Chapter 1

Conflict of Interest and the Principle of Total Evidence

1.1 Introduction

Randomized controlled trials (RCTs) are often treated as the gold standard of medical research (Sackett et al., 1996). The evidence-based-medicine movement holds that RCTs are at the top of the quality-of-evidence hierarchy. According to the "best evidence synthesis" view (Slavin, 1986) only high-quality studies ought to be taken into account. Yet, RCTs have also been severely criticized for being subject to biases (e.g., small size, inadequate blinding) that too often make them less than perfectly reliable (Worrall, 2002; Ioannidis, 2005). In this light, is it legitimate to consider other sources of evidence? In the literature, it has been argued (Bovens and Hartmann, 2003; Osimani and Landes, 2023) that one ought to consider all sources of evidence, once their reliability is also taken into account. This is in line with the so-called *principle of total evidence* (Carnap, 1947).

Meanwhile, it was observed that most medical trials suffer from a *conflict of interest* (henceforth, CoI), such as sponsorship by pharmaceutical companies (Roseman et al., 2011). Similarly to the reliability case, we may regard the decision to ignore studies with CoI as an application of the best evidence view. Contrary to that case, however, it is unclear how one may demonstrate that it is useful to consider sources of evidence subject to CoI, because CoI has an ambiguous influence on reported results. In fact, available reviews suggest both that CoI raises the probability of biased estimates (Friedman and Richter, 2004; Kjaergard

and Als-Nielsen, 2002) and that studies subject to CoI are more reliable in virtue of their better design/quality (Lexchin et al., 2003). Intuitively, the two considerations pull in different directions, namely privileging evidence from RCTs with CoI vs discarding it. So the question arises, would one here, too, benefit from considering all sources of evidence?

Some authors (Stegenga, 2018; Ioannidis, 2008) have recommended that conclusions from studies subject to CoI be down-adjusted, or "discounted". We see a fundamental problem with this proposal, namely that an unqualified discount, which does not take the dual role of CoI into account, would be epistemically unjust. Recently, Jonathan Fuller (2018) has argued that evidence of studies subject to CoI can be appropriately taken into account—in a Bayesian way—rather than ignored. However, he proposes no concrete model to accomplish that.

Here, we provide a Bayesian model, which shows how one ought to take information on CoI into account without committing an epistemic injustice and in agreement with the principle of total evidence. Crucial to this task is that our model facilitates confirmation of true hypotheses and disconfirmation of false ones, precisely in virtue of representing all available information. Accordingly, we show that (A) studies subject to CoI can improve confirmation despite the ambiguous role of CoI, and that (B) information on CoI is not less relevant than other information.

1.2 Bias and the principle of total evidence

Evidence-based medicine is the nowadays very popular view that there's a natural hierarchy in the quality of the evidence, and that one should prioritize evidence at the top and dismiss evidence at the bottom when assessing the support of medical studies for a hypothesis of interest. The strongest take on this view is that only the best available evidence should be considered ("best evidence synthesis", Slavin 1986), namely RCTs, and reviews and meta-analyses thereof.

RCTs are routinely used in the medical community to establish the efficacy of medical interventions with respect to pathologies of interest. Judgments of efficacy are based on outcome measures obtained in an RCT, typically the difference or the ratio between the effects of the intervention in the treatment group and in the control group. The goal of RCT is thus to estimate the causal effect of administering the treatment instead of the control, and ultimately, to warrant clinicians in using that intervention to treat patients suffering from those pathologies.

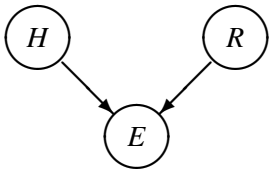

$$H = \begin{cases} h: & \text{Clinically significant benefits.} \\ \bar{h}: & \text{No clinically significant benefits.} \end{cases}$$

$$R = \begin{cases} r: & \text{High reliability of the source.} \\ \bar{r}: & \text{Low reliability of the source.} \end{cases}$$

$$E = \begin{cases} e: & \text{Evidence of significant benefits.} \\ \bar{e}: & \text{No evidence of significant benefits.} \end{cases}$$

Figure 1.1

RCTs have two main virtues. By randomization over the test population, they are meant to eliminate confounding, namely the presence of factors other than the intervention making a difference to the outcome. Thanks to blind assignments of individuals to the treatment and the control (e.g., placebo, alternative treatment) group, they are meant to eliminate biases where beliefs of the experimenters or the subjects may influence the outcome. This design ensures—in principle—a reliable estimate of the treatment's effect size. For this reason, RCTs are often considered the "gold standard" in medical research (Sackett et al., 1996).

In practice, however, RCTs suffer from many limitations. These limitations are widely discussed and have led some to doubt the validity of much contemporary medical research (Ioannidis, 2005). Some limitations are statistical. For instance, small sample size increases the chances of unevenly distributed factors that bias the effect size estimate. Other limitations depend on faults in the experimental design. For instance, without proper blinding, selection bias induces experimenters/subjects to re-allocate. So, the question arises, is it legitimate to consider less-than-ideal evidence, in light of the biases of "best" evidence?

The question has been answered in the affirmative in the philosophical literature: evidence from less-than-ideal sources can improve confirmation if their reliability is properly accounted for (Osimani and Landes, 2023). The simplest scenario to illustrate this claim is a model over three variables, the hypothesis H, the evidence E, and the reliability of the evidential source R. (Figure 1.1). Although the complete model includes a node denoting the reliability of the source, R, if the source is fully reliable, such that the prior of r equals 1, R drops out of the model, such that the confirmation of H (which is a function of its posterior probability relative to its prior) depends entirely on E:

$$P(h|e) = \frac{P(h)P(e|h)}{P(e)} = \frac{P(h)P(e|h)}{P(h)P(e|h) + P(\overline{h})P(e|\overline{h})}$$

By contrast, if the source of the evidence is not fully reliable, namely r has a prior smaller than 1, E has a weaker bearing on H, due to its dependence on R, once the latter is explicitly accounted for:

$$P(h|e) = \frac{P(h)P(r)P(e|hr) + P(h)P(\overline{r})P(e|h\overline{r})}{P(h)P(r)P(e|hr) + P(\overline{h})P(r)P(e|\overline{h}r) + P(h)P(\overline{r})P(e|h\overline{r}) + P(\overline{h})P(\overline{r})P(e|\overline{h}\overline{r})}.$$

Along similar lines, several methodologists advocate consideration of studies of both high- and low- quality when evaluating medical hypotheses. For instance, Welton et al. (2009) and Verde (2021) have claimed that one may benefit from synthesizing evidence from studies randomly drawn from a mixture of studies with low risk of bias (in virtue of, respectively, allocation concealment or randomization) and studies with high risk of bias (for lack of concealment or randomization). Such proposals, like the simple model in Figure 1.1, promote the use of all available evidence over the neglect of part of it, even if it is of a lesser quality.

These considerations are in line with the principle of total evidence (Carnap, 1947), which holds that, when assessing the credibility of hypotheses, we should endeavour to take into account all of the evidence at our disposal instead of just some proper part of it. The rationale behind the principle is that, since truth-conducive scientific inquiry is based on evidence, one ought not ignore any piece of evidence, as this would unnecessarily slow down the inquiry. That holds, note, even if the evidence in question is known to be potentially misleading. In that case, the principle requires that this piece of information—about the reliability of the evidence—be, too, taken into account to ensure maximal truth conduciveness whilst avoiding unjustified conclusions. In what follows, we shall endorse the principle of total evidence and argue that one ought to comply with it when accounting for the evidence, not only when the source is subject to a bias but also when it is subject to CoI.

