

Bounding infection prevalence by bounding selectivity and accuracy of tests: with application to early COVID-19

JÖRG STOYE

Department of Economics, Cornell University, Ithaca, NY 14853, USA.

Email: stoye@cornell.edu

First version received: 2 February 2021; final version accepted: 16 May 2021.

Summary: I propose novel partial identification bounds on infection prevalence from information on test rate and test yield. The approach utilizes user-specified bounds on (i) test accuracy and (ii) the extent to which tests are targeted, formalized as restriction on the effect of true infection status on the odds ratio of getting tested and thereby embeddable in logit specifications. The motivating application is to the COVID-19 pandemic but the strategy may also be useful elsewhere. Evaluated on data from the pandemic's early stage, even the weakest of the novel bounds are reasonably informative. Notably, and in contrast to speculations that were widely reported at the time, they place the infection fatality rate for Italy well above the one of influenza by mid-April.

Keywords: *Partial identification, bounds, COVID-19.*

JEL codes: *C14, C25.*

1. INTRODUCTION

Prevalence of a novel infection like SARS-CoV-2 (the virus causing COVID-19 disease) is a quintessential missing data problem. Only a small subset of the population has been tested, this subset is almost certainly selective; we do not even know the accuracy of tests, and our understanding of the pandemic is vague enough so that we might not want to overly rely on heavily parameterized models. This is a natural application for partial identification analysis, i.e., the analysis of bounds on parameter values that can be inferred from imperfect data and weak but credible assumptions, without forcing statistical identifiability of a model.¹ This paper proposes a general framework for analysing partial identification of prevalence, assuming that one has partially identifying information on the *selectivity* and *sensitivity* of diagnostic tests.

The obvious precedent for this is Manski and Molinari (2021, MM henceforth). I agree with that paper's overall thrust, but propose a considerably different implementation, refining worst-case bounds by bounding test sensitivity and selectivity but not predictive value (all terms will be defined later). These restrictions are readily related to other literatures, and—unlike with predictive values—nonvacuous prior bounds on them can be asserted without implying informative prior bounds on prevalence itself. In the empirical application, bounds that only restrict the direction of selectivity are considerably more informative than the analogous bounds emphasized in MM, and yet I will argue that assumptions became more compelling. The difference

¹ See Manski (2003) for an early monograph and Molinari (2020) for an extensive survey.

matters: the novel bounds on the infection fatality rate exclude ‘flu-like’ values that were the subject of speculation at the time. They become much tighter, though at the cost of reduced credibility, if one substantially restricts selectivity.

2. THE IDENTIFICATION PROBLEM

2.1. Basic setting and worst-case bounds

Consider first the problem of bounding prevalence of an infection in a stylized example where one has observed test rate and test yield for one population. I will call the disease COVID-19 henceforth, but the ideas are more general. For readability, I also follow common parlance and refer to COVID-19, though, strictly speaking, I investigate SARS-CoV-2 infection as opposed to COVID-19 disease.

Thus, let C indicate true infection status (with $C = 1$ indicating infection), T test status (with $T = 1$ indicating having been tested) and R test result (with $R = 1$ a positive test result; we observe R only conditionally on $T = 1$). In particular, define the testing rate $\tau := \Pr(T = 1)$ and the test yield $\gamma := \Pr(R = 1|T = 1)$. These objects are directly identified from the data, and we initially assume that they are known; indeed, inference will turn out to be a secondary concern. We also maintain the assumption that (PCR-)tests have specificity (=true negative rate $\Pr(R = 0|T = 1, C = 0)$) of 1; thus, $\Pr(C = 1|R = 1) = 1$. Generalizing away from this simplification would be straightforward.

Worst-case bounds on the true infection rate $\rho := \Pr(C = 1)$ can then be derived from the Law of Total Probability and the logical bound of $[0,1]$ on any unknown probability. In particular, write

$$\Pr(C = 1) = \Pr(C = 1|R = 1)\Pr(R = 1) + \Pr(C = 1|R = 0)\Pr(R = 0),$$

and observe that $\Pr(C = 1|R = 0) \in [0, 1]$, whereas $\Pr(R = 1) = \gamma\tau$ and $\Pr(C = 1|R = 1) = 1$ by maintained assumption. Thus, without any further assumption,

$$\rho \in [\gamma\tau, 1]. \tag{2.1}$$

These bounds go back to Manski (1989) in spirit and are also the starting point of MM. I next lay out novel ways to refine them.

2.2. Introducing bounds on sensitivity and selectivity of tests

Consider injecting prior information on *test sensitivity* (i.e., true positive rate $\Pr(R = 1|C = 1)$) and on *test selectivity* (i.e., the relation of $\Pr(T = 1|C = 1)$ to $\Pr(T = 1|C = 0)$) but not either of these two probabilities by itself). I do not claim that any of these are context-independent, much less known; hence, the prior information will itself be in the form of bounds. However, test sensitivity relates directly to a large medical literature, and test selectivity is readily related to statistical models of binary response. I next explain the approach and work out its implications.

Refinement: allow for measurement error through bounding sensitivity. Test sensitivity is the target parameter in much medical research. Thus, consider:

ASSUMPTION 2.1. *Sensitivity of the test is bounded by*

$$\Pr(R = 1|C = 1, T = 1) =: \pi \in [\underline{\pi}, \bar{\pi}].$$

Two remarks on Assumption 2.1 are in order. First, it is in principle testable because, with specificity of 1, one must have $\pi \leq \gamma$. Thus, if test yield γ exceeds $\underline{\pi}$, it refines the bounds, and if it exceed $\bar{\pi}$, it contradicts them. I will henceforth assume that $\underline{\pi} \geq \gamma$, a restriction that is far from binding in the empirical application and that can always be enforced before commencing the analysis.

