

**Tipo de documento:** preprint

# Investigating symptom duration using current status data: a case study of post-acute COVID-19 syndrome

**Autoría ditelliana:** *Rotnitzky, Andrea (Universidad Torcuato Di Tella. Departamento de Economía)*

**Otras autorías:** *Wolock, Charles J.; Jacob, Susan; Bennett, Julia C.; Elias-Warren, Anna; O'Hanlon, Jessica; Kenny, Avi; Jewell, Nicholas P.; Weil, Ana A.; Chu, Helen Y.; Carone, Marco*

**Año de publicación:** 2024

## ¿Cómo citar este trabajo?

Wolock, C., et al. "Investigating symptom duration using current status data: a case study of post-acute COVID-19 syndrome". [Preprint. Universidad Torcuato Di Tella]. Repositorio Digital Universidad Torcuato Di Tella  
<https://repositorio.utdt.edu/handle/20.500.13098/12888>

El presente documento se encuentra alojado en el Repositorio Digital de la Universidad Torcuato Di Tella, bajo una licencia internacional Creative Commons Atribución No Comercial CompartirIgual 4.0.

# Investigating symptom duration using current status data: a case study of post-acute COVID-19 syndrome

Charles J. Wolock<sup>\*1</sup>, Susan Jacob<sup>\*2</sup>, Julia C. Bennett<sup>2</sup>, Anna Elias-Warren<sup>2</sup>, Jessica O'Hanlon<sup>2</sup>,  
Avi Kenny<sup>3,4</sup>, Nicholas P. Jewell<sup>5</sup>, Andrea Rotnitzky<sup>8</sup>, Ana A. Weil<sup>6,7</sup>, Helen Y. Chu<sup>2</sup> & Marco  
Carone<sup>8</sup>

<sup>1</sup>Department of Biostatistics, Epidemiology and Informatics, University of Pennsylvania

<sup>2</sup>Department of Epidemiology, University of Washington

<sup>3</sup>Department of Biostatistics and Bioinformatics, Duke University

<sup>4</sup>Global Health Institute, Duke University

<sup>5</sup>Department of Medical Statistics, London School of Hygiene and Tropical Medicine

<sup>6</sup>Department of Medicine, University of Washington

<sup>7</sup>Department of Global Health, University of Washington

<sup>8</sup>Department of Biostatistics, University of Washington

## Abstract

For infectious diseases, characterizing symptom duration is of clinical and public health importance. Symptom duration may be assessed by surveying infected individuals and querying symptom status at the time of survey response. For example, in a SARS-CoV-2 testing program at the University of Washington, participants were surveyed at least 28 days after testing positive and asked to report current symptom status. This study design yielded current status data: Outcome measurements for each respondent consisted only of the time of survey response and a binary indicator of whether symptoms had resolved by that time. Such study design benefits from limited risk of recall bias, but analyzing the resulting data necessitates specialized statistical tools. Here, we review methods for current status data and describe a novel application of modern nonparametric techniques to this setting. The proposed approach is valid under weaker assumptions compared to existing methods, allows use of flexible machine learning tools, and handles potential survey nonresponse. From the university study, we estimate that 19% of participants experienced ongoing symptoms 30 days after testing positive, decreasing to 7% at 90 days. Female sex, history of seasonal allergies, fatigue during acute infection, and higher viral load were associated with slower symptom resolution.

*Keywords:* Interval censoring, long COVID, machine learning, nonparametric, survival analysis

---

<sup>\*</sup>The first two authors contributed equally to this work.

## Introduction

The persistence of COVID-19 symptoms is a feature of the natural history of SARS-CoV-2 infection. Early in the COVID-19 pandemic, it became clear that many infected individuals suffered symptoms continuing beyond the acute phase of the disease [1–3]. Post-acute sequelae are not unique to SARS-CoV-2; many viruses, including the Ebola, Epstein-Barr, chikungunya, dengue, polio, and SARS viruses, among others, have been linked to post-acute infection syndromes [4]. There is now a substantial body of literature on the post-acute sequelae of COVID-19 (PASC), colloquially known as long COVID, but there is no universally agreed upon case definition for PASC. Recently, the National Academies proposed defining PASC as an infection-associated chronic condition persisting for at least three months after acute SARS-CoV-2 infection [5]. The World Health Organization instead suggests that a patient can be considered to experience PASC if, three months after infection, they have experienced symptoms lasting at least two months that are not explainable by an alternative diagnosis [6].

Since there is no consensus on the duration of symptoms warranting a PASC diagnosis, it may be informative to study the distribution of time from SARS-CoV-2 infection — or, more feasibly, a suitable proxy such as onset of symptoms or confirmation of infection by testing — until symptom resolution. In this article, we are motivated by the problem of studying the duration of COVID-19 symptoms using data collected as part of the Husky Coronavirus Testing (HCT) initiative. From September 2020 to June 2023, the HCT study enrolled 40,000 participants at the University of Washington (UW) in Washington state [7]. Beginning in December 2021, roughly at the beginning of the initial wave of Omicron variant infections, participants who reported a SARS-CoV-2 infection confirmed by either polymerase chain reaction (PCR) test or at-home rapid antigen test were invited to complete a follow-up survey regarding the persistence of symptoms in the months following infection. The aim of this survey was to estimate how common persistent post-acute COVID-19 symptoms are, and to identify risk factors for persistence in a population of university affiliates. Notably, this population is predominantly young and highly vaccinated, with relatively few preexisting comorbidities. In the survey, participants were asked whether, at the time of survey completion, they continued to experience symptoms associated with their recent infection.

The exact time from infection until symptom resolution was not observed for any participant in the HCT long COVID study. Instead, for each survey respondent, in addition to baseline covariate information collected at or before the test date, the available data consist of the time elapsed from positive SARS-CoV-2 test until survey response — referred to as the survey response time — and a binary variable indicating whether the respondent’s symptoms had resolved as of the survey response time — referred to as the event indicator. Each participant was observed only at a single survey response time. Response times in this study ranged from 28 days (the earliest time after a positive test at which surveys were sent) to 286 days. Because by design investigators recorded only the current status of a participant’s symptom persistence at their response time, this sampling scheme is referred to as current status sampling. Current status data, which represent an extreme form of interval censoring, arise in a number of epidemiological settings, including partner studies of HIV infection and cross-sectional studies for estimating the age of onset of non-fatal diseases, and require the use of specialized statistical techniques [8]. Study designs analogous to that of the HCT long COVID survey may be attractive because they are simple to implement, economical, and most importantly, less susceptible to recall bias than retrospective designs in which participants are asked to recall symptom status at a specific date in the past [9]. As with other survey designs, nonresponse is a potential concern: When the study concluded in June 2023, a large fraction of participants had not yet responded to the long COVID survey.

Interest lies in (1) estimating the distribution of the time-to-symptom-resolution variable and (2) identifying baseline patient characteristics associated with time to resolution, that is, risk factors for delayed symptom resolution. To achieve (1), we show here that recent methodological advances we have made in causal inference under shape constraints can be adapted to make statistical inference on the time-to-symptom-resolution distribution. In particular, this approach allows us to use flexible machine learning tools and thereby minimize the need for restrictive modeling assumptions. To our knowledge, the link between the methodology we build upon here, termed causal isotonic regression (CIR), and the analysis of current status data has not been made before; we aim to highlight the utility of this approach. To achieve (2), we make use of existing regression methods for current status data, which we briefly describe in the *Methods* section below, with further details given in the Supplementary Material. Our chosen approaches to both (1) and (2) rely on identical assumptions regarding the survey response process, namely, that survey response times are

uninformative for symptom resolution within strata of measured baseline covariates. In addition to illustrating statistical tools tailored to the analysis of current status data, we provide substantive results regarding persistent COVID-19 symptoms in the HCT population during the early Omicron period of the SARS-CoV-2 pandemic.

## Methods

### *Current status data structure*

We begin by describing in general the current status data structure we are considering, using the HCT study for concreteness. For each recipient of the long COVID survey, a collection  $W$  of baseline covariates, observed at or before the time of a positive SARS-CoV-2 test, were recorded. Estimating the distribution of symptom duration  $T$ , the time elapsed between infection and symptom resolution, is challenging because (1)  $T$  is never observed directly and (2) the study is terminated before all participants have responded to the survey.

In an ideal study with unlimited follow-up, we would observe for each participant the survey response time  $Y_*$  — the time elapsed between positive SARS-CoV-2 test and survey response — and the event indicator  $\Delta_*$  — a participant’s symptom status at the time of survey response — with  $\Delta_* = 1$  representing symptom resolution by time  $Y_*$  and  $\Delta_* = 0$  otherwise. We denote by  $c_0$  the investigator-determined time (since the positive test) at which follow-up ends for a participant, and consider as a nonrespondent any participant without a response by time  $c_0$ . We set  $Y = \min(Y_*, c_0)$  as the smaller of the survey response time and  $c_0$ , and denote by  $\Delta = \mathbb{1}(Y_* < c_0)\Delta_*$  the observed response indicator, which equals 1 if a participant responded and reported resolution and 0 otherwise. The available data are considered to be  $n$  independent and identically distributed observations  $(W_1, Y_1, \Delta_1), (W_2, Y_2, \Delta_2), \dots, (W_n, Y_n, \Delta_n)$ .

The two primary goals of our analysis, as described earlier, can be stated symbolically as using the available data to:

1. estimate the survival function of  $T$ ;
2. estimate the association between baseline covariates  $W$  and symptom duration  $T$ .

Below, we describe how these two goals may be achieved.