1.3 Conflict of interest

An important but under-appreciated fact about RCTs is that most of them are subject to CoI (Roseman et al., 2011). CoI may be defined as "anything that may influence professional judgment", be that a financial interest by a pharmaceutical company, a personal interest, an academic interest, a political interest, or else (Kjaergard and Als-Nielsen, 2002, 2).

Assume that sufficiently many studies are present, some with CoI, some without CoI. To understand whether the difference between the two subgroups is significant, one might employ a *meta-regression* as a diagnostic tool.[1] Assume that the effect size may be predicted as a linear function, for instance $f_i(x, y) = \alpha_i + \beta_i x + \gamma y$, of sampled effects x and study characteristics y, as to minimize overall variance. A meta-regression aims to test the statistical significance of additional covariates. In our case, one may treat CoI as a categorical (binary) covariate c in $f_i(x, c) = \alpha_i + \beta_i x + \gamma c$. Conditional on the value of c (i.e., subgroup with or without CoI), the regression line may shift. To check whether the shift is significant, one should test whether the null hypothesis $\gamma = 0$ can be rejected (Fagerland, 2015, 457). As it turns out, financial interest is a *good proxy* for differences induced by CoI (Lundh et al., 2017, 21), while other factors such as the affiliation of the authors play a negligible role. So, the availability of information on financial sponsorship makes it in principle possible to test whether CoI induces a significant difference.

Now, it has been observed that trials that are subject to CoI are 4 times as likely to produce positive outcomes (Lexchin et al., 2003); 1.32 times as likely to report favourable efficacy results, and 1.87 times as likely to report favourable harms (Lundh et al., 2017); more likely to produce favourable conclusions, if not favourable results (Yank et al., 2007). At the same time, this observation alone does not not establish which subgroup of studies is most trustworthy in producing more accurate estimates, since there are several possible explanations for the difference.

The observation that industry-funded research is more likely to report significant beneficial effects is compatible with studies subject to CoI being *more* rather than less accurate. After all, industry-funded studies may have more power and better design, simply because they can exploit larger financial resources. This could fully explain the observed divergence in effects between the two subgroups. Indeed, while some studies report that industry funding is associated with poorer methodology (Jørgensen et al., 2008; Montgomery et al., 2004), other studies report no difference in methodological quality between industry and non-industry funded research or even that industry sponsorship is associated with higher quality (Lexchin et al., 2003).

What is striking, though, and in need of explanation, is the fact that the correlation between CoI and a difference in estimates is *robust*, even

[1] On the difference between meta-regression and simple regression, see Higgins and Green (2011, §9.6.4).

after controlling for a number of potential explanatory factors, such as methodological quality, statistical power, type of intervention or medical specialty (Kjaergard and Als-Nielsen, 2002), sample size, study design, country of primary authors (Friedman and Richter, 2004), etc. This suggests that CoI may be a latent source of bias, even if the nature of the induced bias is hard to pin down exactly. Indeed, Lexchin (2012) indicate that CoI can introduce biases in a subtle way—from selective outcome reporting (e.g., standalone or repeated publication of study with favourable outcome), to poor design (e.g., inappropriate choice of doses, dosing intervals, comparators), inadequate analysis (e.g., "p-hacking"), and fraud (e.g., data fabrication). Typically, one cannot verify whether all of the latter factors are absent from a given study by directly inspecting the report of the study and of the data collected in the study. As a result, popular scores (e.g., Jadad), which are designed to evaluate the quality of the design of RCTs, are insensitive to such biases.

The above considerations should lead one to conclude that the effect of CoI is *ambiguous*—it can both improve quality *and* induce bias. How should one deal with this piece of information?

Ideally, if the sample size of studies without CoI were large enough, one could dispense with ambiguous evidence coming from potentially biased RCTs altogether. This is in line with the view that only the best available evidence should be considered. The problem with this view, however, is that large studies without CoI are often unavailable. For instance, phase III trials, although they may be conducted under the supervision of regulatory bodies, are typically sponsored by the pharmaceutical industry (Reynolds, 2001), which has a clear interest in demonstrating positive efficacy results. Drug approval thus has to rely on only few RCTs, which are normally all subject to CoI.

In light of this observation, some authors have recommended that conclusions from studies subject to CoI be down-adjusted, or "discounted". For instance, Stegenga (2018) claims that "our confidence in medical interventions ought to be low, or at least much lower than is now the case". Ioannidis (2008) advocates a "rational down-adjustment of effect sizes" and the use of "analytical methods that correct for anticipated inflation" (644). Fuller (2018) maintains that, although multiple explanations for the association between CoI and divergent results remain possible, and although it is difficult to quantify the bias, "on net the plausible interpretations compel us to lower our confidence in therapies, at least qualitatively" (778). Yet, how to concretely go about discounting evidence subject to CoI is non-trivial. Given the association between CoI and bias,

some (West et al., 2002; Shea et al., 2007) propose to add CoI to the categories that determine the quality score of a study. The problem with such proposals is that they may commit an *epistemic injustice.* In light of the evidence about the quality of individual studies subject to CoI, it would seem unjust to systematically punish studies with equal or superior design only because they belong to a "suspect" reference class. An unqualified discount of the evidence of studies subject to CoI tantamounts to neglecting its ambiguous role as both a promoter of biases and a preventer of them, and de facto to reducing CoI to a bias.

The lack of a concrete and well-motivated proposal on how, if at all, to discount evidence subject to CoI has important practical implications. For instance, meta-analyses are routinely used to improve effect size estimates by pooling together multiple RCTs and solve potential biases due to their individual (small) sample size. However, as observed by Roseman et al. (2011), almost 70% of the RCTs included in meta-analyses are subject to CoI. Since CoI may decrease their reliability and skew their results, it poses a threat to the validity of the method, such that a prima facie more accurate estimate may hide biases induced by CoI, which actually lead to a less accurate estimate. In spite of this threat, the issue is not explicitly addressed by existing protocols on how to perform a meta-analysis.[2] As a result, current meta-analyses tend to omit any reference to CoI, let alone solve the possible problems due to them.

In sum, from the principle of total evidence, it follows that neglecting evidence from larger and more powerful studies is irrational, and so is neglecting evidence on the relevance of CoI to biased estimates of causal effects. At the same time, without a concrete model, and in light of CoI's ambiguous role, it's unclear whether discounting evidence subject to CoI is *justified* or how evidence from CoI-laden source is *confirmatory*. The rest of the paper is devoted to developing one such model. This, in turn, promises to be a first step towards addressing cognate problems, such as ensuring that meta-analyses actually improve effect size estimates by properly accounting for information on CoI.

[2]For instance, the Cochrane Institute, which promotes the production of meta-analyses, does recommend methods of bias detection such as funnel plots and sensitivity analyses (Higgins and Green 2011, §10.4), but provides no special recommendations on what to do with RCTs that report a CoI.

1.4 Fuller's proposal

A recent proposal by Jonathan Fuller (2018) on how to discount evidence subject to CoI shall serve as a useful starting point for our proposal.

The proposal is based on the distinction between two different types of evidence. There's *first-order evidence*—evidence E that bears directly on the probability of a hypothesis H—and *higher-order evidence*, which is evidence *about the evidence* E. In particular, in the latter category falls evidence E' of how the evidence E was generated. This sort of evidence is the product of "meta-research", which may be characterized as

> [...] an evolving scientific discipline that aims to evaluate and improve research practices. It includes thematic areas of methods, reporting, reproducibility, evaluation, and incentives [...] helping science progress faster by conducting scientific research on research itself. (Ioannidis et al., 2015, 1-2)

In particular, meta-research evidence is "not evidence about a particular agent's reasoning; it concerns the public evidence from which many agents reason" (Fuller, 2018, 774). Based on this distinction, a particular study result is first-order evidence; whereas evidence on how the evidence is generated in the class of studies to which that particular study belongs is higher-order.