Second, the assumption takes a notational shortcut: it seemingly implies that π is constant across true prevalence and test rates. This textbook view (Zhou et al., 2002, chapter 2) has been challenged (Leeflang et al., 2008). In the specific case of COVID-19, both prevalence and testing rate might influence sensitivity through the distribution of viral load among the tested and infected; the corresponding conjecture that asymptomatic surveillance might be characterized by lower sensitivity than symptomatic surveillance has some empirical corroboration (Mohammadi et al., 2020; Zhang et al., 2021). However, bounds derived below do not exploit constancy of π beyond the fact that $\pi \in [\underline{\pi}, \bar{\pi}]$. So this is best thought of as a shortcut to avoid further subscripts, though users should keep the consideration in mind when specifying $(\underline{\pi}, \bar{\pi})$.

The effect of Assumption 2.1 on prevalence bounds is as follows.

PROPOSITION 2.1. *Suppose Assumption 2.1 holds. Then prevalence is sharply bounded by*

$$\rho \in [\gamma\tau/\bar{\pi}, \gamma\tau/\underline{\pi} + 1 - \tau].$$

Proof. Write

$$\Pr(C = 1) = \Pr(C = 1|T = 1)\Pr(T = 1) + \Pr(C = 1|T = 0)\Pr(T = 0). \tag{2.2}$$

While no informative bound on $\Pr(C = 1|T = 0)$ is available, we have

$$\begin{aligned} & \Pr(R = 1|T = 1) \\ &= \Pr(R = 1|C = 1, T = 1)\Pr(C = 1|T = 1) + \underbrace{\Pr(R = 1|C = 0, T = 1)\Pr(C = 0|T = 1)}_{=0 \text{ by assumption}} \\ &= \underbrace{\Pr(R = 1|T = 1, C = 1)\Pr(C = 1|T = 1)}_{=\pi}, \end{aligned}$$

implying (in the notation introduced above) that $\Pr(C = 1|T = 1) = \gamma/\pi \in [\gamma/\bar{\pi}, \gamma/\underline{\pi}]$. The bounds follow by substituting into (2.2). \square

REMARK 2.1. Proposition 2.1 embeds the worst-case bounds (2.1) because these are generated by setting $(\underline{\pi}, \bar{\pi}) = (\gamma, 1)$, the widest bounds on π that are consistent with the data.

REMARK 2.2. This result is easily extended to allow for specificity (=true negative rate $\Pr(R = 0|C = 0, T = 1)$) to differ from 1. Indeed, the bounds simply adjust prevalence in the tested population through the well-known formula ‘prevalence=(yield+specificity-1)/(sensitivity+specificity-1)’ and leave prevalence in the untested population unconstrained. This is not worked out to economize on notation.

Refinement: a ‘logit bound’ on test selectivity. Consider also the following:

ASSUMPTION 2.2. *The factor $\kappa \geq 0$ in*

$$\frac{\Pr(T = 1|C = 1)}{1 - \Pr(T = 1|C = 1)} = \kappa \frac{\Pr(T = 1|C = 0)}{1 - \Pr(T = 1|C = 0)},$$

can be bounded as $\kappa \in [\underline{\kappa}, \bar{\kappa}]$.

Assumption 2.2 resembles sensitivity analysis for treatment effects in Rosenbaum (2002). It bounds the relative odds ratio of being tested between true positives and true negatives. Of course, this is only one of many possible ways to constrain the extent to which tests are targeted. However, it is easily related to binary response models of selection. In particular, bounding κ in the above is equivalent to bounding it in the logit model

$$\Pr(T = 1|C = c) = \frac{\exp(\alpha + c \ln \kappa)}{1 + \exp(\alpha + c \ln \kappa)}.$$

Logit models are well understood in econometrics and medical statistics, so this connection generates an interface to natural estimation strategies and maybe researcher intuitions about plausible parameter values.

For example, Canning et al. (2020) model the age-dependent effect of COVID-19 symptoms on social distancing behaviour through a logit; a similar model could in principle be used to model self-selection into symptomatic surveillance. If such a model were applied to the propensity to get tested, true infection status could be treated as hidden covariate. If one were furthermore willing to bound the coefficient on this covariate—where the bounds may depend on the values of other covariates—then conditionally on any realization of observed covariates, one would have the setting of Assumption 2.2. To give but one idea, the propensity of noninfected subjects to get tested could then be tied to the age-specific frequency of influenza-like symptoms.

The selectivity factor κ could be bounded from both above and below. For this paper’s application, I will impose throughout that $\kappa \geq 1$, thus there is at least weak selection of infected subjects into testing, and I will consider values of κ that force strict selection. Bounding selectivity from above, or also allowing for a lower bound below 1, may be interesting in other contexts, for example, if getting tested is stigmatized or tests are targeted but not at the at-risk population.

The implications of bounding κ are slightly more involved.

