683 ± 0.013	0.654 ± 0.014	0.822 ± 0.035
	Oracle	0.511 ± 0.003	0.974 ± 0.001	1.000 ± 0.000	1.000 ± 0.000

Table 5: Performance of different methods for screening high-risk patients with $P(T > 12) < 0.80$ in a challenging synthetic data scenario (Setting 1 from Table 3), using various censoring and survival models. The training sample size is fixed at 1000. Other details are as in Figure 2.

Survival Model	Method	Screened	Survival	Precision	Recall
Censoring Model: grf					
grf	Model	0.755 ± 0.005	0.035 ± 0.002	0.965 ± 0.002	0.955 ± 0.006
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.743 ± 0.022	0.040 ± 0.002	0.959 ± 0.002	0.934 ± 0.028
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.109 ± 0.004	0.000 ± 0.000	1.000 ± 0.000	0.142 ± 0.006
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.137 ± 0.006	0.000 ± 0.000	1.000 ± 0.000	0.179 ± 0.008
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.819 ± 0.004	0.080 ± 0.003	0.921 ± 0.003	0.988 ± 0.001
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.817 ± 0.024	0.094 ± 0.005	0.904 ± 0.005	0.969 ± 0.028
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.713 ± 0.005	0.071 ± 0.003	0.929 ± 0.003	0.868 ± 0.005
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.680 ± 0.015	0.073 ± 0.003	0.928 ± 0.003	0.826 ± 0.018
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Censoring Model: Cox					
grf	Model	0.755 ± 0.005	0.035 ± 0.002	0.965 ± 0.002	0.955 ± 0.006
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.718 ± 0.034	0.039 ± 0.003	0.959 ± 0.002	0.903 ± 0.043
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.109 ± 0.004	0.000 ± 0.000	1.000 ± 0.000	0.142 ± 0.006
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.134 ± 0.007	0.000 ± 0.000	1.000 ± 0.000	0.175 ± 0.009
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.819 ± 0.004	0.080 ± 0.003	0.921 ± 0.003	0.988 ± 0.001
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.790 ± 0.036	0.089 ± 0.007	0.909 ± 0.006	0.938 ± 0.042
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.713 ± 0.005	0.071 ± 0.003	0.929 ± 0.003	0.868 ± 0.005
	KM	1.000 ± 0.000	0.238 ± 0.003	0.763 ± 0.003	1.000 ± 0.000
	CSB	0.660 ± 0.026	0.070 ± 0.004	0.929 ± 0.004	0.802 ± 0.031
	Oracle	0.763 ± 0.003	0.001 ± 0.000	1.000 ± 0.000	1.000 ± 0.000

Table 6: Performance of different methods for screening low-risk patients with $P(T > 2) > 0.80$ in a relatively challenging data scenario (Setting 2 from Table 3), using various censoring and survival models. The training sample size is fixed at 1000. Other details are as in Figure 2.

Survival Model	Method	Screened	Survival	Precision	Recall
Censoring Model: grf					
grf	Model	0.968 ± 0.005	0.765 ± 0.005	0.658 ± 0.004	1.000 ± 0.000
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.362 ± 0.082	0.932 ± 0.017	0.762 ± 0.018	0.422 ± 0.094
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.951 ± 0.003	0.776 ± 0.004	0.669 ± 0.003	1.000 ± 0.000
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.055 ± 0.044	0.987 ± 0.010	0.682 ± 0.004	0.060 ± 0.048
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.880 ± 0.005	0.834 ± 0.004	0.723 ± 0.005	1.000 ± 0.000
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.556 ± 0.066	0.908 ± 0.011	0.822 ± 0.017	0.693 ± 0.077
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.596 ± 0.011	0.961 ± 0.003	0.967 ± 0.006	0.904 ± 0.011
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.496 ± 0.030	0.969 ± 0.006	0.969 ± 0.010	0.747 ± 0.039
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Censoring Model: Cox					
grf	Model	0.969 ± 0.006	0.765 ± 0.005	0.657 ± 0.004	1.000 ± 0.000
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.599 ± 0.072	0.889 ± 0.015	0.766 ± 0.015	0.704 ± 0.081
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.951 ± 0.003	0.776 ± 0.004	0.669 ± 0.003	1.000 ± 0.000
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.880 ± 0.005	0.834 ± 0.004	0.723 ± 0.005	1.000 ± 0.000
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.776 ± 0.037	0.870 ± 0.009	0.773 ± 0.013	0.929 ± 0.040
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.596 ± 0.011	0.961 ± 0.003	0.967 ± 0.006	0.904 ± 0.011
	KM	1.000 ± 0.000	0.741 ± 0.003	0.636 ± 0.002	1.000 ± 0.000
	CSB	0.569 ± 0.021	0.959 ± 0.005	0.955 ± 0.011	0.846 ± 0.023
	Oracle	0.636 ± 0.002	0.960 ± 0.001	1.000 ± 0.000	1.000 ± 0.000

