

Technical Considerations in the Use of the E-value

Tyler J. VanderWeele*

Peng Ding[†]

Maya Mathur[‡]

*Harvard University, tvanderw@hsph.harvard.edu

[†]University of California - Berkeley, pengdingpku@berkeley.edu

[‡]Harvard University, maya.z.mathur@gmail.com

This working paper is hosted by The Berkeley Electronic Press (bepress) and may not be commercially reproduced without the permission of the copyright holder.

<https://biostats.bepress.com/harvardbiostat/paper215>

Copyright ©2018 by the authors.

Technical Considerations in the Use of the E-value

Tyler J. VanderWeele, Peng Ding, and Maya Mathur

Abstract

The E-value is defined as the minimum strength of association on the risk ratio scale that an unmeasured confounder would have to have with both the exposure and the outcome, conditional on the measured covariates, to explain away the observed exposure-outcome association. We have elsewhere proposed that the reporting of E-values for estimates and for the limit of the confidence interval closest to the null become routine whenever causal effects are of interest. A number of questions have arisen about the use of E-value including questions concerning the interpretation of the relevant confounding association parameters, the nature of the transformation from the risk ratio scale to the E-value scale, inference for and using E-values, and the relation to Rosenbaum's notion of design sensitivity. Here we bring these various questions together and provide responses that we hope will assist in the interpretation of E-values and will further encourage their use.

Technical Considerations in the Use of the E-value

Abstract. The E-value is defined as the minimum strength of association on the risk ratio scale that an unmeasured confounder would have to have with both the exposure and the outcome, conditional on the measured covariates, to explain away the observed exposure-outcome association. We have elsewhere proposed that the reporting of E-values for estimates and for the limit of the confidence interval closest to the null become routine whenever causal effects are of interest. A number of questions have arisen about the use of E-value including questions concerning the interpretation of the relevant confounding association parameters, the nature of the transformation from the risk ratio scale to the E-value scale, inference for and using E-values, and the relation to Rosenbaum's notion of design sensitivity. Here we bring these various questions together and provide responses that we hope will assist in the interpretation of E-values and will further encourage their use.

Key Words: Bias Analysis, Causal Inference, Covariate Adjustment, Design Sensitivity, Sensitivity Analysis, Treatment Effects, Unmeasured Confounding

Introduction

In 2017, we introduced the E-value metric to help assess sensitivity of results to potential unmeasured confounding.¹ The E-value was defined as the minimum strength of association on the risk ratio scale that an unmeasured confounder would have to have with both the exposure and the outcome, conditional on the measured covariates, to explain away the observed exposure-outcome association.¹ Formulas for computing E-values or approximate E-values in a variety of settings were provided. Software and also an online calculator for E-values have since been provided.² Since its introduction a number of, often more technical, questions have been posed concerning the use and interpretation of E-values. The purpose of this paper is to document and address some of the more common questions that have arisen.

Calculation of E-values and Interpretation of the Parameters

The formal derivation of the E-value relies on two parameters.³ Let E denote an exposure of interest, D the outcome, C the measured covariates, and U one or more unmeasured confounders. The observed exposure-outcome association on the risk ratio scale, conditional on covariates C, is given by

$$RR_{obs} = \frac{P(D = 1|E = 1, c)}{P(D = 1|E = 0, c)}$$

The association, conditional on C, but adjusted also for U would be:

$$RR_{true} = \frac{\sum_u P(D = 1|E = 1, c, u)P(u|c)}{\sum_u P(D = 1|E = 0, c, u)P(u|c)}$$