PROPOSITION 2.2. *Suppose that Assumptions 2.1 and 2.2 hold. Then prevalence is sharply bounded by*

$$\Pr(C = 1) \in \left[\frac{\gamma}{\bar{\pi}} \times \frac{\bar{\pi} + (\bar{\kappa} - 1)\tau(\bar{\pi} - \gamma)}{\bar{\kappa}(\bar{\pi} - \gamma) + \gamma}, \frac{\gamma}{\underline{\pi}} \times \frac{\underline{\pi} + (\underline{\kappa} - 1)\tau(\underline{\pi} - \gamma)}{\underline{\kappa}(\underline{\pi} - \gamma) + \gamma} \right].$$

including by the corresponding limiting expressions as $\underline{\kappa} \rightarrow 0$ or $\bar{\kappa} \rightarrow \infty$. In particular, if $\bar{\kappa} = \infty$ as in the empirical application, we have

$$\Pr(C = 1) \in \left[\frac{\tau\gamma}{\bar{\pi}}, \frac{\gamma}{\underline{\pi}} \times \frac{\underline{\pi} + (\underline{\kappa} - 1)\tau(\underline{\pi} - \gamma)}{\underline{\kappa}(\underline{\pi} - \gamma) + \gamma} \right]. \tag{2.3}$$

Proof. To keep algebra transparent, introduce new notation $\tau_c \equiv \Pr(T = 1|C = c)$. Write

$$\gamma = \frac{\Pr(R = 1, T = 1)}{\Pr(T = 1)} = \frac{\rho\tau_1\pi}{\tau} \Rightarrow \tau_1 = \frac{\gamma\tau}{\rho\pi}. \tag{2.4}$$

Substituting

$$\frac{\Pr(T = 1|C = 1)}{1 - \Pr(T = 1|C = 1)} = \kappa \frac{\Pr(T = 1|C = 0)}{1 - \Pr(T = 1|C = 0)} \Rightarrow \tau_0 = \frac{\tau_1}{\tau_1 + \kappa(1 - \tau_1)},$$

into the accounting identity $\tau = \rho\tau_1 + (1 - \rho)\tau_0$ yields

$$\tau = \rho\tau_1 + (1 - \rho)\frac{\tau_1}{\tau_1 + \kappa(1 - \tau_1)}.$$

Substituting for τ_1 from (2.4) yields the following algebra:

$$\begin{aligned} \tau &= \frac{\gamma\tau}{\pi} + (1 - \rho)\frac{\frac{\gamma\tau}{\rho\pi}}{\frac{\gamma\tau}{\rho\pi} + \kappa\left(1 - \frac{\gamma\tau}{\rho\pi}\right)} \\ \iff \pi &= \gamma + (1 - \rho)\frac{\gamma\pi}{\kappa\rho\pi - (\kappa - 1)\gamma\tau} \\ \iff (\kappa\rho\pi - (\kappa - 1)\gamma\tau)(\pi - \gamma) &= (1 - \rho)\gamma\pi \\ \iff \rho &= \frac{\gamma}{\pi} \times \frac{\pi + (\kappa - 1)\tau(\pi - \gamma)}{\kappa(\pi - \gamma) + \gamma}. \end{aligned}$$

By taking derivatives, one can verify that the r.h. fraction in the last expression, and therefore the entire expression, decreases in both π and κ . Bounds follow by evaluating it at $(\pi, \kappa) = (\underline{\pi}, \underline{\kappa})$ respectively $(\pi, \kappa) = (\bar{\pi}, \bar{\kappa})$. \square

The bounds effectively multiply sample prevalence by an adjustment factor that reflects test selectivity. As would be expected, the implied prevalence decreases in selectivity κ and sensitivity π . Note also that the adjustment factor simplifies to $\rho = \gamma/\pi$ at $\kappa = 1$ (no selectivity would mean we estimate prevalence by prevalence in the tested subpopulation), to $\rho \rightarrow \tau\gamma/\pi$ as $\kappa \rightarrow \infty$ (perfect targeting means we impute zero prevalence in the untested population; compare (2.3)) and also, for the record, $\rho \rightarrow 1 - \tau + \tau\gamma/\pi$ as $\kappa \rightarrow 0$ (perfectly wrong targeting means we impute complete prevalence in the untested population).

REMARK 2.3. While I present their cumulative impact, Assumptions 2.1 and 2.2 can be disentangled: the first one restricts the relation between test yield and prevalence in the tested population, the second one restricts prevalence across tested and untested populations. Readers are encouraged to ‘pick and choose’ and, of course, to propose other approaches. For example, sensitivity adjustment could be combined with MM’s suggestion to restrict the rate of asymptomatic infections.

2.3. Bounds on the negative predictive value

The negative predictive value $NPV = \Pr(C = 0|R = 0, T = 1)$ is the probability that a negative test result is accurate. It is of great importance in medical decision making (Eng and Bluemke, 2020; Manski, 2021; Watson et al., 2020). It can be bounded as follows:

PROPOSITION 2.3. *Suppose Assumption 2.1 holds. Then sharp bounds on the NPV are given by*

$$NPV \in \left[\frac{1 - \gamma/\bar{\pi}}{1 - \gamma}, \frac{1 - \gamma/\underline{\pi}}{1 - \gamma} \right].$$

Proof. Also for later use, denote the NPV as $\eta := \Pr(C = 0|R = 0, T = 1)$, then

$$\begin{aligned} \eta &= \frac{\Pr(C = 0, R = 0|T = 1)}{\Pr(C = 0, R = 0|T = 1) + \Pr(C = 1, R = 0|T = 1)} \\ &= \frac{1 - \Pr(C = 1|T = 1)}{1 - \Pr(C = 1|T = 1) + (1 - \pi)\Pr(C = 1|T = 1)} \\ &= \frac{1 - \gamma/\pi}{1 - \gamma/\pi + (1 - \pi)\gamma/\pi} = \frac{1 - \gamma/\pi}{1 - \gamma}, \end{aligned}$$

where $\gamma = \pi \Pr(C = 1|T = 1)$ was used. This obviously decreases in π . \square

The expression derived in the proof has a straightforward intuition: the numerator is the fraction of true negatives, i.e., adjusting yield by sensitivity, whereas the denominator is the proportion of measured negatives, in the tested population. The result could again be generalized to allow for imperfect specificity. In that case, there would also be nondegenerate bounds on the positive predictive value $\Pr(C = 1|R = 1, T = 1)$, which equals 1 under the current assumptions.