### *Statistical procedure*

For a given time  $t$ , we denote by  $\psi(t) = P(T > t)$  the survival function of  $T$  evaluated at  $t$ . This quantity represents the proportion of individuals testing positive for SARS-CoV-2 who experience symptom persistence beyond time  $t$  days. Our aim is to estimate  $\psi(t)$  over an interval  $[t_0, t_1]$  strictly contained in  $(b_0, c_0)$ , where  $b_0 = 28$  days is the smallest observable response time. Estimating this function using current status data requires dedicated statistical tools, since failure to account for the sampling scheme can lead to misleading conclusions.

Traditionally, under current status sampling, the survival function has been estimated based on the nonparametric maximum likelihood estimator (NPMLE), which can be computed using the pool adjacent violators algorithm used for isotonic (i.e., monotone) regression [10]. The key idea behind this approach is that estimation of the survival function can be reframed as a regression problem under a monotonicity constraint: the probability of symptom resolution by time  $t$  is simply the proportion of individuals who have resolved symptoms among those with survey response time  $t$ , and this probability is necessarily nondecreasing as a function of survey response time. This estimator has the advantage of being distribution-free, in the sense that its validity does not rely on any assumption about the underlying shape of the true survival function, and of enabling the construction of valid confidence intervals derived from its large-sample properties [11]. Nevertheless, an important limitation of the NPMLE approach is that it requires the survey response process to be entirely uninformative, in the sense that the response time itself does not carry information about the event time. This condition may fail to hold in practical settings, as there may exist baseline covariates — analogous to confounders in a causal inference setting — that inform the relationship between the event indicator and the survey response time. For example, participants with substantial comorbidities may be more likely to respond promptly to the survey and to have greater symptom persistence. We argue here that the framework of CIR introduced in Westling et al. [12], which was developed to estimate a monotone causal dose-response curve, can be directly used to estimate the survival function with current status data under weaker conditions. Specifically, as opposed to the traditional NPMLE approach, this method only requires independence between the survey response and event times *within strata* defined by the recorded baseline covariates. Such independence condition is typically much more realistic, and thereby makes the analysis less likely

to yield misleading conclusions due to assumption violations.

To describe how CIR can be used to analyze current status data, we introduce the potential outcome notation  $\Delta_*(t)$  to denote the value that  $\Delta_*$  would take were  $Y_*$  set to value  $t$ , e.g., in a hypothetical setting where the investigator controlled the survey response time. We note that, among survey respondents, the observed event indicator  $\Delta$  corresponds to  $\Delta_*(Y)$ , since by definition  $Y < c_0$  for respondents. The survival function evaluated at time  $t$  can then be written as  $\psi(t) = 1 - E\{\Delta_*(t)\}$ . This function is necessarily monotone since the probability of symptom persistence at survey response time  $Y_* = t$  cannot increase in  $t$ . If  $\Delta_*(t)$  and  $Y_*$  are conditionally independent given  $W$ , then the target parameter evaluated at any  $t$  in  $[t_0, t_1]$  can be written as

$$\begin{aligned}
\psi(t) &= 1 - E\{\Delta_*(t)\} && \text{(by definition)} \\
&= 1 - E[E\{\Delta_*(t) | W\}] && \text{(by law of total expectation)} \\
&= 1 - E[E\{\Delta_*(t) | W, Y_* = t\}] && \text{(by conditional independence)} \\
&= 1 - E[E\{\Delta_*(t) | W, Y = t\}] && \text{(since } Y_* = t \text{ if and only if } Y = t \text{ for } t < c_0) \\
&= 1 - E\{E(\Delta | W, Y = t)\} && \text{(since } \Delta = \Delta_*(Y)) .
\end{aligned}$$

Therefore, under this conditional independence assumption, the task of estimating  $\psi(t)$  is equivalent to estimation of what may be viewed as a monotone causal dose-response curve, where the exposure and potential outcomes are assumed to be independent conditional on measured covariates.

Computation of the NPMLE involves performing an isotonic regression of  $\Delta_1, \Delta_2, \dots, \Delta_n$  on  $Y_1, Y_2, \dots, Y_n$  over the domain  $[t_0, t_1]$ . Computation of the CIR estimator is more complex since it includes a step to disentangle the dependence between event and survey response times. This step involves estimation of two functions depending on baseline covariates and survey response times: (1)  $\mu(y, w) = E(\Delta | Y = y, W = w)$ , the conditional mean of  $\Delta$  given  $Y$  and  $W$ ; and (2)  $g(y, w) = \pi(y | w) / E\{\pi(y | W)\}$ , a standardization of the conditional density  $\pi$  of  $Y$  given  $W$ . These functions play a role akin to those of the outcome regression and (standardized) propensity score in causal inference for continuous exposures. To implement the CIR procedure, estimates  $\mu_n$  and  $g_n$  of  $\mu$  and  $g$ , respectively, are first constructed; this task can be accomplished with any off-the-shelf machine learning algorithm. Pseudo-outcomes  $\Gamma_n(\Delta_1, Y_1, W_1), \Gamma_n(\Delta_2, Y_2, W_2), \dots, \Gamma_n(\Delta_n, Y_n, W_n)$ ,

defined pointwise as

$$\Gamma_n(\delta, y, w) = \frac{\delta - \mu_n(y, w)}{g_n(y, w)} + \sum_{j=1}^n \mu_n(y, W_j) ,$$

are then obtained and regressed against  $Y_1, Y_2, \dots, Y_n$  over the domain  $[t_0, t_1]$  using isotonic regression (see Supplementary Material for technical details). We denote the resulting estimator of  $\psi(t)$  as  $\psi_n(t)$ . Use of these pseudo-outcomes rather than the simple event indicators accounts for possible dependence between  $\Delta_*(t)$  and  $Y_*$ , as long as this dependence is explained entirely by  $W$ . Roughly speaking, if at least one of  $\mu$  and  $g$  is estimated consistently, then  $\psi_n(t)$  is itself consistent for  $\psi(t)$ , and if both  $\mu$  and  $g$  are estimated consistently,  $n^{1/3}\{\psi_n(t) - \psi(t)\}$  tends in distribution to a random variable with a scaled Chernoff distribution [13]. This well-characterized large-sample distribution allows for construction of confidence intervals with correct coverage. A key strength of the CIR approach is that it allows for flexible estimation of  $\mu$  and  $g$  — using machine learning, for example — without sacrificing statistical inference. This reduces the risk of systematic bias due to the use of misspecified parametric models.

We note that performing the CIR procedure using data from survey respondents alone — in other words, performing a complete case analysis — could result in bias, since whether a participant is considered a respondent is determined by their value of  $Y_*$ , which may depend on  $W$  and thus be informative for  $\Delta_*$ . Symbolically, the estimand in a complete case analysis would be  $1 - E\{E(\Delta | W, Y = y) | Y < c_0\}$  rather than  $1 - E\{E(\Delta | W, Y = y)\}$ . Therefore, fully removing nonrespondents — that is, participants with  $Y_* \geq c_0$  — may induce selection bias, as the distribution of  $W$  may differ among survey respondents and nonrespondents. To emphasize the difference between the complete case CIR and our proposed method, which includes all survey recipients, we refer to the latter as *extended CIR*.

It is also of interest to understand the association between baseline covariates  $W$  and symptom duration  $T$  in order to identify potential predictors of delayed symptom resolution. In the Supplementary Material, we review regression models that have been adapted for use with current status data and discuss the difficulties that arise in making inference using these models. Focusing on the Cox model, we conduct a simulation study to evaluate the performance of the nonparametric bootstrap for making inference, finding that Wald-type bootstrap confidence intervals based on a



normal approximation demonstrate good performance in modest sample sizes. The Cox model requires no special adaptation to account for nonresponse, as long as the regression model is correctly specified and includes all covariates that inform the response mechanism [14].

## **Analysis of HCT data**

### *Study setting and participants*

The HCT study was conducted at UW, a large public university in Washington state with a population composed of approximately 60,000 students and 32,000 faculty and staff. Beginning in autumn 2020, the HCT program provided free, voluntary PCR testing for SARS-CoV-2 and symptom surveillance for university affiliates. Individuals were sent daily messages via text or email to ask about new symptoms, exposures, and high-risk behaviors, any of which would trigger an invitation for PCR testing. Study participants were also able to walk-in for testing at any time, and for any reason, without an invitation. Once antigen testing became widely available in March 2022, individuals were asked on their daily messages to report any antigen test results. The inclusion criteria for HCT eligibility are described in the Supplementary Material. Informed consent was signed electronically at the time of enrollment.

For this analysis, we used data collected from December 2021 to June 2023. During this time period, participants who had a positive or inconclusive PCR test or who reported a positive rapid antigen test for SARS-CoV-2 were invited to participate in a follow-up survey to assess for ongoing symptoms of COVID-19 a minimum of four weeks after the positive test.

### *Data collection*

All data were collected electronically by sending participants an email or text notification with a link to a questionnaire. The surveys included:

- (a) *Enrollment questionnaire and quarterly update survey*: The enrollment questionnaire included demographics, COVID-19 vaccination information, and current pre-existing medical conditions, which were captured by selecting zero or more options from the list given in Table 1. A quarterly survey was used to update eligibility and vaccination status.