Fuller claims that one ought to incorporate meta-research evidence into a probabilistic (Bayesian) judgment on the hypothesis:

> Meta-evidence may be irrelevant to the bearing of E on H, but it is entirely relevant to our confidence in E and thus to our confidence in H. (774)

However, Fuller offers no concrete model to accompany his claim. The next section aims to remedy this deficiency.

From Fuller's proposal, we inherit the attention to meta-research evidence and the view that both first- and higher-order evidence bear on the probability of the hypothesis. At the same time, Fuller invites an equivocation between the "directness" of the evidence and its "order" (2018, 770), which notions we prefer to keep distinct. Moreover, we do not share his intuitions on the appropriate (Bayesian) way to model this scenario. He maintains that E' is relevant to H because E fails to screen off (2018, 776), or render conditionally independent, E' and H. However, this suggests that there is a missing edge between E' and H relative to either $E' \longleftarrow E \longrightarrow H$ or $E' \longrightarrow E \longrightarrow H$, both of which structures

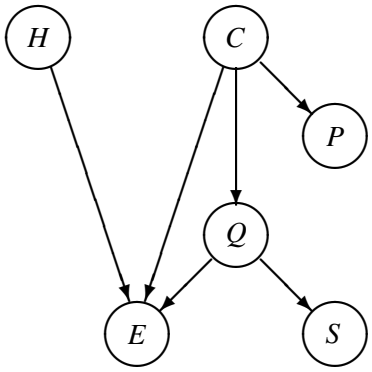

$$H = \begin{cases} h: & \text{Clinically significant benefits.} \\ \bar{h}: & \text{No clinically significant benefits.} \end{cases}$$

$$C = \begin{cases} c: & \text{Conflict of interest.} \\ \bar{c}: & \text{No conflict of interest.} \end{cases}$$

$$P = \begin{cases} p: & \text{For-profit investment.} \\ \bar{p}: & \text{No-profit investment.} \end{cases}$$

$$Q = \begin{cases} q: & \text{High quality.} \\ \bar{q}: & \text{Low quality.} \end{cases}$$

$$S = \begin{cases} s: & \text{High quality score} \\ \bar{s}: & \text{Low quality score.} \end{cases}$$

$$E = \begin{cases} e: & \text{Evidence of significant benefits.} \\ \bar{e}: & \text{No evidence of significant benefits.} \end{cases}$$

Figure 1.2

contradict not only Fuller's statement that there is no direct dependence between meta-research and hypothesis (cf. quote) but also the widely received view that the reliability of the evidence-gathering method bears on the hypothesis via a collider structure, $H \longrightarrow E \longleftarrow E'$ (cf. Figure 1.1, with E' in the place of R). The next section shall introduce a fully fledged Bayesian model that explains how evidence of RCTs subject to CoI may aid confirmation without being subject to the aforementioned problems.

1.5 A Bayesian model

We now introduce a Bayesian model, which we motivate and set up in agreement with the principle of total evidence. Ultimately, the model shall serve to prove two main theses, namely that (A) information on CoI can improve confirmation despite the inherent uncertainty on CoI's opposite effects, and irrespective of the quality of a study, and that (B) it is generally false that CoI has a smaller confirmatory weight than quality, and thus it is unjustified to preferentially neglect CoI over quality.

1.5.1 Variables and dependencies

For simplicity, all of the variables in the model are binary (Figure 1.2). The hypothesis of interest depends on many sources of evidence. The first

obvious difference between these sources is that some are *direct* relative to the hypothesis, H, and others are *indirect*, by which we mean that the former are graphically adjacent to H, and the latter influence H via other nodes. In particular, direct evidence for H are reported effects, E. Indirect evidence for H are conflict of interest C, for-profit investment P, quality Q, and quality scores S, because all of them influence H via E.

A second difference between the sources of evidence is that they are modelled differently: some are modelled as *nodes* (cf. Bovens and Hartmann 2003, ch. 4), whose values are either directly observed, such as the reported effects E, or informed by proxies. For instance, one may use information in the individual studies, scored by quality assessment tools (Stegenga, 2018, ch. 7), S, as a proxy for study validity, or quality, Q. And one may use reports of for-profit sponsorship P as a proxy for CoI (Lundh et al., 2017), C, indicating whether it is in fact present or absent. By contrast, other sources of evidence are conceptualized as *parameters*. Crucially, some such parameters, for instance the likelihood of bias induced by CoI, are constrained by evidence of meta-epidemiological data from appropriate reference classes of studies where as much information as possible is available on factors to which the likelihoods are sensitive (e.g., truth of the hypothesis, quality of the study).

Now we're in a position to better characterize the distinction between first and higher-order evidence for H. First-order evidence for H are facts discovered by observation, viz. E, P, S. Higher-order evidence for H are constraints inferred from meta-research (e.g., meta-regression), such as the (ceteris paribus) residual influence of CoI on the evidence. For instance, S is first-order evidence for H because quality scores aggregate observable characteristics of quality. All criteria in each quality tool are *testable on individual studies independently*. By contrast, the robust dependence of the evidence on CoI is higher-order evidence for H because it denotes a biasing influence of CoI on the evidence via characteristics, such as ghostwriting, fraud, etc., which are unobservable at the individual level, and thus to which the scores are insensitive (Lexchin, 2012), and are *only testable at the population level.* Note that the distinction between first- and higher-order evidence is orthogonal to that between direct and indirect evidence. For instance, C is indirect evidence for H but need *not* be higher-order wrt H.

In this respect, evidence of the residual influence of CoI is analogous to evidence of biases, which may only be assessed at the population level. One example is publication bias (Egger et al., 1997), namely the overestimation of the effect size due to the fact that insignificant results

are less likely to be published or even submitted for publication. The bias becomes visible by the use of so-called funnel plots, which represent the estimates of multiple studies against their size. Without bias, one should expect a symmetric distribution, with estimates of larger studies (with higher precision) being near the average and estimates of smaller studies (with lower precision) being spread evenly on both sides of the distribution. The bias induces an uneven distribution of the estimates of small studies, such that there are more observations on the right side than on the left side. Publication bias, thus, is detectable only at the level of populations of studies.[3] In what sense, then, is CoI analogous to publication bias, and in what sense is it different?

Unlike publication bias, CoI is not a bias. Rather, it is a possible *cause* of biases, some of which are detectable at the individual level and some at the population level. Like evidence of publication bias, however, evidence of CoI is a population-level property, something over and above evidence that may be gathered by looking at individual studies. These features of CoI are directly reflected in our model. In it, CoI is not a bias but a "cause" of biases, which may both promote and hamper confirmation. This is rendered in the model by having C influence E along *two* paths, one where it promotes quality by preventing biases detectable at the individual level, and one where it promotes biases undetectable at the individual level—although they may be detectable at the population level, if suitable proxies are available.[4] Therefore, in our model, evidence subject to CoI may or may not need discounting, depending on higher-order evidence of the strength of CoI on either path, and on first-order evidence, if any, on whether CoI generates biases along the two paths.

1.5.2 Probabilistic constraints

The following constraints, and all of our below results, are relative to the case of Bayesian confirmation of a hypothesis by the evidence from a single study—possibly subject to CoI—in the absence of further evidence. Some of the constraints are quantitative, some are qualitative. Among the former, some are directly informed by meta-research evidence. In

[3] One way to correct for the bias is to re-estimate the effect size not based on the actual, unevenly distributed population, but on the counterfactual, evenly distributed one, to which the (putatively) missing studies have been added.