2.4. Comparison to bounds that start from NPV

Assumption 2.1 contrasts with MM's strategy of inputting ex ante bounds on the NPV. In each case, bounds on the respective other quantity become an output of the model, so the direction of logical inference is reversed. Notating their input bounds as $\eta \in [\underline{\eta}, \bar{\eta}]$, MM establish that

$$\Pr(C = 1) \in [\tau(\gamma + (1 - \gamma)\underline{\eta}), \gamma + (1 - \gamma)\bar{\eta}]. \quad (2.5)$$

I will add to this the observation that in conjunction with empirical test yield, prior bounds on NPV restrict test sensitivity. Specifically, simple algebra building on Proposition 2.3 yields

$$\pi \in \left[\frac{\gamma}{1 - (1 - \gamma)\bar{\eta}}, \frac{\gamma}{1 - (1 - \gamma)\underline{\eta}} \right]. \quad (2.6)$$

The following are some methodological considerations as to why one might want to rather start from sensitivity and selectivity.

First, by bounding the NPV, one necessarily directly restricts prevalence in the tested population. This is because $\Pr(C = 1|T = 1) = \gamma + (1 - \gamma)(1 - \eta)$, so the lower and upper bound on $\Pr(C = 1|T = 1)$ necessarily exceed the corresponding bound on $(1 - \eta)$. Since also the upper bound on overall prevalence is just the upper bound on prevalence in the tested population, the effect can be large.

In addition, (2.6) reveals that bounds on NPV do (in conjunction with test yield, which is observable) imply bounds on sensitivity, but, because these are not made explicit, an opportunity to check empirical plausibility of assumptions is missed.

To see how all of this can play out, consider MM's prior bounds of [0.6, 0.9] on the NPV. With this input, no data can move the upper bound on prevalence below 0.4 and no data—including a test yield of 0—can reduce the lower bound to 0. For example, if test yield is 0.1, then prevalence in the tested population is bounded by [0.19, 0.46]; the upper bound of 0.46 also applies to overall prevalence; and test sensitivity is restricted to [0.22, 0.53], far below any plausible range of values. This example is stark but not hypothetical; compare the first entry in table 2 in MM,

replicated in the first line of Table 1. About half of the upper bounds in that table are below 0.5, so it is important to understand that they cannot be below 0.4 by construction.²

In contrast, prior bounds on sensitivity do not directly restrict prevalence. This does not mean that plausible bounds on the former are completely independent of the latter; see the discussion after Assumption 2.1. However, not directly restricting prevalence matters: because any value of NPV strictly below 1 bounds prevalence away from 0 by assumption, in the early stages of a pandemic, it may inject ‘incredible certitude’ (Manski, 2011) into the analysis. The above, implied bounds on sensitivity are arguably a case in point.

Next, inputting sensitivity (and possibly specificity) generates an interface with the literature on diagnostic tests because it is the focus of this literature. For a general example, see table 1 in Paules and Subbarao (2017). With regard to COVID-19, practitioners’ guides (Eng and Bluemke, 2020; Watson et al., 2020) emphasize the importance of NPV for decisions but treat sensitivity and specificity as scientific input and NPV as jointly determined by those and prevalence. The literature on diagnostic testing (Arevalo-Rodriguez et al., 2020; Yang et al., 2020) is explicitly about sensitivity.

MM seem to disagree when they write: ‘Medical experts have been cited as believing that the rate of false-negative test findings is at least 0.3. However, it is not clear whether they have in mind one minus the NPV or one minus test sensitivity.’ As the technical definition of false-negative rate is not in doubt, the concern must be about informal usage and ultimately base-rate neglect. While the latter is doubtlessly empirically relevant, I have not encountered it in the academic literature on COVID-19.³

Finally, asserting bounds on NPV without taking targeting of tests into account may ignore constraining information that could lead to tighter bounds. Specifically, relatively low values of the NPV (i.e., a large fraction of negative test results being false) will be more plausible if one believes the test to be efficiently targeted. But in that same case one would conclude that the constraint $\Pr(C = 1|T = 1) \geq \Pr(C = 1|T = 0)$ is far from binding. Therefore, the degree of targeting informally enters NPV-based bounds twice, in different directions, but derivation of (2.5) does not force the value to be the same in both appearances. Assumption 2.2 is intended to allow for targeting of tests to affect bounds in a disciplined manner.

3. EMPIRICAL APPLICATION

These bounds are mainly designed to process the information that is available early in a pandemic. For this reason and to highlight some important differences, I illustrate them on MM’s data, i.e., daily counts of tests, test results and fatalities for Illinois, New York and Italy in March and April, and extend them only to early analysis of subsequent hot spots. For credible application to prevalence data at a much later stage of a pandemic, one would have to take into account multiple testing and other factors.

The leftmost columns of Tables 1–2 present bounds that set $\kappa \in [1, \infty)$; that is, they only restrict the direction of selectivity. I also restrict sensitivity to be in $[0.7, 0.95]$. This sensitivity interval

² That bounds on NPV presuppose bounds on prevalence is clearly expressed in Manski (2021), who reverses the direction of logical inference and bounds the NPV of serological tests by inputting MM’s prevalence bounds, which, in turn, inputted assumed (albeit for PCR tests) NPV bounds.

³ The footnote accompanying the cited sentence links to a news piece that attributes an estimated false-negative rate of 0.3 to Yang et al. (2020). While the news piece has vague language, Yang et al. (2020) unambiguously estimate one minus sensitivity.

Table 1. Bounds from Propositions 2.2 ($\bar{\kappa} = \infty$) and 2.3 assuming that sensitivity is bounded by [0.7, 0.95]; bounds on prevalence and sensitivity assuming that negative predictive value is in [0.6, 0.9].