Category	Pre-existing conditions
Cardiometabolic	Diabetes, heart disease, high blood pressure
Respiratory	Asthma/reactive airway disease, chronic bronchitis, COPD/emphysema
Immunological	Cancer, immunosuppression
Seasonal allergies	—

**Table 1:** Pre-existing medical conditions on which baseline information was collected.

Category	Symptoms
Upper respiratory	Cough, ear pain/discharge, eye pain, loss of smell or taste, rash, runny/stuffy nose, sore/itchy/scratchy throat, trouble breathing
Systemic	Chills or shivering, feeling feverish, headache, muscle or body aches, sweats
Gastrointestinal	Diarrhea, nausea/vomiting
Fatigue	—

**Table 2:** List of symptoms during the acute infection phase.

- (b) *Daily attestation survey and symptom attestation survey at the time of SARS-CoV-2 test:* Participants were prompted to fill out a daily attestation survey that collected information on date of known exposures to COVID-19, and date(s) and results of rapid antigen test(s) performed at home. New or worsening symptoms in the past 24 hours were captured by selecting zero or more items from the list given in Table 2.
- (c) *Long COVID survey:* Starting in December 2021, all participants who tested positive or inconclusive for SARS-CoV-2 by PCR testing or those who attested to having had a positive rapid antigen test were contacted a minimum of 28 days after their first positive test date and invited to participate in the long COVID survey. Each participant was allowed to fill out the survey only once per academic year. For participants with multiple responses, only the first is included in this analysis. The survey included the question “Are you still experiencing symptoms related to your COVID-19 illness (e.g., brain fog, fatigue, headaches, etc.)?” Each participant was free to choose when to complete the survey. Data collection was terminated on June 23, 2023.

### *Risk factors*

For the statistical analyses, a subset of baseline covariates collected on participants was used. Demographic risk factors included self-reported sex (male or female) and age (categorized as 25 years and under, 26 to 40 years, 41 to 60 years, and over 60 years). Baseline comorbidities were categorized as cardiometabolic, respiratory, immunological, and/or seasonal allergies (see Table 1). Symptoms at baseline were categorized as upper respiratory symptoms, systemic symptoms, gastrointestinal symptoms, and/or fatigue (see Table 2). Baseline comorbidities and symptoms were included in analyses as binary (presence/absence). As a proxy for viral load, participants were categorized according to the manner in which they tested positive for SARS-CoV-2 (rapid antigen test versus PCR), and among those who tested positive via PCR, according to the cycle threshold (Ct) value. Higher Ct values correspond to a lower viral load. The categories used were rapid antigen test, PCR with Ct below 30, and PCR with Ct greater than 30. Participants were also categorized according to the time elapsed between most recent vaccination and positive test: less than three months, between three and six months, and more than six months. We restricted our analysis to vaccinated participants because there were few unvaccinated participants — less than 2% of the study cohort — and because vaccination was required by UW policy during the data collection period, rendering unvaccinated individuals possibly different from most of the UW population in important ways beyond vaccination status.

### *Implementation*

We implemented the extended CIR method using flexible machine learning tools to estimate  $\mu$  and  $g$ . To estimate  $\mu$ , we used the Super Learner [15, 16], a form of stacked regression, with the optimal linear combination of learners selected to minimize 10-fold cross-validated mean squared error — see Supplementary Material for algorithm library. To estimate  $\pi$ , we used the `haldensify` software package [17] with 10-fold cross-validation used to select optimal tuning parameters. The derivative of the survival curve, which is required for confidence interval construction, was estimated by numerical differentiation of a monotone Hermite spline approximation of the estimated survival curve. Confidence intervals with nominal 95% pointwise coverage were constructed using the quantiles of the Chernoff distribution.

We aimed to estimate the survival function for symptom resolution between  $t_0 = 30$  and  $t_1 = 90$  days after the positive test. We therefore chose the maximum follow-up time to be  $c_0 = 120$ ; participants with positive tests after February 24, 2023 were excluded to ensure all individuals in the analysis cohort had at least 120 days of follow-up. On one hand, large values of  $c_0$  result in more individuals excluded from the study period but make  $t_1$  further away from the largest possible observed survey response times, where the statistical performance of isotonic regression and related procedures is known to degrade. On the other hand, small values of  $c_0$  result in fewer individuals excluded from the study period but make  $t_1$  closer to the largest possible observed survey response times. Our choice of  $c_0$  was meant to achieve a balance between these competing considerations, and resulted in  $t_1$  lying at approximately the 97.5th percentile of observed survey response times; such choice performed well in a simulation study — see Supplementary Material for details.

To investigate the association between baseline risk factors and time from infection to symptom resolution, we fit a Cox proportional hazards regression model accounting for current status sampling using the `icenReg` software package [18]. Main terms were included in the model for each baseline covariate described above. Standard error estimates were calculated using a nonparametric bootstrap with 1000 bootstrap replicates, and used to construct 95% Wald-type confidence intervals and perform two-sided tests of the null regression coefficient hypothesis at a 0.05 level. For categorical risk factors, multivariate Wald tests were used to test the joint null hypothesis that all of the corresponding regression coefficients were zero.


### *Results*

After the removal of 284 individuals due to missing covariate information, the analysis cohort consisted of 3434 HCT study participants with positive SARS-CoV-2 tests prior to February 24, 2023. Of these, 1432 (41.7%) responded to the survey. Characteristics of the analysis cohort are summarized in Table 3. The median age was 23 years (IQR: 20–36), and approximately 64% of the participants reported female sex. Slightly over 33% of participants reported at least one baseline comorbidity, and roughly 52% were symptomatic at the time of testing. Most participants were not recently vaccinated. Survey respondents tended to be older than nonrespondents, with markedly higher prevalence of symptoms during the acute infection phase.

Figure 1 shows the estimated proportion of participants with persistent COVID-19 symptoms at

	Respondents <i>n</i> = 1432	Nonrespondents <i>n</i> = 2002	Overall analysis cohort <i>n</i> = 3434
Age, years			
25 and under	597 (41.7%)	1385 (69.2%)	1982 (57.7%)
26–40	361 (25.2%)	379 (18.9%)	740 (21.5%)
41–60	375 (26.2%)	183 (9.1%)	558 (16.2%)
61 and over	99 (6.9%)	55 (2.7%)	154 (4.5%)
Male	494 (34.5%)	746 (37.3%)	1240 (36.1%)
Symptoms, upper respiratory	1027 (71.7%)	685 (34.2%)	1712 (49.9%)
Symptoms, systemic	715 (49.9%)	473 (23.6%)	1188 (34.6%)
Symptoms, gastrointestinal	122 (8.5%)	107 (5.3%)	229 (6.7%)
Symptoms, fatigue	465 (32.5%)	324 (16.2%)	789 (23.0%)
Comorbidities, cardiometabolic	96 (6.7%)	49 (2.4%)	145 (4.2%)
Comorbidities, respiratory	147 (10.3%)	193 (9.6%)	340 (9.9%)
Comorbidities, immunological	29 (2.0%)	26 (1.3%)	55 (1.6%)
Comorbidities, seasonal allergy	433 (30.2%)	439 (21.9%)	872 (25.4%)
Time since last vaccination			
Under 3 months	307 (21.4%)	27 (1.3%)	334 (9.7%)
3–6 months	204 (14.2%)	123 (6.1%)	327 (9.5%)
More than 6 months	921 (64.3%)	1852 (92.5%)	2773 (80.8%)
Viral load			
PCR, over 30 Ct	273 (19.1%)	438 (21.9%)	711 (20.7%)
PCR, under 30 Ct	545 (38.1%)	897 (44.8%)	1442 (42.0%)
Rapid antigen test	614 (42.9%)	667 (33.3%)	1281 (37.3%)

**Table 3:** Baseline characteristics of the analysis cohort.



**Figure 1:** Covariate-adjusted survival curves for time from positive test until symptom resolution, estimated using the proposed extended causal isotonic regression method (top panel), along with the distribution of survey response times between 28 and 95 days (bottom panel). The solid black line shows the estimated proportion with ongoing symptoms at times from 30 days to 90 days after the positive test. The black points show the estimated proportion with ongoing symptoms at 30, 60, and 90 days. The dashed lines represent a pointwise 95% confidence interval. The gray vertical lines denote 30, 60, and 90 days.

times from 30 to 90 days following the positive test, along with the distribution of survey response times. From the extended CIR method, estimated symptom persistence proportions at 30, 60, and 90 days after the positive test were 19.0% (95% CI: 13.6–25.0), 13.0% (9.4–16.7), and 6.9% (2.9–10.9), respectively.

Table 4 contains the results of the Cox proportional hazards regression, with hazard ratios larger than 1 indicating faster symptom resolution. Participants identifying as male tended to experience faster symptom resolution, with an estimated hazard ratio for resolution of 1.21 (1.04–1.40,  $P = 0.018$ ). Conversely, participants reporting fatigue at baseline and participants with history of seasonal allergies tended to have longer lasting symptoms, with estimated hazard ratios of 0.77 (0.65–0.92,  $P = 0.004$ ) and 0.86 (0.74–1.00,  $P = 0.050$ ), respectively. In terms of viral load, participants with Ct values under 30 and participants who tested positive via rapid antigen tests had more persistent symptoms compared to participants with Ct values over 30, with estimated hazard ratios of 0.78 (0.64–0.96) and 0.75 (0.60–0.95), respectively ( $P = 0.033$ ).