If covariates (C,U) suffice to control for confounding of the effect of E on D, then the latter expression RR_{true} can be interpreted as the causal risk ratio of U on D conditional on C. More formally, let D_e denote the counterfactual outcome if E is set to e, and let $X \perp\!\!\!\perp Y \mid Z$ be used to denote that X is independent of Y given Z. The effect of E on D is said to be unconfounded conditional on Z if $D_e \perp\!\!\!\perp E \mid Z$ for all e. We have that if the effect of E on D is unconfounded given (C,U) then

$$\frac{P(D_1|c)}{P(D_0|c)} = \frac{\sum_u P(D = 1|E = 1, c, u)P(u|c)}{\sum_u P(D = 1|E = 0, c, u)P(u|c)}$$

and hence we denote the expression $\frac{\sum_u P(D=1|E=1,c,u)P(u|c)}{\sum_u P(D=1|E=0,c,u)P(u|c)}$ as RR_{true} .

Consider now the following to sensitivity analysis parameters^{1,3}:

$$RR_{UD} = \max \left\{ \frac{\max_u P(D = 1|E = 1, c, u)}{\min_u P(D = 1|E = 1, c, u)}, \frac{\max_u P(D = 1|E = 0, c, u)}{\min_u P(D = 1|E = 0, c, u)} \right\}$$

$$RR_{EU} = \max_u \frac{P(U = u|E = 1, c)}{P(U = u|E = 0, c)}$$

Essentially, RR_{UD} is the maximum effect that U can have on D, conditional on C=c, comparing any two categories of U, for either the exposed or unexposed; and RR_{EU} is the maximum risk ratio relating the exposure to any particular level of U, conditional on C=c. We showed that³:

$$\frac{RR_{obs}}{RR_{true}} \leq \frac{RR_{UD} \times RR_{EU}}{RR_{UD} + RR_{EU} - 1}$$

so that $\frac{RR_{UD} \times RR_{EU}}{RR_{UD} + RR_{EU} - 1}$ was the maximum bias (comparing the ratio of the observed association adjusted for C, to the true association adjusted also for U) that could be generated by such an unmeasured confounder. We then further derived that for the unmeasured confounder(s) to shift the observed risk ratio to the null of 1, if one wanted both RR_{UD} and RR_{EU} to be as small as possible, then the minimum they could both be (which was what we called the E-value) was^{1,3}:

$$E - value = RR_{obs} + \sqrt{RR_{obs}(RR_{obs} - 1)}$$

The E-value is thus straightforward to calculate from the observed risk ratio.

Questions have come up concerning the consequences of having a potentially continuous, many-valued, or multivariate unmeasured confounder U. In such cases, because many pairwise comparisons of categories of U are possible, it may be more plausible than it is with a binary U that the maxima of these numerous pairwise comparisons produce RR_{UD} and RR_{EU} exceeding the

E-value. Hence, it might the case that a large E-value still in fact does not imply all that much evidence for a causal effect. This is a reasonable inquiry and concern. Several interpretative points here, however, are important. First, the confounding associations RR_{UD} and RR_{EU} are both conditional on the measured covariates C so that the confounding associations RR_{UD} and RR_{EU} reflect residual confounding not captured by the measured covariates C . It is the association between U and both D and E , *independent of C* , that is relevant here. Second, the inequality holds for any U and thus the results are relevant for any set of covariates U such that the effect of E on D is unconfounded conditional on (C,U) . One could thus define the parameters RR_{UD} and RR_{EU} for each possible U such that (C,U) suffice to control for confounding and then take the minimum over U of the resulting bias $\frac{RR_{UD} \times RR_{EU}}{RR_{UD} + RR_{EU} - 1}$. One would then have

$$\frac{RR_{obs}}{RR_{true}} \leq \min_{U: D_e \parallel E | (C,U)} \left\{ \frac{RR_{UD} \times RR_{EU}}{RR_{UD} + RR_{EU} - 1} \right\}$$

The E-value calculated as $RR_{obs} + \sqrt{RR_{obs}(RR_{obs} - 1)}$ is then the minimum strength of association on the risk ratio scale that *any and every* unmeasured confounder, that suffices along with C to control for confounding, would have to have with both the exposure and the outcome, above and beyond the measured covariates, to explain away the observed exposure-outcome association.