	Date	New bounds ($\bar{\kappa} = \infty$)		NPV-based bounds	
		Prevalence	NPV	Prevalence	Sensitivity
Illinois	3/16	[0.000, 0.131]	[0.957, 0.995]	[0.000, 0.455]	[0.202, 0.503]
	3/23	[0.000, 0.186]	[0.936, 0.992]	[0.000, 0.478]	[0.272, 0.599]
	3/30	[0.000, 0.237]	[0.915, 0.990]	[0.000, 0.500]	[0.332, 0.666]
	4/06	[0.001, 0.279]	[0.896, 0.987]	[0.001, 0.517]	[0.377, 0.708]
	4/13	[0.002, 0.297]	[0.887, 0.986]	[0.002, 0.525]	[0.396, 0.724]
	4/20	[0.003, 0.303]	[0.885, 0.986]	[0.003, 0.527]	[0.402, 0.729]
	4/24	[0.003, 0.299]	[0.887, 0.986]	[0.004, 0.525]	[0.398, 0.725]
New York	3/16	[0.000, 0.191]	[0.934, 0.992]	[0.000, 0.480]	[0.279, 0.607]
	3/23	[0.001, 0.400]	[0.833, 0.980]	[0.002, 0.568]	[0.493, 0.795]
	3/30	[0.004, 0.527]	[0.749, 0.969]	[0.005, 0.621]	[0.594, 0.854]
	4/06	[0.007, 0.583]	[0.705, 0.964]	[0.008, 0.645]	[0.633, 0.873]
	4/13	[0.011, 0.579]	[0.708, 0.964]	[0.012, 0.643]	[0.630, 0.872]
	4/20	[0.013, 0.554]	[0.728, 0.967]	[0.015, 0.633]	[0.613, 0.864]
	4/24	[0.015, 0.519]	[0.756, 0.970]	[0.017, 0.618]	[0.588, 0.851]
Italy	3/16	[0.000, 0.290]	[0.891, 0.987]	[0.001, 0.522]	[0.389, 0.718]
	3/23	[0.001, 0.331]	[0.871, 0.984]	[0.002, 0.539]	[0.430, 0.751]
	3/30	[0.002, 0.304]	[0.884, 0.986]	[0.002, 0.528]	[0.404, 0.730]
	4/06	[0.002, 0.263]	[0.903, 0.988]	[0.003, 0.510]	[0.361, 0.693]
	4/13	[0.003, 0.217]	[0.923, 0.991]	[0.004, 0.491]	[0.309, 0.642]
	4/20	[0.003, 0.186]	[0.936, 0.992]	[0.005, 0.478]	[0.272, 0.599]
	4/24	[0.003, 0.169]	[0.943, 0.993]	[0.006, 0.471]	[0.251, 0.572]

was used by Frazier et al. (2020) in the analysis on which Cornell University's autumn reopening plans were based and was supported by the medical literature at the time.⁴ For comparison, the third column of Table 1 presents bounds that restrict NPV to be in [0.6, 0.9] as well as the direction of selectivity. These replicate MM's table 2.⁵ Table 2 extends these same bounds to new data.⁶ To facilitate further comparisons, the tables also illustrate the bounds on NPV implied by bounds on sensitivity (second column, corresponding to Proposition 2.3) and the converse bounds if one starts from restricting NPV (last column, corresponding to (2.6)).

All bounds are rather wide and, with hindsight, likely include values far above true prevalence.⁷ This reflects the limited information used. However, the bounds differ, and the differences matter.

⁴ UCSF (2020) base medical advice on a point estimate of 0.8. Watson et al. (2020) give 0.7 as 'lower end of current estimates from systematic reviews'. Frazier et al. (2020) use a preferred point estimate of 0.9. Recall in particular that the data analysed here were overwhelmingly generated by testing of symptomatic subjects.

⁵ MM's results were independently replicated from their original data. MATLAB code generating all tables is available from the author. To keep the presentation succinct, the tables only show one day per week of data and the last day.

MM refine the bounds by imposing time monotonicity; that is, prevalence (and therefore both bounds on it) cannot decrease over time. I agree with that restriction; it is dropped here solely because those tables have many identical rows.

⁶ The reader should keep in mind that MM might not have asserted the same bounds on NPV there, so the bounds may not be what they would have proposed. On the other hand, it is a feature that one can (in this author's view) comfortably impose the same bounds on π across these contexts.

⁷ The most extreme case may be Illinois, where recent seroanalysis of a representative sample (Kalish et al., 2021) suggests that true prevalence was on the order of 1%.

Table 2. Extension of Table 1 to US pandemic hot spots later in 2020.

	Date	New bounds ($\bar{\kappa} = \infty$)		NPV-based bounds	
		Prevalence	NPV	Prevalence	Sensitivity
Arizona	8/13	[0.028, 0.258]	[0.906, 0.988]	[0.038, 0.508]	[0.355, 0.687]
California	8/13	[0.016, 0.090]	[0.971, 0.996]	[0.037, 0.438]	[0.143, 0.401]
Florida	8/13	[0.027, 0.193]	[0.933, 0.992]	[0.043, 0.481]	[0.281, 0.610]
Texas	8/13	[0.019, 0.173]	[0.941, 0.993]	[0.031, 0.473]	[0.257, 0.580]

The new upper bounds are considerably more restrictive; for the first day of data in Illinois, the comparison is between 13% and 46%. The lower bounds are less restrictive, and, while this effect is less pronounced, it is not negligible in relative terms; for example, for Illinois on 3/23, the rounding error obscures that the NPV-based lower bound is 1.5 times the sensitivity-based one. In sum, all bounds move down.