[4] For simplicity, our model leaves implicit the latter kind of biases as well as any proxy of them. If one were to explicitly model those nodes, they would be graphically analogous to Q and S, only located on, or departing from, the direct edge $C \longrightarrow E$.

the conclusion, we will briefly touch on the issues arising in the case of confirmation by evidence from multiple studies.

$$P(h) \in [0.001, 0.01] \tag{1.1}$$

$$P(c) \in [0.7, 1] \tag{1.2}$$

$$P(p|c) = 0.9 \quad P(p|\bar{c}) = 0.1 \tag{1.3}$$

$$P(q|c) = [0.5, 0.9] \quad P(q|\bar{c}) = 0.5 \tag{1.4}$$

$$P(s|q) = 0.9 \quad P(s|\bar{q}) \in [0.3, 0.7] \tag{1.5}$$

$$\frac{P(e|c)}{P(e|\bar{c})} \in [1.3, 4] \tag{1.6}$$

$$P(e|h) > P(e|\bar{h}) \tag{1.7}$$

$$1 \geq P(e|hqc) > P(e|h\bar{q}c) \quad P(e|hq\bar{c}) > P(e|h\bar{q}\bar{c}) \tag{1.8}$$

$$P(e|\bar{h}qc) < P(e|\bar{h}\bar{q}c) \quad 0 \leq P(e|\bar{h}q\bar{c}) < P(e|\bar{h}\bar{q}\bar{c}) \tag{1.9}$$

$$0 \leq P(e|\bar{h}q\bar{c}) < P(e|\bar{h}qc) \quad P(e|\bar{h}\bar{q}\bar{c}) < P(e|\bar{h}\bar{q}c) \tag{1.10}$$

$$P(e|hq\bar{c}) < P(e|hqc) \leq 1 \quad P(e|h\bar{q}\bar{c}) < P(e|h\bar{q}c) \tag{1.11}$$

Let us examine the above constraints one by one, starting with the quantitative constraints (1.1) to (1.6). (1.1) says that the prior of a (novel) hypothesis being true is low. We assume that major discoveries of treatments with huge benefits have been made already (penicillin, etc.) and still-to-be-made discoveries are fewer and their effect sizes are smaller. Whereas critics such as Ioannidis and Stegenga will be more conservative on the prior of novel findings on clinically relevant benefits (say, 1/1000), the pharmaceutical industry will arguably be less conservative (say, 1/100). (1.2) says the probability of CoI in the overall population of studies is over 70%. This number is suggested by Roseman et al. (2011), after a count of the declarations of CoI in studies included in meta-analyses. This figure, in turn, is arguably representative of the percentage of studies subject to CoI in the overall population of studies, whether or not they are included in meta-analyses. (1.3) says that CoI is very likely to induce a for-profit financial sponsorship, and its absence is unlikely to induce a for-profit financial sponsorship. This is suggested by two considerations: first, Lundh et al. (2017, 21) show that financial sponsorship is a good proxy for differences induced by CoI; and second, although sponsorship can come from both no-profit and for-profit organizations, also at the same time (Kjaergard and Als-Nielsen, 2002, 3), for-profit organizations have arguably more money to invest and more incentive to invest in studies where they can skew the result in their favour. (1.4) reflects that quality is often considered higher (between 30% to 40% higher) in industry-sponsored studies, and sometimes considered

the same as in publicly funded studies (Lexchin et al., 2003, 10). It is seldom considered inferior (Jørgensen et al., 2008; Montgomery et al., 2004). These assessments depend on the quality measure chosen to test for differences, the field of studies considered in the test, etc. By contrast, we assume that there is no distinctive quality expectation if the study is publicly funded. The (scant) research on the topic suggests that publicly funded studies tend to, if anything, be at a higher risk of bias (Gomes and Stavropoulou, 2019). (1.5) says that the likelihood of a high score—by whatever quality measure—given high quality is arguably large, but the likelihood of a high score given low quality tends to be sensitive to the chosen measure, given that each measure is unable to spot at least some internal validity problem, and some measures are worse than others, at least according to Stegenga (2018). Still, it is difficult to extract intervals, let alone reliable ones, from Stegenga's sources, viz. Hartling et al. (2009) and Jüni et al. (1999). Finally, the crucial quantitative piece of information for our subsequent argument is given by (1.6), which says that CoI is expected to render favourable results from 1.3 times to 4 times more likely, depending on the survey (Lexchin et al., 2003; Lundh et al., 2017). A likelihood ratio larger than 1 was verified to be robust to controls for design features and statistical power, among other characteristics (Kjaergard and Als-Nielsen, 2002; Friedman and Richter, 2004).

Next follow a number of qualitative constraints. (1.7) says that (ceteris paribus) the truth of the hypothesis makes positive evidence e more likely. (1.8) and (1.9) say that (ceteris paribus) lowering the quality of the study design makes the study worse at tracking the truth. If $\bar{h}$ is the case, then a low quality study will make e more likely than a higher quality one. If h is the case, then a low quality study will make e less likely than a higher quality one. (1.10) and (1.11) say that (ceteris paribus) CoI makes e more likely. Note that combining (1.8) and (1.11) gives

$$1 = P(e|hqc) > P(e|h\bar{q}c), P(e|hq\bar{c}) > P(e|h\bar{q}\bar{c}) \ . \tag{1.12}$$

That is, when h is the case, CoI and high quality raise the probability of e. Moreover, combining (1.9) and (1.10) gives

$$0 = P(e|\bar{h}q\bar{c}) < P(e|\bar{h}qc), P(e|\bar{h}\bar{q}\bar{c}) < P(e|\bar{h}\bar{q}c) \ . \tag{1.13}$$

That is, when $\bar{h}$ is the case, CoI and *low* quality raise the probability of e. By inspecting (1.12) and (1.13), one may see that (1.8) to (1.11) entail that c raises the probability of e, no matter the fixed values of H and Q. In particular, CoI increases the probability of e both if H is true—because

CoI raises quality, which makes it more likely to get evidence confirming h—and if H is false—because CoI decreases the probability of $\bar{e}$, which makes it less likely to get evidence disconfirming h. In sum, one may conclude that CoI has both a positive and a negative influence on truth-tracking. Which of the two prevails depends on the strengths of its effects on e and q.

1.5.3 Meta-research and the dual role of conflict of interest

Even if it is accepted by a rational agent who endorses (1.8) to (1.11) that CoI has both a positive and a negative influence on truth-tracking, neither those qualitative constraints alone, nor the addition of the quantitative ones in (1.1) to (1.5)—but, crucially, not (1.6)—would force them to conclude that C has a *significant* biasing influence on E, which in turn translates into a *non-negligible* parameter of the influence along the direct edge $C \longrightarrow E$. In fact, while (1.10) and (1.11) maintain that C is (ceteris paribus) relevant to E, they are consistent with this influence being extremely small given Q, to the point that the existence of a direct edge $C \longrightarrow E$ may be neglected without detriment.