Third, and perhaps most importantly when combined with the second observation above, the reality of our estimates and attempts at confounding control are at best approximate. Often we would be content, and indeed very pleased, if our estimates were only a few percent away from the truth. Let S denote the set of all possible covariates U such that adjustment for (C,U) would bring the observed association between E and D , conditional on C and adjusted for U , within a factor of say 1.03 (i.e. 3%) of the actual causal effect i.e.

$$S = \left\{ U: \frac{1}{1.03} \leq \frac{\frac{\sum_u P(D = 1 | E = 1, c, u) P(u|c)}{\sum_u P(D = 1 | E = 0, c, u) P(u|c)}}{\frac{P(D_1|c)}{P(D_0|c)}} \leq 1.03 \right\}$$

One could then define the parameters RR_{UD} and RR_{EU} for each possible U in the set S such that (C,U) would suffice to bring the bias within 3% of the causal effect. One could then further take the minimum over U of the resulting bias $\frac{RR_{UD} \times RR_{EU}}{RR_{UD} + RR_{EU} - 1}$. One would then have

$$\frac{RR_{obs}}{P(D_1|c)/P(D_0|c)} \leq (1.03) \times \min_{U \in S} \left\{ \frac{RR_{UD} \times RR_{EU}}{RR_{UD} + RR_{EU} - 1} \right\}$$

Once we allow for up to 3% bias say, a considerably coarsened version of the relevant unmeasured confounders may well suffice. For such coarsened versions of the unmeasured confounders, the relevant confounding association parameters RR_{UD} and RR_{EU} may be considerably smaller for the coarsened unmeasured confounder than for the original underlying unmeasured confounder. Once we further take the minimum over all possible unmeasured

confounders and all possible coarsenings that would result in at most 3% bias, the confounding association parameters RR_{UD} and RR_{EU} will be yet smaller still.

The E-value calculated as $RR_{obs} + \sqrt{RR_{obs}(RR_{obs} - 1)}$ could then be interpreted as the minimum strength of association on the risk ratio scale that *any and every* unmeasured confounder or coarsening thereof, that suffices along with the observed covariates to bring the observed association within 3% of the true causal effect, would have to have with both the exposure and the outcome, above and beyond the measured covariates, to explain away the observed exposure-outcome association. Thus, even if an unmeasured confounder that completely eliminated bias had very large confounding association parameters RR_{UD} and RR_{EU} , the E-value may arguably still be a useful metric for robustness to unmeasured confounding as it can be applied, in an approximate sense as above, to coarsenings and approximate confounding control as well.

A somewhat related issue that pertains to the definition of the confounding parameters concerns the possibility of multiple unmeasured confounders being needed to eliminate confounding. The bias analysis and E-value calculations above are in fact applicable to the setting of multiple unmeasured confounders.³ The confounding parameters RR_{UD} and RR_{EU} are then simply interpreted respectively as the maximum effect that U can have on D, conditional on $C=c$, comparing any two categories of the entire vector of unmeasured confounders U, for either the exposed or unexposed; and RR_{EU} , is the maximum risk ratio relating the exposure to any particular level of the entire vector U, conditional on $C=c$. In such settings large values of RR_{UD} and RR_{EU} may not be particularly implausible. While an E-value of 5 say, may seem, when considering a single confounder, to require very substantial confounding associations and it is perhaps unlikely a single unmeasured confounder could increase the probability of the outcome by 5-fold above and beyond the measured covariates, an increase of that magnitude may not be as implausible if one is considering a whole group of potential unmeasured confounders. The effect comparing the most favorable values of a set of confounders U to the least favorable values of that set U might plausibly increase the probability of the outcome by 5-fold, perhaps even above and beyond the measured covariates. However if this is indeed so, one should perhaps question whether the data available are in fact adequate to get a reasonable estimate of the causal effect at all. If it is known in advance that there are not just one, but numerous known unmeasured confounders, strongly associated with the outcome and exposure and independent of the measured covariates, then arguably this is not a good study setting in which to attempt to draw conclusions. If it is thought plausible that a 5-fold increase in the probability of the outcome could be generated by the unmeasured confounders above and beyond the measured covariates, then it is perhaps time to leave that study data alone and pursue other more adequate data sources. The initial estimate for the causal effect should of course be adjusted for as many measured confounders as possible. A large E-value is only strong evidence for a true causal effect if the set of measured covariates adjusted for plausibly controls for much of the confounding. Said another way, the design of the study, and the collection of data on measured and known confounders, is essential in whether an estimate is plausible or not. The E-value calculation is useful in assessing quantitatively the potential robustness of results to different magnitudes of unmeasured confounding. But if the magnitude of unmeasured confounding is itself thought to be very substantial, possibly being generated by numerous unmeasured confounders, then it is perhaps best to move on.