Risk factor	Hazard ratio	95% confidence interval	<i>P</i> -value
Age, years			0.631
25 and under	(ref)	(ref)	(ref)
26–40	0.922	(0.767, 1.109)	—
41–60	0.921	(0.754, 1.124)	—
61 and over	1.078	(0.785, 1.480)	—
Male	<b>1.206</b>	<b>(1.039, 1.401)</b>	<b>0.014</b>
Symptoms, upper respiratory	1.078	(0.885, 1.314)	0.453
Symptoms, systemic	0.967	(0.803, 1.163)	0.719
Symptoms, gastrointestinal	0.887	(0.679, 1.160)	0.382
Symptoms, fatigue	<b>0.771</b>	<b>(0.645, 0.921)</b>	<b>0.004</b>
Comorbidities, cardiometabolic	0.808	(0.599, 1.088)	0.160
Comorbidities, respiratory	0.868	(0.694, 1.085)	0.214
Comorbidities, immunological	0.839	(0.489, 1.440)	0.525
Comorbidities, seasonal allergy	<b>0.859</b>	<b>(0.737, 1.000)</b>	<b>0.050</b>
Time since last vaccination			0.813
Under 3 months	(ref)	(ref)	(ref)
3–6 months	0.953	(0.752, 1.209)	—
More than 6 months	0.934	(0.759, 1.150)	—
Viral load			<b>0.033</b>
PCR, over 30 Ct	(ref)	(ref)	(ref)
PCR, under 30 Ct	<b>0.783</b>	<b>(0.637, 0.963)</b>	—
Rapid antigen test	<b>0.754</b>	<b>(0.597, 0.953)</b>	—

**Table 4:** Results of Cox proportional hazards regression for time to symptom resolution. Boldface font indicates statistical significance at a 0.05 level. For categorical variables, *P*-values correspond to a multivariate test of null regression coefficients. Ct: cycle threshold.

## Discussion

In this article, we presented modern methodology for the analysis of current status data arising in epidemiological studies, with a particular focus on the novel extension and application of causal isotonic regression to this setting. When it is of interest to estimate the distribution of disease duration under current status sampling, CIR offers a flexible approach accommodating the use of off-the-shelf machine learning tools. Unlike the commonly used nonparametric maximum likelihood estimator, the validity of CIR relies only on conditional independence of survey response times and event indicators within strata of measured baseline covariates. In addition, we have shown how CIR can be implemented to properly account for survey nonresponse.

We demonstrated the use of CIR, along with existing survival analysis regression methods, to analyze current status data from the long COVID survey study conducted at UW between December 2021 and June 2023. Participants were queried at least 28 days after testing positive, and even for those queried exactly at 28 days, many did not respond until days or weeks later. This variation in response times motivates the use of the statistical methods described in this paper. Using the long COVID survey data, we estimated the prevalence of persistent COVID symptoms in the UW population to be 19% at 30 days after a positive test, dropping to 7% at 90 days.

A retrospective study in a comparable university population found that 36% of survey respondents reported symptoms lasting at least 28 days from the end of the recommended 10-day isolation period [9]. Our results provide further evidence that, even in a predominantly young and vaccinated population, a substantial proportion of infected individuals are affected by persistent COVID-19 symptoms. Notably, because participants were asked only about persistent symptoms and not about new symptoms arising in the post-acute phase, our analysis may still represent an underestimate of the frequency of PASC, which may also be characterized by symptoms that differ from initial acute symptoms and arise after their resolution.

We also investigated potential risk factors for persistent symptoms using Cox proportional hazards regression, and found evidence that female sex, fatigue during the acute phase, history of seasonal allergies, and higher semi-quantitative viral load at time of testing are associated with slower resolution of symptoms. This is in line with previous studies that have identified female sex [19], severe fatigue during the acute infection phase [20], allergic rhinitis [21], and higher peak



semi-quantitative viral load in the four weeks following diagnosis [22] as potential risk factors for PASC.

The validity of extended CIR hinges on the untestable assumption that survey response times and event indicators are conditionally independent given measured baseline covariates, which is more likely to be satisfied through the inclusion of potential risk factors as covariates in the analysis. If the survey response time depends on symptom duration even within strata defined by baseline covariates, the extended CIR procedure could be susceptible to bias. The development of sensitivity analysis methods for this setting is an important avenue of future research. Further work is also needed to better understand the practical implications of the poor boundary performance of CIR in order to provide a principled approach to choose appropriate values of  $t_0$ ,  $t_1$  and  $c_0$ .

### **Acknowledgements**

We would like to thank the participants of the Husky Coronavirus Testing study. This work was supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE-2140004, the National Heart, Lung, and Blood Institute under grant R01-HL137808, and the CARES act under University of Washington Grant 624611.

### **Software and code**

The extended CIR method is implemented in the `survML` R package, available at <https://github.com/cwolock/survML>. Code to reproduce all results is available online at [https://github.com/cwolock/currstat\\_CIR\\_supplementary](https://github.com/cwolock/currstat_CIR_supplementary).

### **Conflict of interest statement**

HYC has consulted for Bill and Melinda Gates Foundation and Ellume, and has served on advisory boards for Vir, Merck and Abbvie. She has received research funding from Gates Ventures, and support and reagents from Ellume and Cepheid outside of the submitted work.

## References

- [1] Yvonne MJ Goërtz, Maarten Van Herck, Jeannet M Delbressine, Anouk W Vaes, Roy Meys, Felipe VC Machado, Sarah Houben-Wilke, Chris Burtin, Rein Posthuma, Frits ME Franssen, et al. Persistent symptoms 3 months after a SARS-CoV-2 infection: the post-COVID-19 syndrome? *ERJ Open Research*, 6(4), 2020.
- [2] Limei Liang, Bohan Yang, Nanchuan Jiang, Wei Fu, Xinliang He, Yaya Zhou, Wan-Li Ma, and Xiaorong Wang. Three-month follow-up study of survivors of coronavirus disease 2019 after discharge. *Journal of Korean Medical Science*, 35(47), 2020.
- [3] Bram van den Borst, Jeannette B Peters, Monique Brink, Yvonne Schoon, Chantal P Bleeker-Rovers, Henk Schers, Hieronymus WH van Hees, Hanneke van Helvoort, Mark van den Boogaard, Hans van der Hoeven, et al. Comprehensive health assessment 3 months after recovery from acute coronavirus disease 2019 (COVID-19). *Clinical Infectious Diseases*, 73(5): e1089–e1098, 2021.
- [4] Jan Choutka, Viraj Jansari, Mady Hornig, and Akiko Iwasaki. Unexplained post-acute infection syndromes. *Nature Medicine*, 28(5):911–923, 2022.
- [5] National Academics of Science, Engineering, and Medicine. A Long COVID Definition. <https://nap.nationalacademies.org/catalog/27768/a-long-covid-definition-a-chronic-systemic-disease-state-with>, 2024. Accessed June 16, 2024.
- [6] Joan B Soriano, Srinivas Murthy, John C Marshall, Pryanka Relan, and Janet V Diaz. A clinical case definition of post-COVID-19 condition by a Delphi consensus. *The Lancet Infectious Diseases*, 22(4):e102–e107, 2022.
- [7] Ana A Weil, Sarah L Sohlberg, Jessica A O’Hanlon, Amanda M Casto, Anne W Emanuels, Natalie K Lo, Emily P Greismer, Ariana M Magedson, Naomi C Wilcox, Ashley E Kim, et al. SARS-CoV-2 epidemiology on a public university campus in Washington State. *Open Forum Infectious Diseases*, 8(11), 2021.
- [8] Nicholas P Jewell and Mark van der Laan. Current status data: Review, recent developments and open problems. *Handbook of Statistics*, 23:625–642, 2003.
- [9] Megan Landry, Sydney Bornstein, Nitasha Nagaraj, Gary A Sardon Jr, Amanda Castel, Amita Vyas, Karen McDonnell, Mira Agneshwar, Alyson Wilkinson, and Lynn Goldman. Postacute sequelae of SARS-CoV-2 in university setting. *Emerging Infectious Diseases*, 29(3):519, 2023.
- [10] Miriam Ayer, H Daniel Brunk, George M Ewing, William T Reid, and Edward Silverman. An empirical distribution function for sampling with incomplete information. *The Annals of Mathematical Statistics*, pages 641–647, 1955.
- [11] Piet Groeneboom. Asymptotics for interval censored observations. Technical report, University of Amsterdam, 1987.
- [12] Ted Westling, Peter Gilbert, and Marco Carone. Causal isotonic regression. *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(3):719–747, 2020.
- [13] Piet Groeneboom and Jon A Wellner. Computing Chernoff’s distribution. *Journal of Computational and Graphical Statistics*, 10(2):388–400, 2001.

- [14] Roderick JA Little and Donald B Rubin. *Statistical Analysis with Missing Data*, volume 793. John Wiley & Sons, 2019.
- [15] Mark J van der Laan, Eric C Polley, and Alan E Hubbard. Super learner. *Statistical applications in genetics and molecular biology*, 6(1), 2007.
- [16] Sherri Rose. Mortality risk score prediction in an elderly population using machine learning. *American Journal of Epidemiology*, 177(5):443–452, 2013.
- [17] Nima S Hejazi, Mark J van der Laan, and David Benkeser. haldensify: Highly adaptive lasso conditional density estimation in R. *Journal of Open Source Software*, 7(77):4522, 2022.
- [18] Clifford Anderson-Bergman. icenreg: regression models for interval censored data in R. *Journal of Statistical Software*, 81:1–23, 2017.
- [19] Yapeng Su, Dan Yuan, Daniel G Chen, Rachel H Ng, Kai Wang, Jongchan Choi, Sarah Li, Sunga Hong, Rongyu Zhang, Jingyi Xie, et al. Multiple early factors anticipate post-acute COVID-19 sequelae. *Cell*, 185(5):881–895, 2022.
- [20] Qing Shen, Emily E Joyce, Omid V Ebrahimi, Maria Didriksen, Anikó Lovik, Karen Sól Sævarsdóttir, Ingibjörg Magnúsdóttir, Dorte Helenius Mikkelsen, Anna Bára Unnarsdóttir, Arna Hauksdóttir, et al. COVID-19 illness severity and 2-year prevalence of physical symptoms: an observational study in Iceland, Sweden, Norway and Denmark. *The Lancet Regional Health–Europe*, 35, 2023.
- [21] Eugene Merzon, Margaret Weiss, Beth Krone, Shira Cohen, Gili Ilani, Shlomo Vinker, Avivit Cohen-Golan, Ilan Green, Ariel Israel, Tzipporah Schneider, et al. Clinical and socio-demographic variables associated with the diagnosis of long COVID syndrome in youth: a population-based study. *International Journal of Environmental Research and Public Health*, 19(10):5993, 2022.
- [22] Scott Lu, Michael J Peluso, David V Glidden, Michelle Davidson, Kara Lugtu, Jesus Pineda-Ramirez, Michel Tassetto, Miguel Garcia-Knight, Amethyst Zhang, Sarah A Goldberg, et al. Early biological markers of post-acute sequelae of SARS-CoV-2 infection. *medRxiv*, pages 2023–07, 2023.
- [23] Ted Westling and Marco Carone. A unified study of nonparametric inference for monotone functions. *Annals of Statistics*, 48(2):1001, 2020.
- [24] Johann Pfanzagl and Johann Pfanzagl. *Estimation in semiparametric models*. Springer, 1990.
- [25] J Pfanzagl. *Contributions to a general asymptotic statistical theory*. Springer, 1982.
- [26] Jian Huang. Efficient estimation for the proportional hazards model with interval censoring. *The Annals of Statistics*, 24(2):540–568, 1996.
- [27] AJ Rossini and AA Tsiatis. A semiparametric proportional odds regression model for the analysis of current status data. *Journal of the American Statistical Association*, 91(434):713–721, 1996.
- [28] Daniel Rabinowitz, Anastasios Tsiatis, and Jorge Aragon. Regression with interval-censored data. *Biometrika*, 82(3):501–513, 1995.