Table 7: Performance of different methods for screening high-risk patients with $P(T > 3) < 0.25$ in a relatively challenging data scenario (Setting 2 from Table 3), using various censoring and survival models. The training sample size is fixed at 1000. Other details are as in Figure 2.

Survival Model	Method	Screened	Survival	Precision	Recall
Censoring Model: grf					
grf	Model	0.999 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	0.999 ± 0.000
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.979 ± 0.028	0.029 ± 0.001	1.000 ± 0.000	0.979 ± 0.028
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.726 ± 0.006	0.008 ± 0.001	1.000 ± 0.000	0.726 ± 0.006
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.738 ± 0.022	0.010 ± 0.001	1.000 ± 0.000	0.738 ± 0.022
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.907 ± 0.014	0.033 ± 0.001	1.000 ± 0.000	0.907 ± 0.014
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.898 ± 0.030	0.031 ± 0.002	1.000 ± 0.000	0.898 ± 0.030
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.964 ± 0.002	0.026 ± 0.001	1.000 ± 0.000	0.964 ± 0.002
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.951 ± 0.019	0.025 ± 0.001	1.000 ± 0.000	0.951 ± 0.019
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Censoring Model: Cox					
grf	Model	0.999 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	0.999 ± 0.000
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.999 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	0.999 ± 0.000
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.726 ± 0.006	0.008 ± 0.001	1.000 ± 0.000	0.726 ± 0.006
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.753 ± 0.006	0.010 ± 0.001	1.000 ± 0.000	0.753 ± 0.006
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.907 ± 0.014	0.033 ± 0.001	1.000 ± 0.000	0.907 ± 0.014
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.918 ± 0.014	0.032 ± 0.001	1.000 ± 0.000	0.918 ± 0.014
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.964 ± 0.002	0.026 ± 0.001	1.000 ± 0.000	0.964 ± 0.002
	KM	0.990 ± 0.020	0.030 ± 0.001	1.000 ± 0.000	0.990 ± 0.020
	CSB	0.962 ± 0.002	0.026 ± 0.001	1.000 ± 0.000	0.962 ± 0.002
	Oracle	1.000 ± 0.000	0.030 ± 0.001	1.000 ± 0.000	1.000 ± 0.000

Table 8: Performance of different methods for screening low-risk patients with $P(T > 3) > 0.90$ in an easier synthetic data scenario (Setting 3 from Table 3), using various censoring and survival models. The training sample size is fixed at 1000. Other details are as in Figure 2.

Survival Model	Method	Screened	Survival	Precision	Recall
Censoring Model: grf					
grf	Model	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	1.000 ± 0.000	0.701 ± 0.002	0.002 ± 0.000	1.000 ± 0.000
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.046 ± 0.004	0.702 ± 0.016	0.003 ± 0.003	0.060 ± 0.038
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000
Censoring Model: Cox					
grf	Model	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	1.000 ± 0.000	0.701 ± 0.002	0.002 ± 0.000	1.000 ± 0.000
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.046 ± 0.004	0.702 ± 0.016	0.003 ± 0.003	0.060 ± 0.038
	KM	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
	Oracle	0.002 ± 0.000	0.920 ± 0.039	1.000 ± 0.000	1.000 ± 0.000

Table 9: Performance of different methods for screening high-risk patients with $P(T > 10) < 0.50$ in an easier synthetic data scenario (Setting 3 from Table 3), using various censoring and survival models. The training sample size is fixed at 1000. Other details are as in Figure 2.