Lastly, it is to be remembered that the E-value is conservative insofar as, if the parameters RR_{UD} and RR_{EU} are in fact as large as the E-value, then it is possible to construct scenarios in which an unmeasured confounder U with those parameters would suffice to bring the observed association down to the null.³ However, there are also many other scenarios in which an unmeasured confounder has confounding parameters RR_{UD} and RR_{EU} that are equal to the E-value and yet the unmeasured confounder would not suffice to reduce the observed association to the null. The inequality for the maximum bias $\frac{RR_{obs}}{RR_{true}} \leq \frac{RR_{UD} \times RR_{EU}}{RR_{UD} + RR_{EU} - 1}$ is an inequality, not an equality. The inequality is sharp in that it is always possible to construct a variable U with those confounding associations that attains the bound, but, with an actual unmeasured confounder, the bias will often be less. This is especially the case when, for example, the unmeasured confounder is rare.^{1,3} The E-value essentially assumes that the distribution of U is as least favorable as possible. Indeed when it is known in advance that the unmeasured confounder is rare, this is one scenario in which the E-value calculation is perhaps of less use, and is perhaps to be avoided, as it will, in that setting, be exceedingly conservative.

The E-value as a Transformation of the Estimate and Confidence Interval

In our paper, we recommend reporting the E-value for the estimate and for the limit of the confidence interval closest to the null. The former E-value reports how much unmeasured confounding would be needed to shift the estimate itself (one's best guess given the data) to the null. The latter E-value is perhaps a more adequate measure related to the actual strength of the evidence for an effect, since a large E-value for the limit of the confidence closest to the null suggests that even allowing for uncertainty in the estimation of the observed association, the entire range of plausible values for the estimate are all relatively robust to potential unmeasured confounding. We will return more explicitly to issues of inference for and with the E-value in the following section. However, with regard to our recommended practices of reporting the E-value for the estimate and for the limit of the confidence interval closest to the null, another question that has sometimes arisen concerns the E-value simply being a transformation of the estimate and confidence interval itself and thus not really providing any additional information beyond that estimate and confidence interval.

While it certainly is the case that the E-value for the estimate is just a transformation of the observed risk ratio; and the E-value for the limit of the confidence interval closest to the null is just a transformation of that limit, we still believe the reporting of these metrics is useful for interpretative purposes. The E-value gives the interpretation of the estimate and confidence interval with respect to the minimum strength of confounding associations that would be needed to explain away the estimate. It is a more intuitive transformation to the confounding association scale, and one which we believe makes it easier to evaluate the robustness of results to potential unmeasured confounding. Most people cannot simply compute E-values in their head, nor necessarily have a clear sense as to how much confounding would be needed to explain away an estimate of a given magnitude. While the E-value, simply taken as a number, conveys nothing that is not already there in the estimate itself, we think the reporting of the E-value may assist substantially in the actual practice of science, in interpretation, and in the assessment of the robustness of conclusions.