- [29] Wei Pan. Extending the iterative convex minorant algorithm to the Cox model for interval-censored data. *Journal of Computational and Graphical Statistics*, 8(1):109–120, 1999.
- [30] Jerome H Friedman. Multivariate adaptive regression splines. *The Annals of Statistics*, 19(1): 1–67, 1991.
- [31] Trevor J Hastie. Generalized additive models. In *Statistical models in S*, pages 249–307. Routledge, 2017.
- [32] Marvin N Wright and Andreas Ziegler. ranger: A fast implementation of random forests for high dimensional data in C++ and R. *arXiv preprint arXiv:1508.04409*, 2015.
- [33] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

## Supplementary Material

### S1. Causal isotonic regression for current status data

#### S1.1 Data structure and target of estimation

We denote by  $T$  the unobserved symptom duration after testing positive for SARS-CoV-2. For each study participant, we observe a vector  $W \in \mathbb{R}^p$  of baseline covariates. In the context of current status sampling, the *ideal data* unit — that is, the data unit that would be observed given unlimited follow-up — consists of  $W$ , the survey response time  $Y_* \in [b_0, \infty)$ , and the event indicator  $\Delta_* \in \{0, 1\}$ . We denote by  $c_0$  the maximum follow-up time of the study. In reality, we observe  $Y := \min(Y_*, c_0)$  and  $\Delta := \mathbb{1}(Y_* < c_0)\Delta_*$ . For individuals who have not responded to the survey by time  $c_0$ , we do not observe the event indicator. Thus, the *observed data* unit is  $Z := (W, Y, \Delta)$ . In this setting, survey non-response corresponds to having failed to respond to the survey prior to time  $c_0$ . The available data consist of  $n$  independent observations  $Z_1, Z_2, \dots, Z_n$  draw from distribution  $P_0$ , which is assumed to lie in a nonparametric model  $\mathcal{M}$ . Here and after, we use the subscript 0 to denote a functional of the true data-generating distribution  $P_0$ , e.g.,  $E_0\{f(W)\} := E_{P_0}\{f(W)\}$ . For simplicity of notation, in what follows we consider estimation of the cumulative distribution function  $t \mapsto \psi_0(t) := P_0(T \leq t)$ . (In the main text,  $\psi$  was used to denote the survival function, which is simply equal to one minus the distribution function).

For a generic distribution  $P \in \mathcal{M}$ , we define the mapping

$$\theta_P : y_0 \mapsto E_P \{E_P(\Delta | W, Y = y_0)\}. \quad (1)$$

For all  $y_0 \in (b_0, c_0)$ , the target parameter  $\psi_0(y_0)$  is identified by the observed data as  $\theta_0(y_0)$ . If  $y_0 < c_0$ , then the event  $\{Y_* = y_0\}$  is equivalent to the event  $\{Y = y_0\}$ , and so this parameter is equal to the distribution function of  $T$  evaluated at  $y_0$  under the conditional independence assumption described in the main text.

We note that the outer expectation in (1) is taken with respect to the marginal distribution of  $W$  under sampling from  $P_0$ , as opposed to the distribution of  $W$  given  $Y < c_0$ . This suggests that performing a standard causal isotonic regression (CIR) procedure among survey respondents would

not in general yield a consistent estimator, as the marginal distribution of  $W$  and the conditional distribution of  $W$  given  $Y < c_0$  are not necessarily equal unless  $W$  and  $Y$  are independent.

### *S1.2 Estimation and inference*

Without making strong modeling assumptions, the target regression function  $\theta_0(y_0)$  is not pathwise differentiable and is therefore not estimable at  $n^{\frac{1}{2}}$ -rate. However, it is known to be monotone as a function of  $y_0$  and can be viewed as the target of an isotonic regression procedure. Using isotonic regression to estimate  $\theta_0(y_0)$  allows for the application of established statistical theory [23] to characterize the large-sample properties of the resulting estimator. We outline this isotonic regression procedure here.

For a generic  $P \in \mathcal{M}$ , we first define the outcome regression function  $\mu_P : (y, w) \mapsto E_P(\Delta | W = w, Y = y)$  and density ratio  $g_P : (y, w) \mapsto \pi_P(y | w) / f_P(y)$  with  $\pi_P$  the conditional density of  $Y$  given  $W$  and  $f_P$  the marginal density of  $Y$ . These nuisance functions play a key role in accounting for the possible dependence between  $\Delta_*$  and  $Y_*$ . We use  $F_P$  to denote the marginal distribution function of  $Y$ . We use  $F_P$  as a domain transformation function in the CIR procedure. While the CIR procedure could be performed without the use of this domain transformation, its use results in a procedure that is invariant to strictly increasing transformations of  $Y$  [12]. In other words, changing the scale on which response times are measured will have no impact on the estimated survival function.

We next define  $\psi_P := \theta_P \circ F_P^{-1}$  and corresponding primitive function  $\Psi_P(t) := \int_0^t \psi_P(u) du$ . The causal isotonic regression (CIR) procedure entails estimation of the parameter  $\Gamma_P := \Psi_P \circ F_P$ , which takes the form

$$\Gamma_P : y_0 \mapsto \iint \mathbb{1}(u \leq y_0) \mu_P(u, w) F_P(dw) Q_P(dw)$$

with  $Q_P$  the marginal distribution function of  $W$  under sampling from  $P$ . Unlike  $\theta_0(y_0)$ , the parameter  $\Gamma_0(y_0)$  is pathwise differentiable relative to the nonparametric model  $\mathcal{M}$  and can therefore be estimated at  $n^{\frac{1}{2}}$ -rate [24]. Per the results of Westling et al. [12], the efficient influence function

of  $P \mapsto \Gamma_P(y_0)$  relative to  $\mathcal{M}$  evaluated at  $P_0$  is given by

$$D_{0,y_0}^* : (\delta, y, w) \mapsto \mathbb{1}(y \leq y_0) \left\{ \frac{\delta - \mu_0(y, w)}{g_0(y, w)} \right\} + \int \mathbb{1}(u \leq y_0) \mu_0(u, w) F_0(du) + \mathbb{1}(y \leq y_0) \theta_0(y) - 2\Gamma_0(y_0) .$$

Knowledge of this efficient influence functions enables standard debiasing approaches, such as the one-step correction [25], to be used in estimation of  $\Gamma_0(y_0)$ . This yields a debiased, asymptotically linear estimator  $\Gamma_n(y_0)$  of  $\Gamma_0(y_0)$ , which is used in the isotonic regression procedure to construct an estimator of  $\theta_0(y_0)$ . The procedure is as follows.

- (1) Select domain bounds  $t_0$  and  $t_1$  satisfying  $[t_0, t_1] \subset (b_0, c_0)$ .
- (2) Construct an estimate  $P_n$  of  $P_0$  using estimators  $(\mu_n, g_n, F_n, Q_n)$  of  $(\mu_0, g_0, F_0, Q_0)$ . For example,  $\mu_n$  and  $g_n$  can be constructed using machine learning algorithms, while  $F_n$  and  $Q_n$  can be constructed using the empirical distributions of  $Y$  and  $W$ , respectively, based on the observed data  $Z_1, Z_2, \dots, Z_n$ .
- (3) Let  $\mathcal{R}_n$  denote the set of unique values of  $Y_1, Y_2, \dots, Y_n$ . For each  $y_0 \in \mathcal{R}_n$  satisfying  $t_0 \leq y \leq t_1$ , compute the one-step debiased estimator  $\Gamma_n(y_0)$  of  $\Gamma_0(y_0)$ , given by  $\Gamma_n(y_0) := \Gamma_{P_n}(y_0) + \frac{1}{n} \sum_{i=1}^n D_{P_n, y_0}^*(Z_i)$ .
- (4) Compute the greatest convex minorant (GCM) of the set of points

$$\{(0, 0)\} \cup \{(F_n(Y_i), \Gamma_n(Y_i)) : i = 1, 2, \dots, n, t_0 \leq Y_i \leq t_1\}$$

Set  $\psi_n$  to be the left derivative of the GCM.

- (5) Compute  $\theta_n(y_0) := \psi_n(\Phi_n(y_0))$  as an estimator of  $\theta_0(y_0)$ .