Survival Model	Method	Screened	Survival	Precision	Recall
Censoring Model: grf					
grf	Model	0.988 ± 0.010	0.386 ± 0.003	0.953 ± 0.001	0.987 ± 0.010
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.910 ± 0.051	0.370 ± 0.016	0.953 ± 0.001	0.910 ± 0.051
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.000 ± 0.000	0.000 ± 0.000	0.500 ± 0.141	0.000 ± 0.000
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.000 ± 0.000	0.029 ± 0.031	0.972 ± 0.019	0.000 ± 0.000
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.791 ± 0.025	0.386 ± 0.003	0.953 ± 0.001	0.791 ± 0.025
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.771 ± 0.039	0.390 ± 0.005	0.951 ± 0.003	0.771 ± 0.039
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.688 ± 0.017	0.386 ± 0.003	0.953 ± 0.002	0.688 ± 0.017
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.605 ± 0.034	0.388 ± 0.008	0.952 ± 0.004	0.605 ± 0.034
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000
Censoring Model: Cox					
grf	Model	0.988 ± 0.010	0.386 ± 0.003	0.953 ± 0.001	0.987 ± 0.010
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.916 ± 0.049	0.370 ± 0.016	0.953 ± 0.001	0.916 ± 0.049
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000
survreg	Model	0.000 ± 0.000	0.000 ± 0.000	0.500 ± 0.141	0.000 ± 0.000
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.000 ± 0.000	0.029 ± 0.031	0.972 ± 0.019	0.000 ± 0.000
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000
rf	Model	0.791 ± 0.025	0.386 ± 0.003	0.953 ± 0.001	0.791 ± 0.025
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.762 ± 0.043	0.385 ± 0.009	0.953 ± 0.002	0.761 ± 0.042
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000
Cox	Model	0.688 ± 0.017	0.386 ± 0.003	0.953 ± 0.002	0.688 ± 0.017
	KM	0.910 ± 0.058	0.352 ± 0.022	0.953 ± 0.001	0.910 ± 0.058
	CSB	0.625 ± 0.029	0.388 ± 0.008	0.952 ± 0.005	0.625 ± 0.029
	Oracle	0.953 ± 0.001	0.378 ± 0.003	1.000 ± 0.000	1.000 ± 0.000

B.7. Effect of Distribution Shift

Table 10: Performance of different methods for screening high-risk patients with $P(T > 3) < 0.50$ under distribution shift (Setting 4 from Table 3). These results correspond to the same experiments reported in Table 1, but evaluate high-risk screening instead of low-risk. In this case, a misspecified *grf* model primarily reduces the power of the screening rule by selecting fewer patients than desired.

Method	Screened	Survival	Precision	Recall
Low-Quality Model				
Model	0.000 ± 0.000	NA	NA	0.000 ± 0.000
KM	0.030 ± 0.034	0.016 ± 0.019	0.495 ± 0.002	0.030 ± 0.034
CSB	0.000 ± 0.000	NA	NA	0.000 ± 0.000
Oracle	0.501 ± 0.003	0.077 ± 0.002	1.000 ± 0.000	1.000 ± 0.000
High-Quality Model				
Model	0.503 ± 0.004	0.080 ± 0.003	0.998 ± 0.000	1.000 ± 0.000
KM	0.026 ± 0.036	0.014 ± 0.020	0.502 ± 0.001	0.026 ± 0.036
CSB	0.503 ± 0.004	0.080 ± 0.003	0.998 ± 0.000	0.999 ± 0.002
Oracle	0.503 ± 0.004	0.078 ± 0.003	1.000 ± 0.000	1.000 ± 0.000