As an analogy, in many settings, the p-value in fact conveys no additional information beyond the estimate and the confidence interval and can be derived from it.⁴ While the use of the p-value has been at times controversial, it arguably is still a valuable measure of evidence for an association when properly interpreted as a continuous metric (rather than say as being dichotomized at the 0.05 level). While the p-value, as a number, likewise often does not convey any information that is not already there in the confidence interval, it can still be helpful for the practical purposes of trying to understand the strength of the evidence.⁵⁻⁷ Most people cannot simply automatically compute a p-value in their head when given the estimate and confidence interval. The scale on which something is reported does make a difference in trying to understand and interpret, and this is the case with the p-value.^{6,7} As another example, instead of reporting risk ratios, we could report the hundredth root of the risk ratios that were obtained so that a risk ratio of 4 was reported as 1.014 and a risk ratio of 1.6 as 1.0047. As numbers, the information conveyed in these two forms of reporting is exactly the same, but the interpretation of the latter is arguably not very intuitive, nor as useful as the former; and again, most people cannot simply do the conversion in their head.

It is similar with the E-value. The proposed E-value calculations, as numbers, do not provide additional information beyond what is already present in the estimate and limit of the confidence interval closest to the null. However, the transformation of these estimates, carried out by the E-value computation, provides the appropriate scale on which to interpret robustness to confounding. Most people again cannot carry out such computations in their head and will thus have more difficulty in interpreting robustness to potential unmeasured confounding when using the untransformed numbers. What is the E-value for a lower limit of the confidence interval which is 1.12? How much confounding would at the minimum be needed to bring such a risk ratio to the null? Again, without going through the computation it is not entirely easy to see or guess. In this case we obtain an E-value of nearly 1.5.

We believe the E-value computations, if routinely carried out are likely to affect interpretative practices with regard to robustness to unmeasured confounding. Consider two hypothetical estimates of a causal effect from two different studies that have adjusted for similar, and all known, confounders: one study obtains an estimate as $RR=1.18$ (95% CI: 1.04, 1.33) and the other as $RR=1.18$ (95% CI: 1.12, 1.24). In our current set of practices, we believe the evidence for a causal effect in these two studies would be interpreted in a relatively similar manner. Both obtained similar effect sizes; both had confidence intervals somewhat bounded away from the null so that it seemed unlikely that it was simply a matter of “p-hacking” to get the confidence interval just above 1; the p-value in the latter study is smaller, but both are relatively extreme. Current practices for both studies would probably suggest evidence for a causal effect but with the caveat of course that association is not causation and that there may be unmeasured confounding. We would argue, however, that the evidence from these two studies is quite different in strength. The E-value for the confidence interval of the former study is 1.24 and for the latter it is 1.49. While we routinely see risk ratios of 1.24 in the research literature, those of a magnitude of 1.5 are rather rarer, and to have a risk ratio of magnitude 1.5 with both the outcome and the exposure, above and beyond the measured covariates, rarer still. We believe if the E-values for the lower limit of the confidence intervals for these two studies were reported, along with the estimates and confidence intervals themselves, the robustness to potential unmeasured

confounding would be more appropriately evaluated, discussed, and assessed. And this is not simply a matter of also reporting the p-value. Examples have been given of two studies, one with a more extreme p-value, but the other having the more extreme E-value for the confidence interval.¹ The p-value and E-value are assessing different types of evidence: the p-value evidence that there is an association, the E-value evidence that that association if present is causal. Both can, and we believe should, be reported. So while our proposed reporting practices for the E-value are indeed just a transformation of the estimate and the limit of the confidence interval closest to the null, we believe this will prove helpful in interpretation and will improve assessments of robustness.