Why does this matter? First, the new bounds would rather clearly have ruled out speculation of herd immunity being ‘around the corner’ in March. Second, consider the implied bounds on the infection fatality rate (i.e., fatalities divided by true infections; IFR henceforth). The most informative NPV-based lower bounds (i.e., evaluated on 4/24) equal 0.0003 for Illinois, 0.0013 for New York and 0.0010 for Italy. This is close to ‘flu-like’ numbers, and one might have concluded that credible partial identification analysis did not exclude those. Yet it did: for the same data, the bounds from Proposition 2.3 are 0.0005, 0.0016 and 0.0026; for Italy, the lower bound is above 0.001 starting on 3/29. In places where the data admittedly spoke very loudly, these numbers would have cast strong doubts on ‘just the flu’ conjectures in real time.

Tighter bounds are an unambiguous improvement only if assumptions did not become less credible. I would argue that credibility might, if anything, have improved. A symptom of this is that the NPV-based bounds frequently imply sensitivity below 0.7, whereas the sensitivity-based bounds imply NPV mostly close to, and frequently above, 0.9. The former number seems out of step with expert opinion, including at the time, whereas the latter one would not have raised any eyebrows.⁸ Also, the NPV-based bounds for New York fail to overlap, meaning that test sensitivity must have increased. This might be the case, but forcing it by assumption is arguably against the spirit of weak and credible partial identification assumptions. As discussed in Subsection 2.4, these implausibilities could be avoided by relaxing prior bounds on NPV to [0.6, 1]. However, the upper bounds would then just be worst-case bounds, all bounds would (in the present data) be wider than the new ones, and the intriguing feature that MM can exclude the crude case fatality rate (i.e., observed fatalities divided by observed cases) as true IFR would be lost.

Table 2 repeats the exercise for data from subsequent hot spots of the pandemic.⁹ I deliberately restrict attention to states with high test yield because it seems that MM calibrated their input bounds to such places, and also to states that were in their first wave because that is what the

⁸ For the time period in question, Sacks et al. (2020) provide an empirically informed NPV estimate for Indiana of 0.995. This comes with caveats: it almost certainly corresponds to lower prevalence than in the data considered here, so that MM might have inputted different NPV bounds; also, it operationalizes NPV as test-retest validity. UCSF (2020) gives NPV as 0.972 for symptomatic and 0.998 for asymptomatic cases in the Bay Area, although using the sort of point-identifying assumptions that we seek to avoid here.

⁹ Test counts and results were retrieved from the COVID tracking project. State populations are U.S. Census estimates for 7/1/19.

Table 3. Change in upper bounds from Table 1 as test selectivity is increasingly restricted.

	Date	Lower bound	Upper bound with				
			$\kappa \geq 5$	$\kappa \geq 3$	$\kappa \geq 2$	$\kappa \geq 1.5$	$\kappa \geq 1$
Illinois	3/16	0.000	0.029	0.048	0.0700	0.092	0.131
	3/23	0.000	0.044	0.071	0.102	0.132	0.186
	3/30	0.000	0.059	0.094	0.135	0.172	0.237
	4/06	0.001	0.073	0.115	0.162	0.205	0.279
	4/13	0.002	0.080	0.125	0.175	0.220	0.297
	4/20	0.003	0.083	0.129	0.180	0.226	0.303
	4/24	0.003	0.082	0.127	0.177	0.222	0.299
New York	3/16	0.000	0.045	0.073	0.106	0.136	0.191
	3/23	0.001	0.119	0.183	0.251	0.308	0.400
	3/30	0.004	0.186	0.274	0.360	0.427	0.527
	4/06	0.007	0.225	0.322	0.414	0.484	0.583
	4/13	0.011	0.225	0.321	0.411	0.480	0.579
	4/20	0.013	0.211	0.302	0.389	0.457	0.554
	4/24	0.015	0.191	0.274	0.357	0.422	0.519
Italy	3/16	0.000	0.076	0.120	0.170	0.214	0.290
	3/23	0.001	0.091	0.143	0.199	0.249	0.331
	3/30	0.002	0.082	0.129	0.180	0.226	0.304
	4/06	0.002	0.069	0.108	0.153	0.193	0.263
	4/13	0.003	0.055	0.087	0.123	0.157	0.217
	4/20	0.003	0.047	0.073	0.104	0.133	0.186
	4/24	0.003	0.043	0.066	0.094	0.120	0.169

bounds are designed for. NPV-based bounds continue to allow for very high prevalence but also force test sensitivity to be relatively low. Sensitivity-based upper bounds are again at most half—often much less—than their NPV-based counterparts, and other implications of respective bounds are roughly as before.

Table 3 shows the effect of increasingly restricting test selectivity through Assumption 2.2. Reading it from right to left (meant to evoke ‘outside in’), it starts with $\kappa = 1$, i.e., test monotonicity, and progresses through arguably weak restrictions up to $\kappa = 5$, which is restrictive and may be more in the spirit of a sensitivity parameter. Upper bounds respond strongly. This is reflected in the implied lower bounds on the IFR; for Italy, these increase to (in order) 0.0036, 0.0046, 0.0065 and 0.0101.¹⁰ The same exercise, but for current hot spots, is displayed in Table 4.

The lower bounds are driven by the possibility that all true positives got tested. While this paper focuses on upper bounds, one could also use Assumption 2.2 to refine lower bounds away from that scenario. For the record, restricting $\kappa \leq 100$ [$\kappa \leq 10$] would refine the (last period) lower bound on prevalence for Italy from 0.0034 to 0.0047 [0.0170]. The upper bound on IFR would be refined from 0.1282 to 0.0909 [0.0254], a not completely vacuous restriction.