If both  $\mu_n$  and  $g_n$  converge to  $\mu_0$  and  $g_0$ , Theorem 2 of Westling et al. [12] establishes that, for any  $y_0 \in (b_0, c_0)$  with  $F_0(y_0) \in (0, 1)$ ,  $n^{\frac{1}{3}} \{\theta_n(y_0) - \theta_0(y_0)\}$  converges in distribution to  $\tau_0(y_0)^{\frac{1}{3}} \mathbb{W}$ , where  $\mathbb{W}$  follows a standard Chernoff distribution [13] and  $\tau_0(y_0) := 4\theta_0'(y_0)\kappa_0(y_0)/f_0(y_0)$  with

$$\kappa_0(y_0) := \int \frac{E_0[\{\Delta - \mu_0(y_0, w)\}^2 | W = w, Y = y_0]}{g_0(y_0 | w)} Q_0(dw) .$$

In light of the fact that  $\Delta$  is binary, we can rewrite the above display as

$$\kappa_0(y_0) = \int \frac{\mu_0(y_0, w) \{1 - \mu_0(y_0, w)\}}{g_0(y_0 | w)} Q_0(dw) .$$

We refer to  $\tau_0$  as the scale factor. Construction of confidence intervals requires estimation of  $\tau_0$ . We note that both  $\kappa_0$  and  $f_0$  can be estimated by simply plugging in available estimates of  $\mu_0$  and  $g_0$ , with the empirical distribution of  $W_1, W_2, \dots, W_n$  used to estimate  $Q_0$ . To estimate  $\theta'_0(y_0)$ , the derivative of  $y \mapsto \theta_0(y)$  evaluated at  $y_0$ , there are a number of available options. In our implementation, we construct a monotone Hermite spline approximation of  $\theta_n$  and perform numerical differentiation.

## S2. Regression analysis

Many popular regression models used in survival analysis have been adapted for use with current status data, including the Cox proportional hazards model [26], the proportional odds model [27], and the accelerated failure time model [28]. These methods are generally valid under the same assumptions on the survey response process as the adaptation of the CIR method described above. They do, however, make additional assumptions on the relationship between  $T$  and  $W$ .

Fitting regression models using current status data is generally more complicated than with right censored data. For the Cox model, unlike in the right-censored setting, estimation of the regression coefficients cannot be completely separated from estimation of the baseline hazard function through the partial likelihood [26]. Nonetheless, estimation and valid inference is still attainable for the regression coefficients, with the main difficulties being computational in nature. Recent work has focused on developing more efficient algorithms for computing estimates of the regression coefficients via maximum likelihood [18]. Calculation of the observed Fisher information, which is used to compute analytic standard error estimates, is generally impractical. The bootstrap has been proposed as an alternative approach [29], although as of yet there is no theoretical proof of its validity in this context. In Section S3, we present results of a simulation study investigating the performance of the nonparametric bootstrap for construction of confidence intervals using current status data for inference on regression coefficients of a Cox model. Wald-type bootstrap confidence intervals based on a normal approximation demonstrate good performance, while percentile boot-



strap confidence intervals tend to be slightly anticonservative. Based on these findings, we make use of the Wald-type interval in the regression analyses we report.

As with the Cox model, the other regression models described above do not need to be altered to account for survey nonresponse, under the assumption that the regression model is correctly specified and includes all covariates that inform the response mechanism [14].

### S3. Simulation studies

#### S3.1 Causal isotonic regression simulations

In this section, we describe a simulation study conducted to evaluate the performance of the proposed extended CIR accounting for survey nonresponse. While we expect the complete case CIR procedure to perform well when follow-up is unlimited, we expect decreased performance when termination of follow-up precludes survey response for some study participants. In contrast, we expect the extended CIR procedure to perform well in both of these scenarios.

In this experiment, the data were generated according to a proportional hazards model. We simulated a covariate vector  $W = (W_1, W_2, W_3)$  of three independent components, each distributed according to a  $\text{Unif}(\{-1, 1\})$  distribution. Given covariate vector  $W = w$ , we then simulated the survey response time  $Y_*$  and event time  $T$  from independent Weibull distributions with shape parameter 0.75 and scale parameter  $\exp(\frac{2}{5}w_1 - \frac{1}{5}w_2)$ . For computational purposes,  $Y_*$  was set to the closest point on a grid of 50 values evenly spaced on the quantile scale. In other words, the values of  $Y_*$  were coarsened to 50 unique values.

For each simulation replicate, the aim was to estimate the distribution function of  $T$  over the interval  $[t_0, t_1]$  with  $t_0 = 0.02$  (corresponding to the 5th percentile of observed response times) and  $t_1 = 1.5$ . The maximum follow-up time was set to  $c_0 \in \{2.1, 1.8, 1.65\}$ . We refer to the settings corresponding to these values of  $c_0$  as Scenarios 1, 2 and 3, respectively. Under these three choices of  $c_0$ ,  $t_1$  corresponded to the 90th, 95th, or 97.5th quantile of observed survey response times, respectively. This allowed us to explore whether performance of the complete case CIR or extended CIR procedures would degrade if  $c_0$  was chosen such that nonresponse was more prevalent and  $t_1$  was nearer the edge of the support of observed response times. For each value of  $c_0$ , the observed data consisted of  $n$  independent replicates of  $(W, Y, \Delta)$  with  $Y := \min(Y_*, c_0)$  and

$\Delta := \mathbb{1}(Y_* < c_0)\mathbb{1}(T \leq Y)$ . We generated datasets of size  $n = \{500, 1000, 1500, 2000\}$ .

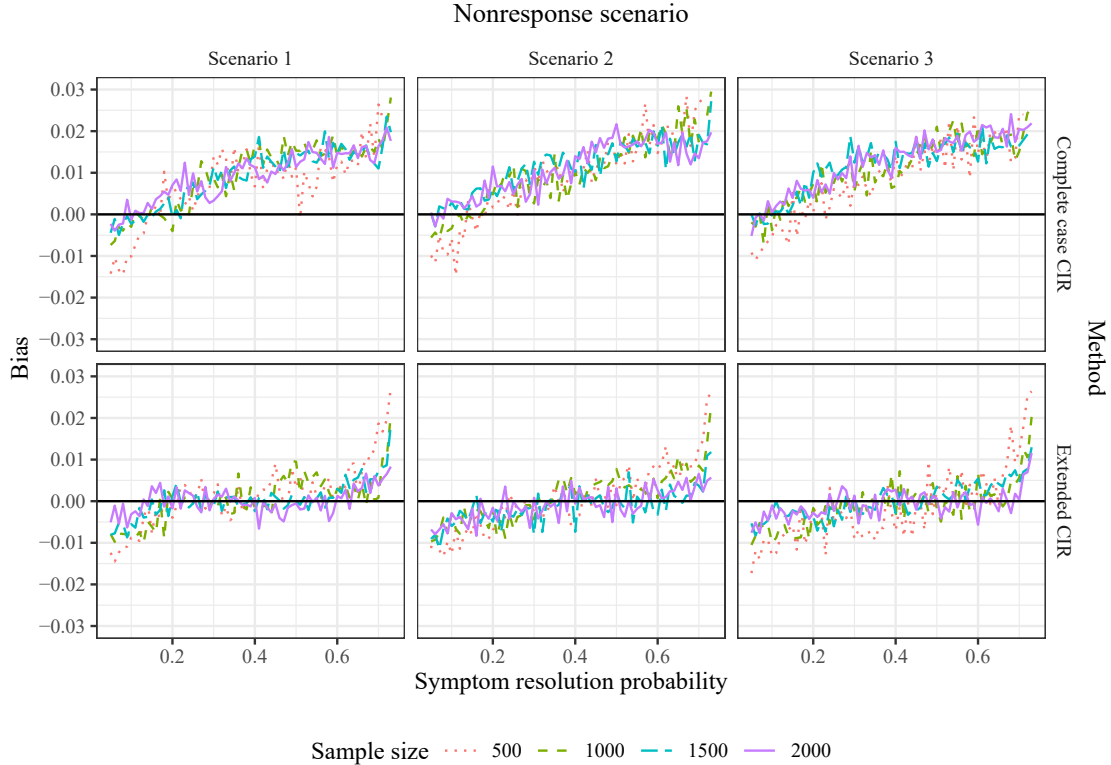
The survival function was estimated either using only complete cases (i.e., survey respondents, or those with  $Y_* < c_0$ ) with the existing CIR procedure described in [12] or using all participants with the extended CIR procedure proposed here. Nominal 95% confidence intervals were constructed by estimating the scale factor and using the quantiles of the Chernoff distribution. The nuisance function  $\mu$  was estimated using Super Learner [15] with a library consisting of the marginal mean, logistic regression, multivariate adaptive regression splines [30], generalized additive models [31], and random forests [32]. The conditional density  $\pi$  was estimated using `haldensify` [17] with 10-fold cross-validation used to select optimal tuning parameters. The derivative of the survival curve was estimated by numerical differentiation of a monotone Hermite spline approximation of the estimated survival curve. For each simulation replicate, the survival function of  $T$  was estimated at 100 time-points, evenly spaced on the quantile scale, between  $t_0 = 0.02$  and  $t_1 = 1.5$ . For each setting and method, we conducted 1000 simulation replicates. Pointwise and integrated (over all 100 time-points) performance was quantified using the empirical bias and empirical confidence interval coverage.

From Figures S1 and S2, we see that the extended CIR generally has smaller bias compared to the complete case CIR, although both methods have generally inflated bias near the estimation boundaries. From Figures S3 and S4, we observe that the extended CIR demonstrates strong confidence interval coverage across nonresponse scenarios, while the complete case CIR tends to be somewhat anticonservative on average.