Now, let us consider adding (1.6) to one's total evidence. To begin with, note that the ratio in (1.6) may be unpacked as follows:

$$\frac{P(e|c)}{P(e|\bar{c})} = \frac{P(ec)\cdot P(\bar{c})}{P(e\bar{c})\cdot P(c)} = \frac{P(\bar{c})}{P(c)}\cdot\frac{P(hecq)+P(hec\bar{q})+P(\bar{h}ecq)+P(\bar{h}ec\bar{q})}{P(he\bar{c}q)+P(he\bar{c}\bar{q})+P(\bar{h}e\bar{c}q)+P(\bar{h}e\bar{c}\bar{q})}$$
$$= \frac{P(e|hcq)\cdot P(h)\cdot P(q|c)+P(e|hc\bar{q})\cdot P(h)\cdot P(\bar{q}|c)+P(e|\bar{h}cq)\cdot P(\bar{h})\cdot P(q|c)+P(e|\bar{h}c\bar{q})\cdot P(\bar{h})\cdot P(\bar{q}|c)}{P(e|h\bar{c}q)\cdot P(h)\cdot P(q|\bar{c})+P(e|h\bar{c}\bar{q})\cdot P(h)\cdot P(\bar{q}|\bar{c})+P(e|\bar{h}\bar{c}q)\cdot P(\bar{h})\cdot P(q|\bar{c})+P(e|\bar{h}\bar{c}\bar{q})\cdot P(\bar{h})\cdot P(\bar{q}|\bar{c})}$$

Given that $P(h)$ is by assumption very small (1.1), that the first two terms in numerator and denominator are negligible, and that $P(\bar{h})$ in the other terms can be factored out and simplified, the expression approximates to

$$\frac{P(e|c)}{P(e|\bar{c})} \approx \frac{P(e|\bar{h}cq)\cdot P(q|c)+P(e|\bar{h}c\bar{q})\cdot P(\bar{q}|c)}{P(e|\bar{h}\bar{c}q)\cdot P(q|\bar{c})+P(e|\bar{h}\bar{c}\bar{q})\cdot P(\bar{q}|\bar{c})} = \frac{P(e|\bar{h}cq)\cdot P(q|c)+P(e|\bar{h}c\bar{q})\cdot P(\bar{q}|c)}{[P(e|\bar{h}\bar{c}q)+P(e|\bar{h}\bar{c}\bar{q})]/2}\,.$$

A conservative yet rational agent, who believes that C is almost irrelevant to E given Q, will judge the terms $\frac{P(e|\bar{h}cq)}{P(e|\bar{h}\bar{c}q)}$ and $\frac{P(e|\bar{h}c\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})}$ close to 1, such that the ratio will ultimately be driven by $P(q|c)$ and $P(q|\bar{c})$. In that case,

$$\frac{P(e|c)}{P(e|\bar{c})} \approx \frac{P(e|\bar{h}cq)\cdot P(q|c)+P(e|\bar{h}c\bar{q})\cdot P(\bar{q}|c)}{[P(e|\bar{h}\bar{c}q)+P(e|\bar{h}\bar{c}\bar{q})]/2}$$
$$= \frac{\frac{P(e|\bar{h}cq)}{2}+P(e|\bar{h}c\bar{q})\cdot P(\bar{q}|c)}{[P(e|\bar{h}\bar{c}q)+P(e|\bar{h}\bar{c}\bar{q})]/2} + \frac{(0.9-P(c|q))\cdot P(e|\bar{h}cq)}{[P(e|\bar{h}\bar{c}q)+P(e|\bar{h}\bar{c}\bar{q})]/2}$$

$$\leq \frac{[P(e|\bar{h}cq) + P(e|\bar{h}c\bar{q})]/2}{[P(e|\bar{h}\bar{c}q) + P(e|\bar{h}\bar{c}\bar{q})]/2} + \frac{0.4 \cdot P(e|\bar{h}cq)}{[P(e|\bar{h}\bar{c}q) + P(e|\bar{h}\bar{c}\bar{q})]/2}$$
$$\leq 1 + 0.8 \cdot \frac{P(e|\bar{h}cq)}{P(e|\bar{h}\bar{c}q) + P(e|\bar{h}\bar{c}\bar{q})} < 1 + 0.8 \cdot \frac{P(e|\bar{h}cq)}{P(e|\bar{h}\bar{c}q)} = 1.8 \ .$$

That is, a rational agent, who is also a strong supporter of the confirmatory value of C via its positive effect on Q will assign values, which in the most favourable case, $P(q|c) = 0.9$, return an upper bound for interval of the likelihood ratio equal to 1.8.

At the same time, coherence requires avoiding probability assignments violating (1.6), once the latter piece of evidence becomes known. A rational agent will thus have to accommodate their beliefs in light of (1.6) to ensure that the range of possible values of the ratio $\frac{P(e|c)}{P(e|\bar{c})}$ *contains the interval* $[1.3, 4]$. The only probabilities, which have not yet been pinned down, are those of E conditional on its parents H, Q, C. Since $P(h)$ being very small makes almost irrelevant probabilities of e conditional on h, $P(e|h...)$, the most plausible way for a rational agent to adjust their beliefs is to increase the value of either the ratio $P(e|\bar{h}cq)/P(e|\bar{h}\bar{c}q)$, or the ratio $P(e|\bar{h}c\bar{q})/P(e|\bar{h}\bar{c}\bar{q})$, or both. This boils down to acknowledging that C has a relevant effect on the study's outcome e, conditional on its quality Q. In other words, the principle of total evidence dictates that a conservative yet rational agent assign to CoI a non-redundant role for confirmation, given quality.

1.5.4 Conflict of interest and Bayesian confirmation

Let us now come to the illustration of how the model supports our main claims, starting with (A): it is unjustified to generally neglect CoI despite its ambiguity, because CoI can make a difference —whether positive or negative— to the confirmation of the hypothesis.

Theorem 1.1 (Confirmation by CoI-laden study)**.** *A CoI-laden study is confirmatory if the relevant Bayes factor* > 1*:*

$$\text{sign}(P(h|ec) - P(h)) = \text{sign}\Big(\frac{P(e|hcq) \cdot P(q|c) + P(e|hc\bar{q}) \cdot P(\bar{q}|c)}{P(e|\bar{h}cq) \cdot P(q|c) + P(e|\bar{h}c\bar{q}) \cdot P(\bar{q}|c)} - 1\Big) \ .$$

Theorem 1.2 (Confirmation by CoI-laden study, irrespective of quality)**.** *A CoI-laden study is more confirmatory for* $H = h$ *iff* $P(h|e\bar{c}) < P(h|e)$*:*

$$\text{sign}(P(h|ec) - P(h|e\bar{c})) = \text{sign}(P(h|ec) - P(h|e)) = \text{sign}\Big(\frac{P(e|hc)}{P(e|\bar{h}c)} - \frac{P(e|h\bar{c})}{P(e|\bar{h}\bar{c})}\Big)$$
$$= \text{sign}\Big(\frac{P(e|hcq) \cdot P(q|c) + P(e|hc\bar{q}) \cdot P(\bar{q}|c)}{P(e|\bar{h}cq) \cdot P(q|c) + P(e|\bar{h}c\bar{q}) \cdot P(\bar{q}|c)} - \frac{P(e|h\bar{c}q) \cdot P(q|\bar{c}) + P(e|h\bar{c}\bar{q}) \cdot P(\bar{q}|\bar{c})}{P(e|\bar{h}\bar{c}q) \cdot P(q|\bar{c}) + P(e|\bar{h}\bar{c}\bar{q}) \cdot P(\bar{q}|\bar{c})}\Big) \ .$$

In particular, 1.2 entails that CoI can not only decrease but also *increase* confirmation, namely if c is a good predictor of $H = h$, that is, if c makes e much more likely given h, and only slightly more likely given $\bar{h}$.

Next, let us turn to prove (B): it is unjustified to generally (or necessarily, or a priorily) neglect CoI vis-à-vis quality.

