

## Evidence Mapping to Justify Health Interventions

Olaf Dammann

Department of Public Health and Community Medicine, Tufts University School of Medicine, Boston; and Department of Gynecology and Obstetrics, Hannover Medical School, Germany.

Email: [olaf.dammann@tufts.edu](mailto:olaf.dammann@tufts.edu).

Correspondence: Olaf Dammann, Department of Public Health and Community Medicine, Tufts University School of Medicine, Boston, MA 02111.

The author is grateful to Alan Leviton, Nigel Paneth, Alex Broadbent, and two anonymous reviewers for their helpful comments on this paper.

### ABSTRACT

In order to support health interventions, biomedical and population health researchers need to collect solid evidence. This article asks what type of evidence this should be and expands on previous work that focused on etiological explanations, or causal-mechanical explanations of why and how illness occurs. The article proposes adding predictive evidence to the explanatory evidence, in order to form a *joint evidence set*, or  $JES = [A,B,C,D]$ , which consists of four different types of evidence: association [A], biology [B], confirmation [C], and difference-making [D]. The article discusses explanatory coherence as a backbone for this proposal, suggests criteria for each type of evidence, and offers a rubric for multi-evidence mapping.

Health interventions need to be well justified. Arguably, helpfulness is a good justifier. In turn, good justifications need to be supported by solid evidence. While some of the process of generating such evidence is formalized in individual branches of the health sciences, an overarching framework for the integration of multiple *different* types of evidence is lacking. In this essay, I propose a semi-formal methodological framework, *multilevel evidence mapping*, that helps organize and interpret the available evidence in support of health interventions. In brief, I will argue that the collection of evidence from multiple sources in support of health interventions is a two-step process (Table 1).

**Table 1** *Fourfold table of explanatory (top) and predictive (bottom) evidence, provided by observational (left) and interventional (right) research*

	<i>Observation</i>	<i>Intervention</i>
Step 1 <b>EXPLANATION</b>	<b>A.</b> <b>Association</b> <i>“There is something!”</i> Observational epidemiology	<b>B.</b> <b>Biology</b> <i>“It might work like this.”</i> Laboratory science
Step 2 <b>PREDICTION</b>	<b>C.</b> <b>Confirmation</b> <i>“It happens—as expected!”</i> Predicted observations	<b>D.</b> <b>Difference</b> <i>“We can make a difference.”</i> Clinical and population trials

The first step is to gather *explanatory* evidence from epidemiological observations and interventional laboratory experiments. The second step involves making predictions based on these explanations and gathering *predictive* evidence by checking whether these predictions are accurate. My main argument is that it is not necessarily the randomized intervention trial, but the explanatory-predictive coherence of the above evidence that justifies health interventions, even in the absence of randomized trials. This argument dovetails with, for example, Paneth and Joyner’s (2021) discussion of evidence supporting the use of convalescent serum in COVID-19 patients, which includes the warning that “we should not be paralyzed into inaction when [randomized controlled trials] are not available.”

The argument builds on the viewpoints offered by Sir Austin Bradford Hill (1965) in support of causal inference based on observational epidemiological data, and on my previous interpretation of Hill’s viewpoints as an explanatory coherentist framework (Dammann 2018b). In brief, Hill proposed nine viewpoints to be considered before attaching causal meaning to an association identified in observational epidemiological studies: strength of the association, consistency and specificity of findings, cause before effect (temporality), biological gradient (dose-response), plausibility, coherence, experiment (intervention success), and analogy (preexisting knowledge about similar associations). Some, but not all, of these viewpoints are captured in one way or another by the framework I propose in this paper.

Here is a brief roadmap of the sections to follow. I begin by outlining how clinical and epidemiological observations provide evidence that a certain health phenomenon occurs and which risk factors, if any, are statistically associated with its occurrence [A]. Laboratory experiments are then designed to clarify the pathogenetic process that connects risk factors and their consequences mechanistically [B]. Together, [A] and [B] represent the first explanatory step, providing what I have called *etiological explanations* (EE) (Dammann 2020), a term I have borrowed from Salmon (1984), who borrowed it from Wright (1976). Etiological explanations are comprehensive descriptions of the causes of illness and the pathogenetic mechanisms that connect the two. One way to establish the “fit” of EE is explanatory coherence modeling.

The second step consists of two kinds of predictions. The first are predictions of the phenomenon under study [C], as they are to be expected based on knowledge about the distribution of risk factors. The second includes predictions of intervention success in trial settings [D]. [C] and [D] stand in a confirmatory relationship to [A] and [B]. However, they also play an explanatory role together with A and B as members of the *joint evidence set*,  $JES = [A,B,C,D]$ .

I next argue that JES exhibits *explanatory-predictive coherence* if (1) members do not contradict each other (non-contradiction); (2) their joint explanatory force is stronger than that of individual members of JES or of combinations of its members (additive synergism); and (3) the evidence they provide turns out to be successful in research settings (efficacy). If interventions based on explanatory-predictive coherence also turn out to be helpful by successfully improving population health outside the research setting (effectiveness), we would consider this a final confirmation of the entire endeavor.

Finally, I describe how one recently developed kind of research synthesis, *evidence mapping*, can help arrange the available evidence in the form of an explanatory-predictive coherence model. The goal is to have a structured rubric to document the availability or absence of different types of evidence, collected in order to justify interventions in medicine and public health. The concluding section of the article summarizes the previous sections.