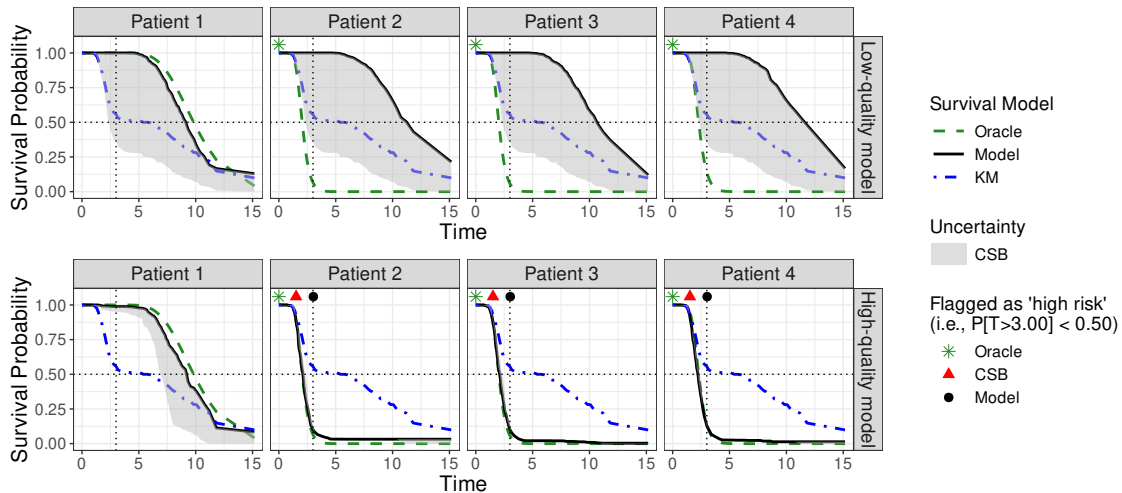


Figure 13: Illustration of the use of conformal survival bands (shaded regions) for screening test patients in a simulated censored dataset under distribution shift, corresponding to the high-risk screening experiments discussed in Table 10. Solid black curves show survival estimates from either an inaccurate (top) or accurate (bottom) survival forest model, while dashed green curves represent the true survival probabilities. The goal is to identify high-risk patients—those with less than 50% probability (horizontal dotted line) of surviving beyond 3 years (vertical dotted line). Other details are as in Figure 1.

Appendix C. Additional Details on Real Data Applications

We apply our method to seven datasets previously utilized by [Sesia and Svetnik \(2025\)](#): the Colon Cancer Chemotherapy (COLON) dataset; the German Breast Cancer Study Group (GBSG) dataset; the Stanford Heart Transplant Study (HEART); the Molecular Taxonomy of Breast Cancer International Consortium (METABRIC) dataset; the Primary Biliary Cirrhosis (PBC) dataset; the Diabetic Retinopathy Study (RETINOPATHY); and the Veterans’ Administration Lung Cancer Trial (VALCT). Table 11 provides details on the number of observations, covariates, and data sources.

The datasets were obtained from various publicly available sources. COLON, HEART, PBC, RETINOPATHY, and VALCT are included in the `survival` R package. GBSG was sourced from GitHub: <https://github.com/jaredleekatzman/DeepSurv/>. METABRIC was accessed via https://www.cbioportal.org/study/summary?id=brca_metabrig.

Each dataset underwent a pre-processing pipeline to ensure consistency and prepare the data for analysis, as in [Sesia and Svetnik \(2025\)](#). Survival times equal to zero were replaced with half the smallest observed non-zero time. Missing values were imputed using the median for numeric variables and the mode for categorical variables. Factor variables were processed to merge rare levels (frequency below 2%) into an “Other” category, while binary factors with one rare level were removed entirely. Dummy variables were created for all factors, and redundant features were identified and removed using an alias check. Additionally, highly correlated features (correlation above 0.75) were iteratively filtered.

Table 11: Summary of the publicly available survival analysis datasets used in Section 4.

Dataset	Obs.	Vars.	Source	Citation
COLON	1858	11	survival	Moertel et al. (1990)
GBSG	2232	6	github.com	Katzman et al. (2018)
HEART	172	4	survival	Crowley and Hu (1977)
METABRIC	1981	41	cbioportal.org	Curtis et al. (2012)
PBC	418	17	survival	Therneau et al. (2000)
RETINOPATHY	394	5	survival	Blair et al. (1980)
VALCT	137	6	survival	Kalbfleisch and Prentice (2002)

Table 12: Detailed screening results for low-risk selection using the *grf* model, with threshold rule $P(T > t_1) > 0.80$ and t_1 set to the 0.1 quantile of observed times in each dataset. Shown are the screened proportion and bounds on the survival rate among selected patients, aggregated over 100 repetitions.