Theorem 1.3 (Confirmation by low-quality & no-CoI vs CoI & high--quality)**.**

$$\begin{aligned}\operatorname{sign}(P(h|e\bar{c}\bar{q}) - P(h|ecq)) &= \operatorname{sign}\Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big)\\ &= \operatorname{sign}\Big(\frac{P(e|\bar{h}cq)}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hcq)}{P(e|h\bar{c}\bar{q})}\Big)\ .\end{aligned}$$

In general, neither one between neglect of CoI (given high quality) and neglect of quality (in the absence of CoI) is less relevant to confirmation, and thus more justifiable on the ground that it is a smaller violation of the principle of total evidence.

Theorem 1.4 (Confirmation by low-quality vs CoI)**.**

$$\begin{aligned}\operatorname{sign}(P(h|e\bar{q}) - P(h|ec) = \operatorname{sign}\Big(&\Big(\frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) P(e|\bar{h}c\bar{q})\, P(e|\bar{h}cq)\, P(c)^2\, P(q|c)\, P(\bar{q}|c)\\ &+ \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) P(e|\bar{h}\bar{c}\bar{q})\, P(e|\bar{h}cq)\, P(c)\, P(\bar{c})\, P(q|c)\, P(\bar{q}|\bar{c})\\ &+ \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})}\Big) P(e|\bar{h}\bar{c}\bar{q})\, P(e|\bar{h}c\bar{q})\, P(c)\, P(\bar{c})\, P(\bar{q}|c)\, P(\bar{q}|\bar{c})\Big)\ .\end{aligned}$$

In general, neither one between neglect of CoI (irrespective of quality) and neglect of quality (irrespective of CoI) is less relevant to confirmation, and thus in principle more justifiable.

Of course, to disentangle CoI's truth-tracking contribution, namely the one via Q, from the "truth-diverting" contribution, namely the one that introduces biases via E directly, one needs (plausible) numerical assignments for all of the probabilities above. If one can precisely assign such values, one can use the theorem to verify which contribution prevails, by looking at the sign of $(P(h|eq) - P(h)) - (P(h|ec) - P(h)) = P(h|eq) - P(h|ec)$. If we employ a notion of rationality where imprecise probabilities are allowed, however, there is no guarantee that the sign is conclusively decidable. If meta-research doesn't provide sufficient information to pin down all probabilities, the model may fail to eliminate the uncertainty on whether there is confirmation or not.

At the same time, note that an ideal Bayesian agent assigns exact values to all of the probabilities in the above theorem, which makes it in

principle possible to use our model to definitely determine whether CoI confirms or disconfirms on any given occasion. Moreover, even if—more realistically—meta-research doesn't justify precise assignments, it may still suffice to constrain the probabilities to intervals, which are narrow enough, such that the value of that difference does not span both positive and negative values, and the only uncertainty left concerns the amount of confirmation but not its sign. In this sense, the model should be viewed as a roadmap, which allows (non-ideal) agents to collect as much information as needed for the model to return acceptable answers, given the level of accuracy required by the problem at hand.

1.6 Conclusion

Existing literature on Bayesian confirmation supports the view that less-than-ideal evidence—from, say, medical studies—provides confirmatory benefits if a reliability-based discount is introduced to account for the likelihood of the evidence being produced not only by the truth of the hypothesis but also by underlying biases. This is in line with the principle of total evidence, which recommends that one takes into account all available evidence when determining probabilities.

At the same time, it has been noted that medical studies are subject not only to biases but also to conflicts of interest. It is unclear, however, whether evidence obtained from studies subject to a conflict of interest, too, should be discounted, given that conflict of interest has an ambiguous role. On the one hand, in fact, it promotes larger studies with better designs, which improves their accuracy by making certain biases less likely. On the other hand, it introduces subtle and difficult-to-detect biases, which may tend to skew the results, by for instance overestimating effect sizes. Without a concrete model, however, it is difficult to assess the bearing of conflict of interest on the hypothesis under scrutiny, and more generally, the benefit of complying with the principle of total evidence by taking evidence of conflict of interest into account.

This paper has addressed the issue by providing a Bayesian model, which disambiguates the dual role of conflict of interest as a promoter of bias and of quality along distinct paths, and interprets its influences along the two paths as higher-order constraints obtained from meta-research. The analysis of the model shows that information on conflict of interest can benefit confirmation. More precisely, it is unjustified to neglect conflict of interest, whether or not information on quality is available, or to

prefer the neglect of conflict of interest vis-à-vis quality. Together, these results vindicate the endorsement of the principle of total evidence in the presence of conflict of interest.

Two important issues have not been addressed here. First, our model only considers Bayesian updates by a single study. It can be extended to updates by evidence from multiple studies. In that case, however, the choice of the model will depend on the source(s) of conflict of interest behind the studies, more precisely on whether each study is subject to a different conflict of interest, or all studies are subject to the same conflict of interest or, more realistically, on where the studies are located on an interval between these two extremes. The set up and the analysis of this extended model raises technical issues, which cannot be reviewed here (the interested reader is referred to Casini and Landes 2022).

Second, our model uses binary variables. Only considering the significance of a result is, of course, a gross simplification. Public health decisions depend on numerical estimates of the effect size. The impact of conflict of interest on the veracity of a study depends crucially on the observed effect size. Among significant studies subject to conflict of interest, some may show a larger effect than others, and this is certainly relevant to the bearing of the study on the estimate, and to how the estimate should be revised in light of the evidence. For a Bayesian model to inform such estimates, one must specify how a continuous prior of an effect size (characterized by, say, a given mean and variance) should be updated given the evidence (an empirical distribution with another mean and variance) under the assumed constraints. To this end, however, analytic results are not sufficient and one must resort to numerical simulations (see, e.g., Welton et al. 2009 and Verde 2021). Such simulations, in turn, could be instrumental to revising estimates by not only single studies but also collections of such studies, as is typical done by meta-analyses.

Jointly, these two issues motivate a study of how to devise a meta-analytic revision of the effect size (in the presence of conflicts of interest) in a Bayesian framework. We leave this study to future research.

Proofs

Proof of Theorem 1.1 (Confirmation by CoI-laden study).

$$\begin{aligned}
\operatorname{sign}(P(h|ec) - P(h)) &= \operatorname{sign}(P(hec) - P(h) \cdot P(ec)) \\
&= \operatorname{sign}(P(hec) - P(h) \cdot [P(hec) + P(\bar{h}ec)]) \\
&= \operatorname{sign}(P(\bar{h}) \cdot P(hec) - P(h) \cdot P(\bar{h}ec)]) \\
&= \operatorname{sign}\Big(\frac{P(hecq) + P(hec\bar{q})}{P(\bar{h}ecq) + P(\bar{h}ec\bar{q})} - \frac{P(h)}{P(\bar{h})}\Big) \\
&= \operatorname{sign}\Big(\frac{P(e|hcq) \cdot P(cq) + P(e|hc\bar{q}) \cdot P(c\bar{q})}{P(e|\bar{h}cq) \cdot P(cq) + P(e|\bar{h}c\bar{q}) \cdot P(c\bar{q})} - 1\Big) \\
&= \operatorname{sign}\Big(\frac{P(e|hcq) \cdot P(q|c) + P(e|hc\bar{q}) \cdot P(\bar{q}|c)}{P(e|\bar{h}cq) \cdot P(q|c) + P(e|\bar{h}c\bar{q}) \cdot P(\bar{q}|c)} - 1\Big) \ .
\end{aligned}$$

□

Proof of Theorem 1.2 (Confirmation by CoI-laden study, irrespective of quality).