Inference on these bounds has at least two nonstandard aspects: as MM point out, one might think of states and regions as populations of interest rather than samples from meta-populations; at the very least, one might be interested in inference conditionally on the realized population.

¹⁰ Of course, these numbers should not be compared to MM’s table 2; to the contrary, MM reach similar conclusions when restricting the proportion of asymptomatic infections.

Table 4. Change in upper bounds from Table 2 as test selectivity is increasingly restricted.

	Date	Lower bound	Upper bound with				
			$\kappa \geq 5$	$\kappa \geq 3$	$\kappa \geq 2$	$\kappa \geq 1.5$	$\kappa \geq 1$
Arizona	8/13	0.028	0.093	0.126	0.164	0.198	0.258
California	8/13	0.016	0.036	0.046	0.057	0.068	0.090
Florida	8/13	0.027	0.074	0.097	0.123	0.148	0.193
Texas	8/13	0.019	0.060	0.081	0.106	0.130	0.173

In that sense, conventional sampling theory might not apply. However, whether a given subject is tested, and the result of that test, are realizations of well-defined random variables, opening a clear avenue for statistical inference. Separately, such inference might be complicated if the variables interact (e.g., the marginal tested subject is asymptomatic) and will also involve small probabilities, so that Central Limit Theorem-based approximations (including many forms of the bootstrap) would not apply. Questions like this inform an exciting strand of current research (Rothe, 2020; Toulis, 2021). However, they are orthogonal to the thrust of this paper and also less salient in the application because sample sizes are so large, and identified intervals so long, that estimation uncertainty is dominated by identification issues.¹¹

4. CONCLUSION

This paper proposes new methods to bound prevalence of a disease from partially identifying data and assumptions. It is to some extent intended as a ‘think piece’ to alert researchers to the possibly fruitful application of partial identification methods. Domain knowledge may inform further, and better, iterations.

Options for refinement and subsequent analysis abound. Users who are comfortable with injecting more prior information may refine bounds by placing priors on unidentified parameters (Bollinger and van Hasselt, 2020). Users who find this paper’s simple input bounds too coarse, but do not want to commit to a prior, could interpolate between these approaches with sets of priors, i.e., in the spirit of Robust Bayesian Analysis (Ruggeri et al., 2005). The analysis can also be used as input for decision recommendations. However, it will typically only partially identify optimal actions; to arrive at a precise decision, one may also have to commit to a specific decision criterion under ambiguity (Manski, 2000; Stoye, 2012).

The conceptual innovation is to think of test accuracy as (unknown, not necessarily constant, and possibly not even identifiable) technological parameter and of test selectivity as something that econometric or epidemiological models can speak to. Bounds are therefore constructed with these as starting points, deriving bounds on predictive values by implication and not imposing prior bounds on prevalence in the tested population. In the empirical application, it turns out that some of the more audacious speculations floated at the time contradicted credible partial

¹¹ For a paper-and-pencil computation, approximate the distribution of $(\hat{\tau}, \hat{\gamma})$ as independently normal. Then the distribution of estimated bounds follows easily. Because these bounds are ordered by construction, inference would be a direct application of Imbens and Manski (2004) and in particular Stoye (2009, lemma 3) and would practically amount to intersecting one-sided confidence intervals. Simple calculations show that the standard errors on $(\hat{\gamma}, \hat{\tau})$ would be at least two orders of magnitude smaller than the estimators; consequently, the difference between these confidence intervals and the estimated bounds is at most comparable to the tables’ rounding errors.

identification analysis even then. This illustrates the potential utility of such analysis in early stages of a pandemic.

That said, many of this paper's simplifications are a stretch in the current, more advanced stage of the pandemic. For example, one should distinguish between current and past infection and take multiple testing into account. I leave further exploration of such extensions to future work, but would recommend to use restrictions on test sensitivity as starting points rather than implications. Once again, my main hope is to stimulate further research on partial identification in, or in collaboration with, epidemiology.

ACKNOWLEDGEMENTS

I am indebted to Chuck Manski and Francesca Molinari for stimulating discussions and for sharing their data, to Dan Sacks, Coady Wing and Gabriel Ziegler for feedback and literature pointers, to the Editor and two anonymous referees for extremely helpful feedback, and to Cynthia Stoye for special support in crazy times.