Method	Screened	Survival (lower bound)	Survival (upper bound)
COLON			
Model	0.921 ± 0.006	0.917 ± 0.004	0.920 ± 0.004
KM	1.000 ± 0.000	0.902 ± 0.004	0.905 ± 0.004
CSB	0.647 ± 0.121	0.939 ± 0.012	0.941 ± 0.012
GBSG			
Model	0.897 ± 0.006	0.921 ± 0.003	0.931 ± 0.003
KM	1.000 ± 0.000	0.901 ± 0.003	0.912 ± 0.003
CSB	0.905 ± 0.007	0.921 ± 0.003	0.930 ± 0.003
HEART			
Model	1.000 ± 0.000	0.930 ± 0.011	0.965 ± 0.009
KM	1.000 ± 0.000	0.930 ± 0.011	0.965 ± 0.009
CSB	0.919 ± 0.077	0.934 ± 0.012	0.965 ± 0.009
METABRIC			
Model	0.964 ± 0.004	0.911 ± 0.003	0.927 ± 0.004
KM	1.000 ± 0.000	0.905 ± 0.003	0.921 ± 0.003
CSB	0.970 ± 0.005	0.910 ± 0.004	0.926 ± 0.004
PBC			
Model	0.857 ± 0.012	0.961 ± 0.007	0.961 ± 0.007
KM	1.000 ± 0.000	0.904 ± 0.010	0.904 ± 0.010
CSB	0.809 ± 0.061	0.956 ± 0.009	0.956 ± 0.009
RETINOPATHY			
Model	0.982 ± 0.010	0.884 ± 0.008	0.903 ± 0.008
KM	1.000 ± 0.000	0.884 ± 0.008	0.904 ± 0.008
CSB	0.515 ± 0.137	0.934 ± 0.019	0.945 ± 0.016
VALCT			
Model	0.851 ± 0.028	0.914 ± 0.018	0.914 ± 0.018
KM	0.980 ± 0.040	0.905 ± 0.015	0.905 ± 0.015
CSB	0.805 ± 0.068	0.928 ± 0.019	0.928 ± 0.019

Table 13: Detailed screening results for high-risk selection using the *Cox* model, with threshold rule $P(T > t_1) < 0.80$ and t_1 set to the 0.1 quantile of observed times in each dataset. Shown are the screened proportion and bounds on the survival rate among selected patients, aggregated over 100 repetitions. In some cases, no patients are selected.

Method	Screened	Survival (lower bound)	Survival (upper bound)
COLON			
Model	0.069 ± 0.005	0.747 ± 0.030	0.747 ± 0.030
KM	0.000 ± 0.000	NA	NA
CSB	0.014 ± 0.008	0.166 ± 0.090	0.166 ± 0.090
GBSG			
Model	0.028 ± 0.002	0.849 ± 0.027	0.861 ± 0.024
KM	0.000 ± 0.000	NA	NA
CSB	0.000 ± 0.000	NA	NA
HEART			
Model	0.000 ± 0.000	NA	NA
KM	0.000 ± 0.000	NA	NA
CSB	0.000 ± 0.000	NA	NA
METABRIC			
Model	0.036 ± 0.003	0.749 ± 0.031	0.752 ± 0.031
KM	0.000 ± 0.000	NA	NA
CSB	0.000 ± 0.000	NA	NA
PBC			
Model	0.125 ± 0.013	0.597 ± 0.054	0.597 ± 0.054
KM	0.000 ± 0.000	NA	NA
CSB	0.023 ± 0.015	0.140 ± 0.085	0.140 ± 0.085
RETINOPATHY			
Model	0.012 ± 0.006	0.256 ± 0.115	0.266 ± 0.119
KM	0.000 ± 0.000	NA	NA
CSB	0.002 ± 0.002	0.075 ± 0.073	0.075 ± 0.073
VALCT			
Model	0.139 ± 0.018	0.698 ± 0.077	0.698 ± 0.077
KM	0.020 ± 0.040	0.018 ± 0.036	0.018 ± 0.036
CSB	0.001 ± 0.003	0.010 ± 0.020	0.010 ± 0.020