$$\begin{aligned}
&\operatorname{sign}(P(h|e\bar{c}) - P(h|ec)) = \operatorname{sign}([P(he\bar{c}) \cdot P(\bar{h}ec)] - [P(hec) \cdot P(\bar{h}e\bar{c})]) \\
&= \operatorname{sign}([(P(he\bar{c}q) + P(he\bar{c}\bar{q})) \cdot (P(\bar{h}ecq) + P(\bar{h}ec\bar{q}))] \\
&\quad - [(P(hecq) + P(hec\bar{q})) \cdot (P(\bar{h}e\bar{c}q) + P(\bar{h}e\bar{c}\bar{q}))]) \\
&= \operatorname{sign}([P(h) \cdot (P(e|h\bar{c}q) \cdot P(\bar{c}q) + P(e|h\bar{c}\bar{q}) \cdot P(\bar{c}\bar{q})) \cdot P(\bar{h}) \cdot (P(e|\bar{h}cq) \cdot P(cq) + P(e|\bar{h}c\bar{q}) \cdot P(c\bar{q}))] \\
&\quad - [P(h) \cdot (P(e|hcq) \cdot P(cq) + P(e|hc\bar{q}) \cdot P(c\bar{q})) \cdot P(\bar{h}) \cdot (P(e|\bar{h}\bar{c}q) \cdot P(\bar{c}q) + P(e|\bar{h}\bar{c}\bar{q}) \cdot P(\bar{c}\bar{q}))]) \\
&= \operatorname{sign}([(P(e|h\bar{c}q) \cdot P(q|\bar{c}) + P(e|h\bar{c}\bar{q}) \cdot P(\bar{q}|\bar{c})) \cdot (P(e|\bar{h}cq) \cdot P(q|c) + P(e|\bar{h}c\bar{q}) \cdot P(\bar{q}|c))] \\
&\quad - [(P(e|hcq) \cdot P(q|c) + P(e|hc\bar{q}) \cdot P(\bar{q}|c)) \cdot (P(e|\bar{h}\bar{c}q) \cdot P(q|\bar{c}) + P(e|\bar{h}\bar{c}\bar{q}) \cdot P(\bar{q}|\bar{c}))]) \\
&= \operatorname{sign}(\frac{P(e|h\bar{c}q) \cdot P(q|\bar{c}) + P(e|h\bar{c}\bar{q}) \cdot P(\bar{q}|\bar{c})}{P(e|\bar{h}\bar{c}q) \cdot P(q|\bar{c}) + P(e|\bar{h}\bar{c}\bar{q}) \cdot P(\bar{q}|\bar{c})} - \frac{P(e|hcq) \cdot P(q|c) + P(e|hc\bar{q}) \cdot P(\bar{q}|c)}{P(e|\bar{h}cq) \cdot P(q|c) + P(e|\bar{h}c\bar{q}) \cdot P(\bar{q}|c)}) \ .
\end{aligned}$$

□

Proof of Theorem 1.3 (Confirmation by low-quality & no-CoI vs CoI & high-quality).

$$\begin{aligned}
&\operatorname{sign}(P(h|e\bar{c}\bar{q}) - P(h|ecq)) = \operatorname{sign}(P(he\bar{c}\bar{q}) \cdot P(ecq) - P(hec\bar{q}) \cdot P(e\bar{c}\bar{q})) \\
&\quad = \operatorname{sign}(P(he\bar{c}\bar{q}) \cdot [P(hecq) + P(\bar{h}ecq)] - P(hecq) \cdot [P(he\bar{c}\bar{q}) + P(\bar{h}e\bar{c}\bar{q})]) \\
&\quad = \operatorname{sign}(P(he\bar{c}\bar{q}) \cdot P(\bar{h}ecq) - P(hecq) \cdot P(\bar{h}e\bar{c}\bar{q})) \\
&\quad = \operatorname{sign}(P(h) \cdot P(\bar{h}) \cdot P(cq) \cdot P(\bar{c}\bar{q}) \cdot [P(e|h\bar{c}\bar{q}) \cdot P(e|\bar{h}cq) - P(e|hcq) \cdot P(e|\bar{h}\bar{c}\bar{q})]) \\
&\quad = \operatorname{sign}(P(e|h\bar{c}\bar{q}) \cdot P(e|\bar{h}cq) - P(e|hcq) \cdot P(e|\bar{h}\bar{c}\bar{q})) \\
&\quad = \operatorname{sign}\Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) \ .
\end{aligned}$$

□

Proof of Theorem 1.4 (Confirmation by low-quality vs CoI).

$$
\begin{aligned}
\operatorname{sign}(P(h|e\bar{q}) - P(h|ec)) &= \operatorname{sign}(P(he\bar{q}) \cdot [P(hecq) + P(hec\bar{q}) + P(\bar{h}ecq) + P(\bar{h}ec\bar{q})] \\
&\quad - P(hec) \cdot [P(hec\bar{q}) + P(he\bar{c}\bar{q}) + P(\bar{h}ec\bar{q}) + P(\bar{h}e\bar{c}\bar{q})]) \\
&= \operatorname{sign}([P(hec\bar{q}) + P(he\bar{c}\bar{q})] \cdot [P(hecq) + P(hec\bar{q}) + P(\bar{h}ecq) + P(\bar{h}ec\bar{q})] \\
&\quad - [P(hecq) + P(hec\bar{q})] \cdot [P(hec\bar{q}) + P(he\bar{c}\bar{q}) + P(\bar{h}ec\bar{q}) + P(\bar{h}e\bar{c}\bar{q})]) \\
&= \operatorname{sign}(P(hec\bar{q}) \cdot P(\bar{h}ecq) + P(he\bar{c}\bar{q}) \cdot P(\bar{h}ecq) + P(he\bar{c}\bar{q}) \cdot P(\bar{h}ec\bar{q}) \\
&\quad - [P(hecq) \cdot P(\bar{h}ec\bar{q}) + P(hecq) \cdot P(\bar{h}e\bar{c}\bar{q}) + P(hec\bar{q}) \cdot P(\bar{h}e\bar{c}\bar{q})]) \\
&= \operatorname{sign}([P(e|hc\bar{q}) \cdot P(e|\bar{h}cq) - P(e|hcq) \cdot P(e|\bar{h}c\bar{q})] \cdot P(h) \cdot P(\bar{h}) \cdot P(c\bar{q}) \cdot P(cq) \\
&\quad + [P(e|h\bar{c}\bar{q}) \cdot P(e|\bar{h}cq) - P(e|hcq) \cdot P(e|\bar{h}\bar{c}\bar{q})] \cdot P(h) \cdot P(\bar{h}) \cdot P(\bar{c}\bar{q}) \cdot P(cq) \\
&\quad + [P(e|h\bar{c}\bar{q}) \cdot P(e|\bar{h}c\bar{q}) - P(e|hc\bar{q}) \cdot P(e|\bar{h}\bar{c}\bar{q})] \cdot P(h) \cdot P(\bar{h}) \cdot P(\bar{c}\bar{q}) \cdot P(c\bar{q})) \\
&= \operatorname{sign}\Big(\Big(\frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) \cdot P(e|\bar{h}c\bar{q}) \cdot P(e|\bar{h}cq) \cdot P(c\bar{q}) \cdot P(cq) \\
&\quad + \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) \cdot P(e|\bar{h}\bar{c}\bar{q}) \cdot P(e|\bar{h}cq) \cdot P(\bar{c}\bar{q}) \cdot P(cq) \\
&\quad + \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})}\Big) \cdot P(e|\bar{h}\bar{c}\bar{q}) \cdot P(e|\bar{h}c\bar{q}) \cdot P(\bar{c}\bar{q}) \cdot P(c\bar{q})\Big) \\
&= \operatorname{sign}\Big(\Big(\frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) \cdot \frac{P(e\bar{h}c\bar{q})}{P(\bar{h}c\bar{q})} \cdot \frac{P(e\bar{h}cq)}{P(\bar{h}cq)} \cdot P(c\bar{q}) \cdot P(cq) \\
&\quad + \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) \cdot \frac{P(e\bar{h}\bar{c}\bar{q})}{P(\bar{h}\bar{c}\bar{q})} \cdot \frac{P(e\bar{h}cq)}{P(\bar{h}cq)} \cdot P(\bar{c}\bar{q}) \cdot P(cq) \\
&\quad + \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})}\Big) \cdot \frac{P(e\bar{h}\bar{c}\bar{q})}{P(\bar{h}\bar{c}\bar{q})} \cdot \frac{P(e\bar{h}c\bar{q})}{P(\bar{h}c\bar{q})} \cdot P(\bar{c}\bar{q}) \cdot P(c\bar{q})\Big) \\
&= \operatorname{sign}\Big(\Big(\frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) \cdot P(e|\bar{h}c\bar{q}) \cdot P(e|\bar{h}cq) \cdot P(c\bar{q}) \cdot P(cq) \\
&\quad + \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hcq)}{P(e|\bar{h}cq)}\Big) \cdot P(e|\bar{h}\bar{c}\bar{q}) \cdot P(e|\bar{h}cq) \cdot P(\bar{c}\bar{q}) \cdot P(cq) \\
&\quad + \Big(\frac{P(e|h\bar{c}\bar{q})}{P(e|\bar{h}\bar{c}\bar{q})} - \frac{P(e|hc\bar{q})}{P(e|\bar{h}c\bar{q})}\Big) \cdot P(e|\bar{h}\bar{c}\bar{q}) \cdot P(e|\bar{h}c\bar{q}) \cdot P(\bar{c}\bar{q}) \cdot P(c\bar{q})\Big) \ .
\end{aligned}
$$