Inference For and Using E-values

As noted above, we recommend reporting the E-value for the estimate and for the limit of the confidence interval closest to the null. Questions have arisen as to whether it might be good to provide a confidence interval for the E-value itself. Note that our recommendation is to provide an E-value for the limit of the confidence interval closest to the null; it is not to provide a confidence interval for the E-value itself. The distinction is subtle, but important, and concerns the goal of inference. Our perspective is that, in settings in which the E-value may be of use, the goal of inference is the causal effect itself of the exposure on the outcome. The E-value is a tool, not the goal, of inference. The E-value is a tool, a tool to assess the robustness of one's conclusions to potential unmeasured confounding when trying to draw inferences about causal effects. The goal and object of inference does not concern E-value itself, but rather the causal effect.

The distinction between the E-value for the confidence interval versus the confidence interval for the E-value becomes clearer when we think about the type of inferential statements one is able to make in repeated sampling. Suppose one calculated a 95% confidence interval for the E-value for the confounded association. In that case, one could make statements along the lines of “Across repeated samples, at least 95% of the time, the minimum strength of association on the risk ratio scale that an unmeasured confounder would have to have with both the exposure and the outcome, conditional on the measured covariates, to explain away the actual confounded exposure-outcome association will lie in the confidence interval provided.” Such statements may be of some interest but they are statements concerning, over repeated samples, minimum unmeasured confounding associations, rather than statements directly about the causal effect itself. Suppose instead of calculating a confidence interval for the E-value, one alternatively, as we advocate, calculated the E-value for the limit of the confidence interval closest to the null and did this across samples and settings. One could then make statements along the following lines: “Across repeated samples, at least 95% of the time, if the actual confounding parameters RR_{UY} and RR_{EU} are both less than the E-value for the confidence interval that was calculated, then there will be a true effect in the same direction as the observed association.”^a This is a statement

^a Or more generally one could make statements of the form “Across repeated samples, at least 95% of the time, if the actual confounding parameters RR_{UY} and RR_{EU} are such that the bias factor generated by them is less than that given by having RR_{UY} and RR_{EU} equal to the E-value

more directly about the presence of a true causal effect and for this reason we believe that in most settings it is the type of statement that is of interest. It makes the causal effect, not the E-value, the target of inference.

Again, as above, one could in principle obtain a confidence interval for the E-value for the estimate, perhaps by bootstrapping or by the delta method. It is not difficult to derive an asymptotic standard error using the delta method for the E-value of the estimate, when that E-value is computed by $RR_{obs} + \sqrt{RR_{obs}(RR_{obs} - 1)}$. Typically, estimation and inference for risk ratios are carried out using symmetric confidence interval around $\beta = \log(RR_{obs})$. Suppose we have an estimate $\hat{\beta}$ of β with estimated standard error $\hat{\sigma}$, then the E-value for the estimate is $e^{\hat{\beta}} + \sqrt{e^{\hat{\beta}}(e^{\hat{\beta}} - 1)}$ and its standard error is, by the delta method, $\hat{\sigma} \left\{ e^{\hat{\beta}} + \frac{2e^{2\hat{\beta}} - e^{\hat{\beta}}}{2\sqrt{e^{\hat{\beta}}(e^{\hat{\beta}} - 1)}} \right\}$ and from there

one could obtain an asymptotic 95% confidence interval for the E-value as $e^{\hat{\beta}} + \sqrt{e^{\hat{\beta}}(e^{\hat{\beta}} - 1)} \pm 1.96\hat{\sigma} \left\{ e^{\hat{\beta}} + \frac{2e^{2\hat{\beta}} - e^{\hat{\beta}}}{2\sqrt{e^{\hat{\beta}}(e^{\hat{\beta}} - 1)}} \right\}$. However, as above, it is not clear why the E-value itself, rather than the causal effect should be the target of inference.