REFERENCES

- Arevalo-Rodriguez, I., D. Buitrago-Garcia, D. Simancas-Racines, P. Zambrano-Achig, R. del Campo, A. Ciapponi, O. Sued, L. Martinez-Garcia, A. Rutjes, N. Low, J. A. Perez-Molina and J. Zamora (2020). False-negative results of initial RT-PCR assays for COVID-19: a systematic review. *PLoS One* 15, e0242958.
- Bollinger, C. R. and M. van Hasselt (2020). Estimating the cumulative rate of SARS-cov-2 infection. *Economics Letters* 197, 109652.
- Canning, D., K. Mahesh, D. Rashmi, C. Muqi and D. Bloom (2020). The association between age, COVID-19 symptoms, and social distancing behavior in the United States. Discussion paper, Harvard T.H. Chan School of Public Health, Cambridge, MA.
- Eng, J. and D. A. Bluemke (2020). Imaging publications in the COVID-19 pandemic: applying new research results to clinical practice. *Radiology* 277, E228–31.
- Frazier, P., S. Henderson, D. Shmoys, J. M. Cashore, N. Duan, A. Janmohamed, J. Wan and Y. Zhang (2020). COVID-19 mathematical modeling for Cornell's fall semester. Discussion paper, Cornell University, Ithaca, New York, USA.
- Imbens, G. W. and C. F. Manski (2004). Confidence intervals for partially identified parameters. *Econometrica* 72, 1845–57.
- Kalish, H., C. Klumpp-Thomas, S. Hunsberger, H. A. Baus, M. P. Fay, N. Siripong, J. Wang, J. Hicks, J. Mehalko, J. Travers, M. Drew, K. Pauly, J. Spathies, T. Ngo, K. M. Adusei, M. Karkanitsa, J. A. Croker, Y. Li, B. I. Graubard, L. Czajkowski, O. Belliveau, C. Chairez, K. Snead, P. Frank, A. Shunmugavel, A. Han, L. T. Giurgea, L. A. Rosas, R. Bean, R. Athota, A. Cervantes-Medina, M. Gouzoulis, B. Heffelfinger, S. Valenti, R. Caldararo, M. M. Kolberg, A. Kelly, R. Simon, S. Shafiq, V. Wall, S. Reed, E. W. Ford, R. Lokwani, J.-P. Denson, S. Messing, S. G. Michael, W. Gillette, R. P. Kimberly, S. E. Reis, M. D. Hall, D. Esposito, M. J. Memoli and K. Sadtler (2021). Undiagnosed SARS-CoV-2 seropositivity during the first 6 months of the COVID-19 pandemic in the United States. *Sci Transl Med.* 13(601), eabh3826.
- Leeflang, M., P. Bossuyt and L. Irwig (2008). Diagnostic test accuracy may vary with prevalence: implications for evidence-based diagnosis. *Journal of Clinical Epidemiology* 62, 5–12.
- Manski, C. F. (1989). Anatomy of the selection problem. *Journal of Human Resources* 24, 343–60.

- Manski, C. F. (2000). Identification problems and decisions under ambiguity: empirical analysis of treatment response and normative analysis of treatment choice. *Journal of Econometrics* 95, 415–42.
- Manski, C. F. (2003). *Partial Identification of Probability Distributions*. New York, NY: Springer.
- Manski, C. F. (2011). Policy analysis with incredible certitude. *The Economic Journal* 121, F261–89.
- Manski, C. F. (2021). Bounding the accuracy of diagnostic tests, with application to COVID-19 antibody tests. *Epidemiology* 32(2), 162–7.
- Manski, C. F. and F. Molinari (2021). Estimating the COVID-19 infection rate: anatomy of an inference problem. *Journal of Econometrics* 220, 181–92.
- Mohammadi, A., E. Esmaeilzadeh, Y. Li, R. J. Bosch and J. Z. Li (2020). SARS-cov-2 detection in different respiratory sites: a systematic review and meta-analysis. *EBioMedicine* 59, 102903.
- Molinari, F. (2020). Microeconometrics with partial identification. In S. N. Durlauf, L. P. Hansen, J. J. Heckman and R. L. Matzkin (Eds.), *Handbook of Econometrics, Volume 7A*, pp. 355–486. Amsterdam: Elsevier.
- Paules, C. and K. Subbarao (2017). Influenza. *The Lancet* 390, 697–708.
- Rosenbaum, P. R. (2002). *Observational Studies*. New York, NY: Springer.
- Rothe, C. (2020). Combining population and study data for inference on event rates. Discussion paper, University of Mannheim, Mannheim, Germany.
- Ruggeri, F., D. Ríos Insua and J. Martín (2005). Robust Bayesian analysis. In D. Dey and C. Rao (Eds.), *Bayesian Thinking, Volume 25 of Handbook of Statistics*, pp. 623–67. Amsterdam: Elsevier.
- Sacks, D. W., N. Menachemi, P. Embi and C. Wing (2020). What can we learn about SARS-cov-2 prevalence from testing and hospital data?. Discussion paper, Indiana University, Bloomington, Indiana, Monroe.
- Stoye, J. (2009). More on confidence intervals for partially identified parameters. *Econometrica* 77, 1299–315.
- Stoye, J. (2012). New perspectives on statistical decisions under ambiguity. *Annual Review of Economics* 4, 257–82.
- Toulis, P. (2021). Estimation of COVID-19 prevalence from serology tests: a partial identification approach. *Journal of Econometrics* 220(1), 193–213.
- UCSF (2020). COVID-19 diagnostic testing. UCSF Health Hospital Epidemiology and Infection Prevention, Discussion paper, UCSF, San Francisco, California, USA.
- Watson, J., P. F. Whiting and J. E. Brush (2020). Interpreting a COVID-19 test result. *BMJ* 369, m1808.
- Yang, Y., M. Yang, C. Shen, F. Wang, J. Yuan, J. Li, M. Zhang, Z. Wang, L. Xing, J. Wei, L. Peng, G. Wong, H. Zheng, M. Liao, K. Feng, J. Li, Q. Yang, J. Zhao, Z. Zhang, L. Liu and Y. Liu (2020). Evaluating the accuracy of different respiratory specimens in the laboratory diagnosis and monitoring the viral shedding of 2019-ncov infections. Discussion paper, Third People's Hospital, Shenzhen, China.
- Zhang, Z., Q. Bi, S. Fang, L. Wei, X. Wang, J. He, Y. Wu, X. Liu, W. Gao, R. Zhang, W. Gong, Q. Su, A. S. Azman, J. Lessler and X. Zou (2021). Insight into the practical performance of RT-PCR testing for SARS-cov-2 using serological data: a cohort study. *The Lancet Microbe* 2(2), E79–87.
- Zhou, X., N. Obuchowski and D. McClish (2002). *Statistical Methods in Diagnostic Medicine, Second Edition*. Hoboken, NJ: John Wiley & Sons.

SUPPORTING INFORMATION

Additional Supporting Information may be found in the online version of this article at the publisher's website:

Replication Package

Co-editor Victor Chernozhukov handled this manuscript.