□

Bibliography

Bovens, L. and Hartmann, S. (2003). *Bayesian Epistemology*. Oxford University Press, New York.

Carnap, R. (1947). On the Application of Inductive Logic. *Philosophy and Phenomenological Research*, 8(1):133–148.

Casini, L. and Landes, J. (2022). Confirmation by robustness analysis: A bayesian account. *Erkenntnis*.

Egger, M., Smith, G. D., Schneider, M., and Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal*, 315(629).

Fagerland, M. (2015). Evidence-based medicine and systematic reviews. *Research in Medical and Biological Sciences: From Planning and Preparation to Grant Application and Publication*, pages 431–61.

Friedman, L. S. and Richter, E. D. (2004). Relationship between conflicts of interest and research results. *Journal of General Internal Medicine*, 19:51–56.

Fuller, J. (2018). Meta-research evidence for evaluating therapies. *Philosophy of Science*, 85:767–80.

Gomes, D. and Stavropoulou, C. (2019). The impact generated by publicly and charity-funded research in the united kingdom: a systematic literature review. *Health Research Policy and Systems*, 17:22.

Hartling, L., Ospina, M., Liang, Y., Dryden, D. M., Hooton, N., Krebs Seida, J., and Klassen, T. P. (2009). Risk of bias versus quality assessment of randomised controlled trials: cross sectional study. *BMJ*, 339:b4012.

Higgins, J. and Green, S., editors (2011). *Cochrane Handbook for Systematic Reviews of Interventions*. The Cochrane Collaboration, 2011, version 5.1.0 edition.

Ioannidis, J. (2005). Why most published research findings are false. *PLoS Med*, 8(2):e124.

Ioannidis, J. (2008). Why most discovered true associations are inflated. *Epidemiology*, 19(5):640–48.

Ioannidis, J. P. A., Fanelli, D., Dunne, D. D., and Goodman, S. N. (2015). Meta-research: Evaluation and Improvement of Research Methods and Practices. *PLOS Biology*, 13(10):e1002264.

Jørgensen, A. W., Maric, K. L., Tendal, B., Faurschou, A., and Gøtzsche, P. C. (2008). Industry-supported meta-analyses compared with meta-analyses with non-profit or no support: Differences in methodological quality and conclusions. *BMC Medical Research Methodology*, 8:60.

Jüni, P., Witschi, A., Bloch, R., and Egger, M. (1999). The hazards of scoring the quality of clinical trials for meta-analysis. *JAMA*, 282:1054–60.

Kjaergard, L. L. and Als-Nielsen, B. (2002). Association between competing interests and authors' conclusions: Epidemiological study of randomised clinical trials published in the BMJ. *BMJ*, 325:249.

Lexchin, J. (2012). Those who have the gold make the evidence: How the pharmaceutical industry biases the outcomes of clinical trials of medications. *Sci Eng Ethics*, 18:247–61.

Lexchin, J., Bero, L., Djubegovic, B., and Clark, O. (2003). Pharmaceutical industry sponsored research: Evidence for a systematic bias. *BMJ*, 326:1167–70.

Lundh, A., Lexchin, J., Mintzes, B., Schroll, J. B., and Bero, L. (2017). Industry sponsorship and research outcome. *Cochrane Database Systematic Reviews*, 12:MR000033.

Montgomery, J. H., Byerly, M., Carmody, T., Li, B., Miller, D. R., Varghese, F., and Holland, R. (2004). An analysis of the effect of funding source in randomized clinical trials of second generation antipsychotics forthe treatment of schizophrenia. *Controlled Clinical Trials*, 25(6):598–612.

Osimani, B. and Landes, J. (2023). Varieties of Error and Varieties of Evidence in Scientific Inference. *British Journal for the Philosophy of Science*, 74(1):117–170.

Reynolds, T. (2001). Industry-funded versus publicly funded trials: Are the standards the same? *JNCI: Journal of the National Cancer Institute*, 93(21):1590–1592.

Roseman, M., Milette, K., Bero, L., Coyne, J., Lexchin, J., Turner, E., and Thombs, B. (2011). Reporting of conflicts of interest in meta-analyses of trials of pharmacological treatments. *JAMA*, 305(10):1008–17.

Sackett, D., Rosenberg, W., Gray, J., Haynes, R., and Richardson, W. (1996). Evidence based medicine: what it is and what it isn't. *British Medical Journal*, 312(7023):71–2.

Shea, B. J., Grimshaw, J. M., and G A Wells, e. (2007). Development of amstar: a measurement tool to assess the methodological quality of systematic reviews. *BMC Med Res Methodol*, 7:10.

Slavin, R. E. (1986). Best-evidence synthesis: An alternative to meta-analytic and traditional reviews. *Educational Researcher*, 15(9):5–11.

Stegenga, J. (2018). *Medical nihilism*. Oxford University Press.

Verde, P. E. (2021). A bias-corrected meta-analysis model for combining, studies of different types and quality. *Biometrical Journal*, 63(2):406–22.

Welton, N. J., Ades, A. E., Carlin, J. B., Altman, D. G., and Sterne, J. A. C. (2009). Models for potentially biased evidence in meta-analysis using empirically based priors. *Journal of the Royal Statistical Society*, 172:119–36.

West, S., King, V., and T S Carey, e. (2002). Systems to rate the strength of scientific evidence. techreport 47, Rockville, MD: Agency for Healthcare Research and Quality.

Worrall, J. (2002). What evidence in evidence-based medicine? *Philosophy of Science*, 69:S316–30.

Yank, V., Rennie, D., and Bero, L. (2007). Financial ties and concordance between results and conclusions in meta-analyses: retrospective cohort study. *British Medical Journal*, 335:1202–05.