Relation to Rosenbaum's Design Sensitivity

Questions have also arisen with respect to the relation of the E-value to what Paul Rosenbaum calls design sensitivity.⁸ The two concepts are related but also have a number of important differences. The sensitivity analysis parameter, Γ , in Rosenbaum's design sensitivity is the maximum ratio by which two units with identical covariates C may differ in their odds of receiving the exposure. Under randomization conditional on C, two units with the same covariates would not differ at all in their odds of exposure and thus we would have $\Gamma=1$. If however there were an unmeasured covariate U that affected the odds of exposure, then we may have $\Gamma>1$. For a given population, and a given design, and a proposed method of analysis, the design sensitivity is how large the sensitivity analysis parameter, Γ , would have to be in large samples to change the conclusion. What is similar with design sensitivity and the E-value is that both concern the amount of unmeasured confounding that would be required to alter conclusions or to explain away an observed association as to not being due to a true causal effect of the exposure on the outcome.

However, there are several differences between the design sensitivity and the E-value. First, different associations are used to characterize unmeasured confounding in the two approaches. In Rosenbaum's design sensitivity the strength of the unmeasured confounding relates to how much an unmeasured covariate might increase the odds of exposure. With the E-value, the sensitivity analysis parameters are the associations relating the exposure to the unmeasured confounder, *and*

for the confidence interval that was calculated, then there will be a true effect in the same direction as the observed association"

also relating the unmeasured confounder to the outcome. Rosenbaum's design sensitivity does not make explicit reference to the effect of the unmeasured confounder on the outcome. In further work Rosenbaum and Silber⁹ propose what they call an amplification of the sensitivity analysis that re-expresses the sensitivity analysis parameter Γ in terms of effects of an unmeasured confounder on the exposure and outcome. However, it does so under a particular model for the effect of the confounder on the outcome. In contrast, the sensitivity analysis parameters that are used in the E-value, RR_{UD} and RR_{EU} , do not presuppose a model for the effect of the unmeasured confounder on the outcome, nor for the relation between the exposure and the unmeasured confounder. The sensitivity analysis parameters RR_{UD} and RR_{EU} are defined non-parametrically, as above, using maximums.

A second difference between the approaches is that Rosenbaum's design sensitivity was developed to evaluate the *sharp* null hypothesis of no causal effect for any individual. The E-value can be used to assess the strength of unmeasured confounding that is needed to move the estimate to the null of no *average* causal effect; the E-value can also be used to assess the strength of unmeasured confounding that is needed to move the estimate of the average causal effect to any other value of the causal effect as well, for example to a scientifically meaningful threshold for which a causal effect of lesser magnitude would simply not be of substantive interest.^{1,3}

A third difference between the approaches is that Rosenbaum proposes that the design sensitivity be calculated for large samples, whereas our proposal is that the E-value be calculated for the actual sample. Rosenbaum's design sensitivity is intended to be a property of the design, not the sample size. Using the design sensitivity, one can compare different designs for large sample sizes to determine which designs may be more robust to potential unmeasured confounding. Our proposed approach using E-values is calculated with the actual data and estimates. As above, we propose calculating E-values for both the estimate and for the limit of the confidence interval closest to the null.¹ The E-value for the limit of the confidence interval closest to the null will of course vary across samples and will vary by sample size. There is in principle also an E-value for the actual confounded association between exposure and the outcome conditional on C i.e. the risk ratio one would obtain in an infinite sample size relating the exposure and the outcome, conditional on the measured covariates, but not adjusting for unmeasured covariates U. That E-value for the actual confounded risk in an infinite sample is more closely analogous to Rosenbaum's design sensitivity. It is also what would be the target of inference if one were to calculate a confidence interval for the E-value of the estimated risk ratio. However, as argued in the previous section, this seems of less use in evaluating the actual evidence for a causal effect from a given study than the E-value for the confidence interval itself. Again, as argued above, the target of inference is the causal effect, not the E-value.