### *S3.2 Cox proportional hazards regressions simulations*

In this section, we describe a simulation study performed to evaluate the validity of the nonparametric bootstrap for constructing confidence intervals for regression coefficients of the Cox proportional hazards model when the data are subject to current status sampling and survey nonresponse. While this is not meant to replace rigorous theoretical analysis, it may help inform the practical use of the bootstrap in this setting.


Data were generated as in the previous set of experiments. To evaluate the performance of the bootstrap, we fit a correctly specified Cox model (i.e., an additive model with only first-order terms) on each of  $B$  bootstrap resamples of the dataset. We constructed confidence intervals for



**Figure S1:** Pointwise bias of the CIR procedure for estimating the survival function under current status sampling with nonresponse. From left to right, the columns represent nonresponse Scenarios 1, 2 and 3, ordered from smallest to largest degree of nonresponse. The top and bottom rows denote the standard CIR method using only complete cases and the extended CIR method accounting for nonresponse. The  $x$ -axis displays the true value of the distribution function. The black line denotes zero bias.

each regression coefficient using either a normal approximation with standard deviation equal to the standard deviation of the point estimates across the  $B$  replicates or a percentile bootstrap. We repeated this procedure 1000 times and calculated the empirical confidence interval coverage for each method of interval construction. We generated datasets of size  $n = \{500, 1000, 1500, 2000\}$  and used bootstrap replicates  $B = \{100, 250, 500, 1000\}$ . The empirical coverage for the coefficients corresponding to  $W_1, W_2$  and  $W_3$  in each of these settings is displayed in Figures S5–S7.

From Figures S5–S7, we observe that, under all nonresponse mechanisms, the normal approximation bootstrap performs well, with empirical coverage within Monte Carlo error of the nominal level of 0.95 across both sample sizes and number of bootstrap replicates. The percentile bootstrap produces somewhat anticonservative confidence intervals, although its performance tends to improve with a larger number of bootstrap replicates.



**Figure S2:** Integrated absolute bias of the CIR procedure for estimating the survival function under current status sampling with nonresponse. From left to right, the columns represent nonresponse Scenarios 1, 2 and 3, ordered from lowest to highest nonresponse level. The top and bottom rows denote the standard CIR method using only complete cases and the extended CIR method accounting for nonresponse. The  $x$ -axis corresponds to sample size. The black line denotes zero bias.

#### S4. Supplementary information for HCT data analysis

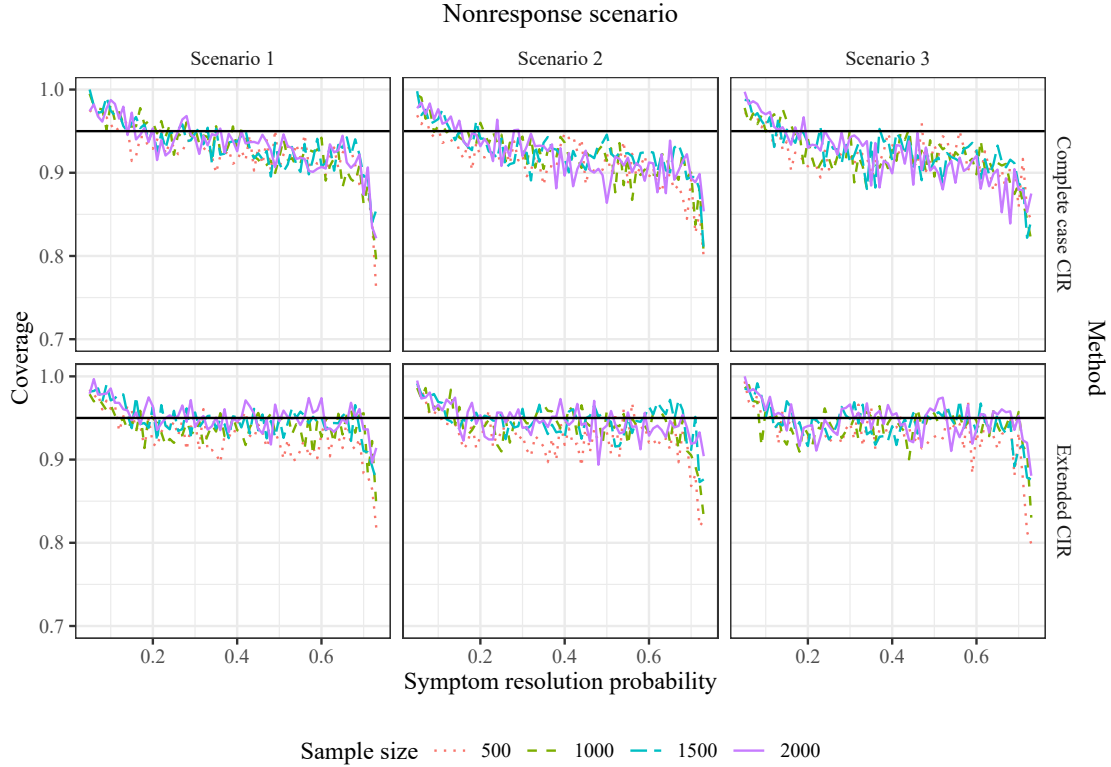
This section contains supplementary information on the analysis of the HCT long COVID survey data.

##### *S4.1 Eligibility criteria for HCT enrollment*

The eligibility criteria were

1. a valid university identification number;
2. age 13 years or older;
3. working or attending classes in-person at the university campus at least once per month
4. living within commuting distance of one of the three campuses
5. ability to provide consent in English.

Any individual who met the inclusion criteria could enroll in the parent HCT study.




**Figure S3:** Confidence interval coverage using the CIR procedure for estimating the survival function under current status sampling with nonresponse. From left to right, the columns represent nonresponse Scenarios 1, 2 and 3, ordered from smallest to largest degree of nonresponse. The top and bottom rows denote the standard CIR method using only complete cases and the extended CIR method accounting for nonresponse. The  $x$ -axis displays the true value of the distribution function. The black line denotes the nominal 95% coverage level.

#### S4.2 Algorithm library for Super Learner

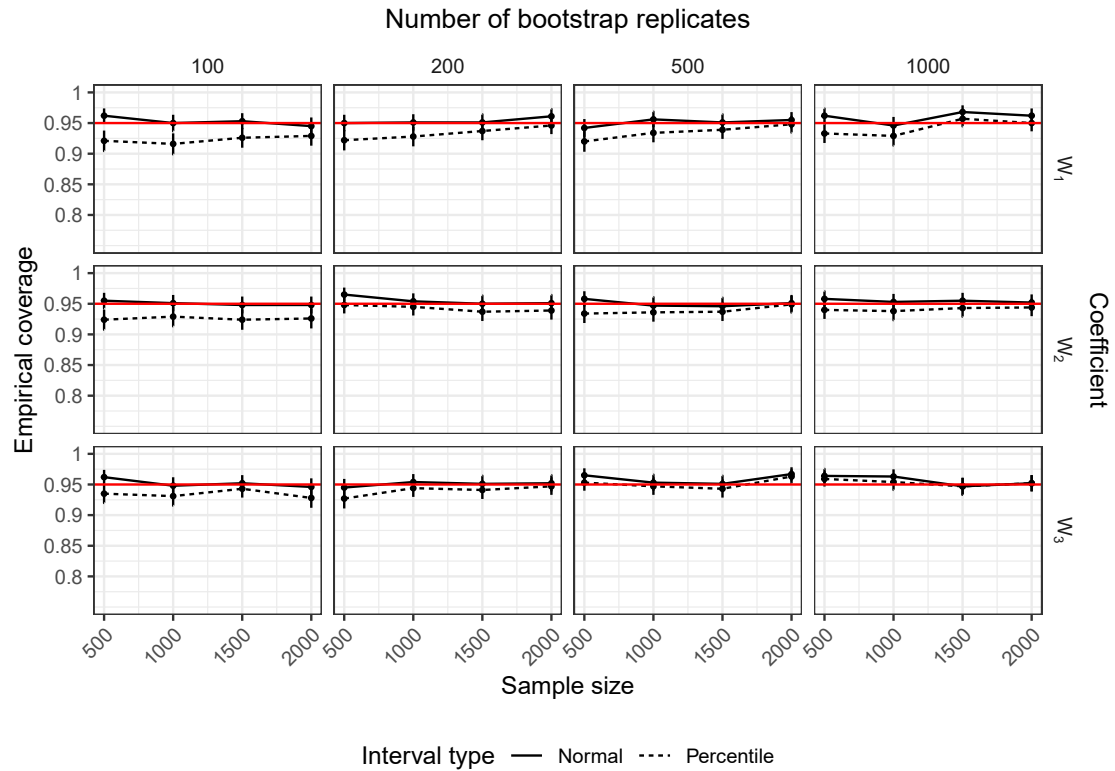
For data analyses and simulations Super Learner [15] was implemented with a library consisting of the marginal mean, logistic regression, multivariate adaptive regression splines [30], generalized additive models [31], random forests [32], and gradient boosted trees [33].