A third difference between design sensitivity and the E-value is the scale used. The design sensitivity is defined on an odds ratio scale. The E-value is defined on the risk ratio scale. In principle, this difference is only a matter of mathematical definition of scale. However, in practice, we think it is often an important difference. In practice, odds ratios are not infrequently interpreted, often inadvertently, as risk ratios. When the variable under consideration is rare, odds ratios in fact approximate risk ratios and this is then unproblematic.⁴ However, when the variable for which the odds is being considered is common, then odds ratios can vastly

overestimate risk ratios. In many scenarios, odds ratios are roughly the square of risk ratios.¹⁰ When the probability of the variable being considered lies in the range 0.2 to 0.8, the odds ratio can exaggerate the risk ratio by a factor as large as 400%!¹⁰ In these cases, interpreting odds ratios as risk ratios is highly problematic. The sensitivity analysis parameters in Rosenbaum's design sensitivity are defined in terms of odds ratios for the exposure. The exposures being examined in many studies are of course often relatively common. Sensitivity analysis using odds ratio scales in these settings can be problematic,¹¹ and one must take due caution. If investigators are careful not to interpret odds ratios as risk ratios, then this need not necessarily be problematic. If, however, in the interpretation of the investigators, or that of the readers of their research reports, odds ratios are interpreted as risk ratios then this may vastly exaggerate the robustness of conclusions to unmeasured confounding since odds ratios can vastly exaggerate risk ratios. We believe such misinterpretation of odds ratios as risk ratios is common in practice and for that reason would in general advocate for using sensitivity analysis parameters on risk ratio scales. We have endeavored to provide a variety of approximations so that E-values, with parameters reported on risk ratio scales, can be obtained regardless of the initial method of analysis or effect measure employed in estimation.^{1,10}

In summary, while design sensitivity is somewhat analogous to the E-value for the actual confounded association between the exposure and outcome, the reporting practices for the E-value that we advocate for¹ differ from design sensitivity in their considerations of the relations between the unmeasured confounder and the outcome; differ in considering the null of no average causal effect rather than the sharp null; differ in considering the actual sample versus an infinite sample; and differ in using risk ratio rather than odds ratio scales.

It is our hope, by addressing these questions concerning the interpretation of the confounding association parameters, the nature of the E-value transformation, questions of statistical inference using the E-value, and distinctions from design sensitivity, that the interpretation of the E-value metric is clearer and that its use will thereby be further facilitated.



References

1. VanderWeele, T.J. and Ding, P. (2017). Sensitivity analysis in observational research: introducing the E-value. *Annals of Internal Medicine*, 167:268-274.
2. Mathur, M.B., Ding, P., Riddell, C.A., and VanderWeele, T.J. (2017). Website and R package for computing E-values. *Epidemiology*, in press.
3. Ding, P., and VanderWeele, T.J. (2016). Sensitivity analysis without assumptions. *Epidemiology*, 27;3:368-377.
4. Rothman, K.J., Greenland, S., Lash, T.L. (2008). *Modern Epidemiology*, 3rd Edition. Lippincott Williams & Wilkins.
5. Wasserstein RL, Lazar NL. (2017). The ASA's statement on p-values: context, process, and purpose. *Am Stat*. 2016;70:129-33.
6. Greenland, S. (2017). Invited commentary: The need for cognitive science in methodology. *Am J Epidemiol*, 186: 639-645.
7. VanderWeele, T.J. (2010). Re: The ongoing tyranny of statistical significance testing in biomedical research. *European Journal of Epidemiology*, 25:843.
8. Rosenbaum, P.R. (2004). Design sensitivity in observational studies. *Biometrika*, 91:153-164.
9. Rosenbaum, P.R. and Silber, J.H. (2009). Amplification of sensitivity analysis in observational studies. *Journal of the American Statistical Association*, 104:1398-1405.
10. VanderWeele, T.J. (2017). On a square-root transformation of the odds ratio for a common outcome. *Epidemiology*, 28:e58-e60.
11. Robins JM. (2002). Comment on “Covariance adjustment in randomized experiments and observational studies.”. by Paul R. Rosenbaum. *Statistical Science*, 17(3):286-327.