#### S4.3 Supplementary figures

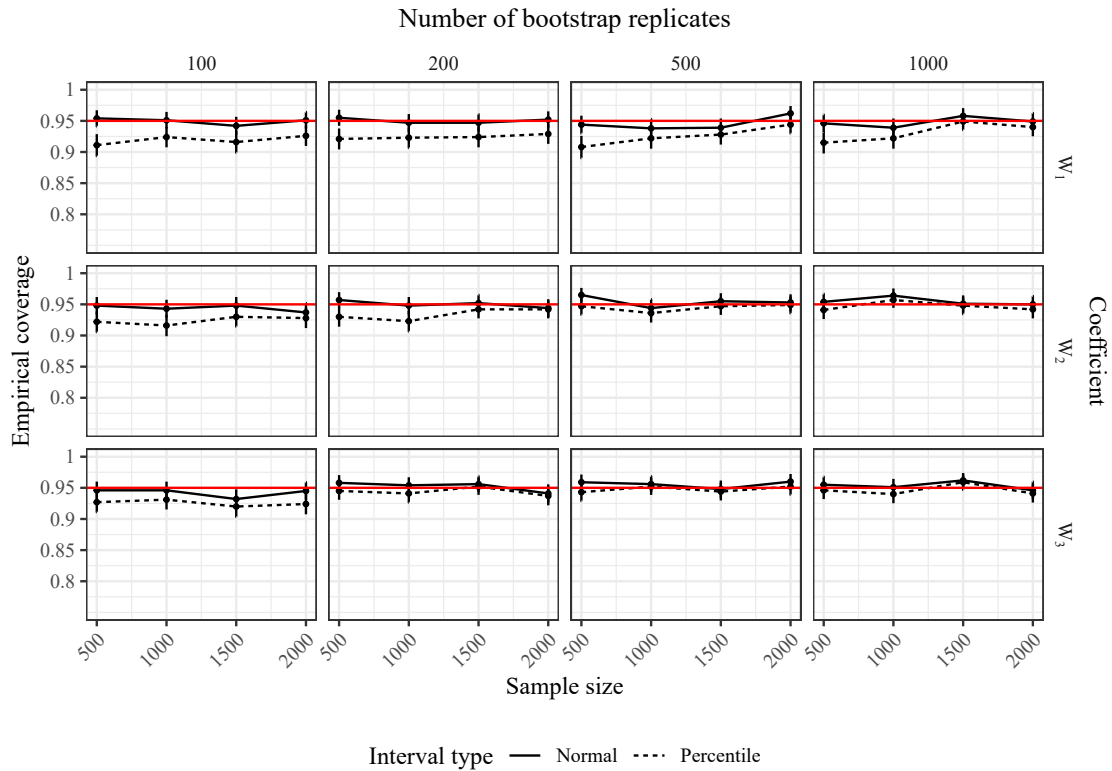
Figure S8 shows the distribution of survey response time  $Y$  among HCT survey respondents.



**Figure S4:** Confidence interval coverage using the CIR procedure for estimating the survival function under current status sampling with nonresponse. From left to right, the columns represent nonresponse Scenarios 1, 2 and 3, ordered from smallest to largest degree of nonresponse. The top and bottom rows denote the standard CIR method using only complete cases and the extended CIR method accounting for nonresponse. The  $x$ -axis corresponds to sample size. The black line denotes the nominal 95% coverage level.

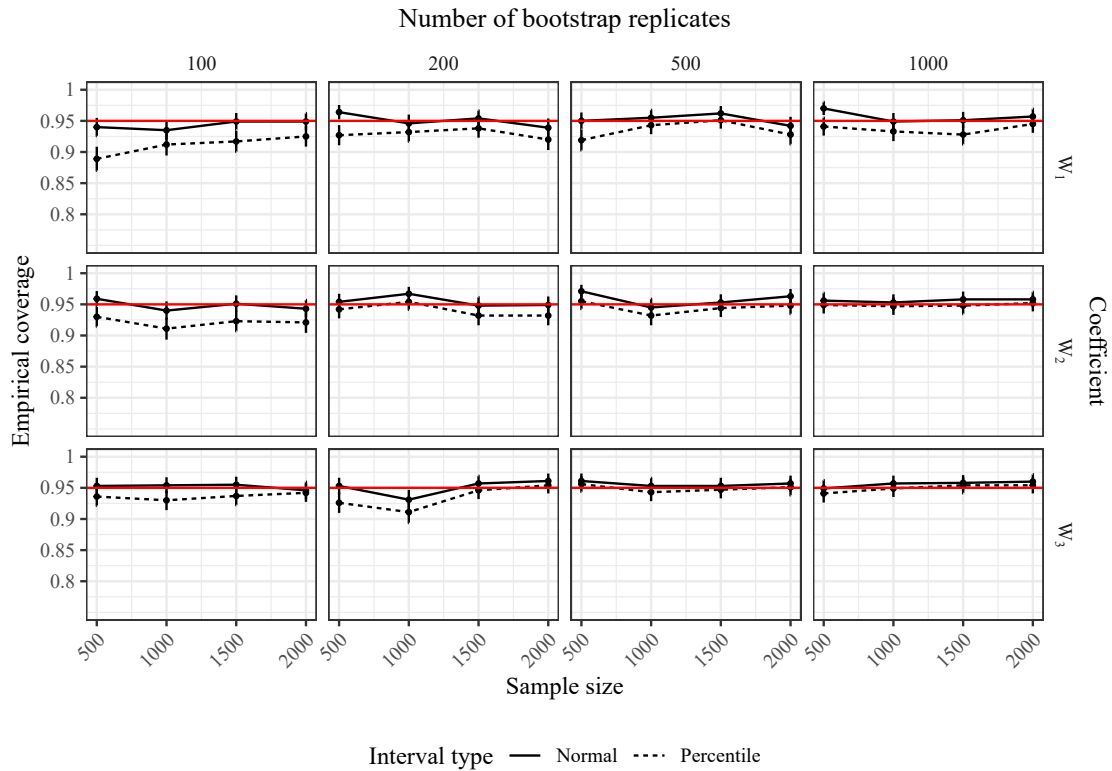


**Figure S5:** Performance of nonparametric bootstrap confidence intervals for the Cox model coefficients under current status sampling and nonresponse in Scenario 1. The columns represent different numbers of bootstrap replicates, and the rows represent the regression coefficients for the three covariates included in the model. In each plot, the solid black line shows the confidence interval coverage of nominal 95% Wald-type confidence intervals based on a Normal approximation, while the dashed black line shows the coverage of nominal 95% percentile confidence intervals. The red line denotes 95% coverage.

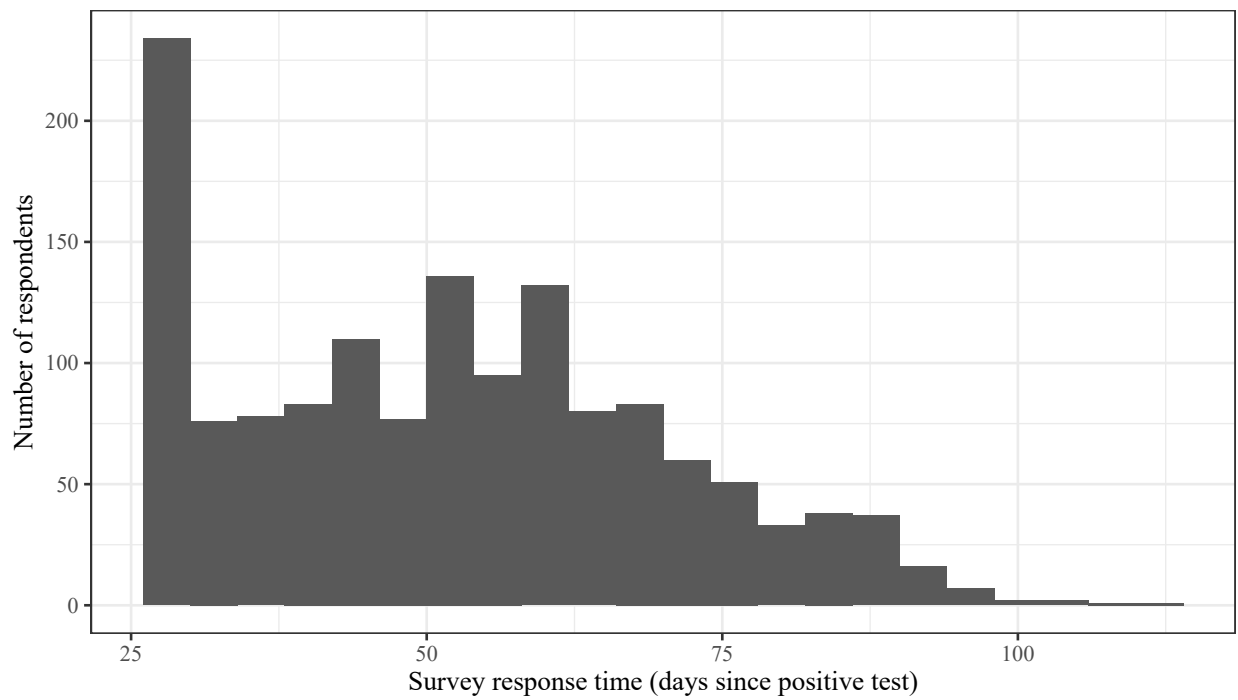


**Figure S6:** Performance of nonparametric bootstrap confidence intervals for the Cox model coefficients under current status sampling and nonresponse in Scenario 2. The columns represent different numbers of bootstrap replicates, and the rows represent the regression coefficients for the three covariates included in the model. In each plot, the solid black line shows the confidence interval coverage of nominal 95% Wald-type confidence intervals based on a Normal approximation, while the dashed black line shows the coverage of nominal 95% percentile confidence intervals. The red line denotes 95% coverage.





**Figure S7:** Performance of nonparametric bootstrap confidence intervals for the Cox model coefficients under current status sampling and nonresponse in Scenario 3. The columns represent different numbers of bootstrap replicates, and the rows represent the regression coefficients for the three covariates included in the model. In each plot, the solid black line shows the confidence interval coverage of nominal 95% Wald-type confidence intervals based on a Normal approximation, while the dashed black line shows the coverage of nominal 95% percentile confidence intervals. The red line denotes 95% coverage.



**Figure S8:** Histogram of survey response times. Surveys were sent at least 28 days after the positive test.