## **Etiological Explanations**

Modern biomedicine is an orchestrated response to the challenge of illness. The main goal is the reduction of suffering by targeting the etiology of illness. The arguably most useful definition of the term *etiology* includes causes that occur before any bodily response and mechanisms that lead from the initial bodily response to the first manifestations of illness (MacMahon and Pugh 1970). In other words, this view conceptualizes illness development as a process with three main components: causal initiators set biological mechanisms in motion, which result in clinical illness.

The cognitive entry point into the system is knowledge about the characteristics of diseases, disorders, or defects. While some kinds of illness have very general, nonspecific features, such as fever, others have specific (pathognomonic) characteristics, such as the strawberry tongue in Kawasaki disease. Whatever the constellation of symptoms, their initial occurrence in a person can be defined as the onset of clinical illness.

Asking the etiological question of illness occurrence is asking for an explanation of why and how a person has become ill. In response, causal and mechanistic explanations are given, alone or together. A causal explanation is a response to the *why* question, and it is given by listing factors that are considered responsible for the illness to occur. For example, infection with a human immunodeficiency virus (HIV) is considered to be the cause of AIDS; needle sharing among drug users is considered to be one major cause of HIV infection; drug addiction is what leads to needle sharing; and so on. To give a mechanistic explanation is to explain *how* exposure to the cause culminates in the onset of illness. This mechanism is also known as the pathogenesis of the illness under consideration. It is the biological cascade of events that connects exposure to causes with illness, the outcome of the pathogenetic process.

Taken together, causal and mechanistic explanations describe the story of illness occurrence in the form of an EE. A good, comprehensive EE should be useful in at least two ways. First, it specifies the causes of illness in order to define targets for interventions that prevent illness occurrence. This is the task at hand for epidemiologists. The idea is that if we eliminate or at least reduce the occurrence of the cause (exposure), we might eliminate or reduce the disease (outcome). Making sure that fewer people are infected with HIV should lead to fewer cases of AIDS and reducing the number of drug addicts that share needles should lead to fewer cases of HIV infection. Thus, some AIDS prevention programs focus on reducing the need for needle sharing among drug users. Second, comprehensive EEs identify target points for medical interventions designed to interfere with the pathogenesis of illness. This pathogenetic research is performed mainly by laboratory scientists. Drugs have been identified that inhibit viral replication or viral spread. Thus, EEs are very helpful in the design of health interventions by virtue of explaining both causal and pathomechanistic factors.

Providing useful EEs requires solid evidence that supports causal-mechanistic claims. Causal claims are supported by good evidence from modern epidemiological risk factor studies. A candidate risk factor is studied in relation to illness occurrence. If certain requirements are fulfilled, the association between risk factor and illness is considered causal. The two classic techniques to arrive at a point where experts consider causation likely are the quantitative *potential outcomes approach* (POA) proposed by Rubin (1974) and the qualitative approach according to Hill (1965). The former involves judgement of the results of randomized or quasi-randomized epidemiologic studies, while the latter requires the consideration of qualitative viewpoints when making causal inferences based on observed associations (Glass et al. 2013).

Pathogenetic mechanisms are studied in the laboratory by means of biomechanistic controlled experimentation. Clever experimental setup and intervention design ensures that the results can be considered causal. The rationale here is the same as in the POA for observational epidemiologic studies: comparing two randomly generated groups of otherwise identical groups of individuals (animals, cell cultures, or other comparatively simple systems) after only one of the groups has been exposed to a specific intervention will reveal differences between groups that can be considered to be caused by that intervention. In sum, EEs provide information about what causes a specific disease and how it does so. In its simplest form, an EE looks like this: “Cause C contributes to the occurrence of disease D by initiating mechanism M which contributes to the occurrence of disease D.” (For more detailed discussions of the concept, see Dammann 2017, 2020.)

Constructing etiological explanations is desirable even if causation cannot be proven. We all know the “association is not causation” paradigm. Although one may have simplistic situations of the sort that someone jumping off a cliff will always land in the water, one can also always create more-or-less clever exceptions that have a boat located underneath or the jumper operating a hang-glider. Here, causation is inferred from what David Hume (1711–1776) called a “constant connexion.”

But in the context of complex systems like those that dominate the health sciences today, causal inference is not this simple. Miguel Hernán (2004) states that in a perfectly designed and conducted randomized trial, association *is* causation. But even in a randomized trial, nothing more than association is observed. What makes this association causal is Hernán’s belief: his causal inference—not some extra ingredient in the data—renders the observed association different from an observation of a similar association in similar data from a prospective (nonrandomized) cohort study. Thus, I do not share the belief that perfect randomization guarantees causal associations. Indeed, despite heroic efforts, I do not believe perfect randomization achievable. There is always the possibility that results of randomized trials are affected by chance events or residual bias.

Instead, I hold that etiological explanations can be based solely on observed associations, if we have a coherent set of error-independent pieces of evidence. It is the coherence among pieces of evidence that provides explanatory vigor. Think of such explanation as comparable to the coherence of evidence in a criminal lawsuit. It is not just the absence of an alibi that leads to the conviction of the gardener, but the additional evidence provided by an eyewitness, who saw the gardener close to his victim’s home at around the time the murder was committed, plus the presence of his DNA on the murder weapon.

The second requirement, error independence, ensures that if individual pieces of evidence turn out to be erroneous, the remainder of the evidence remains unaffected and valid (Stegenga 2009). In other words, evidential error independence makes sure that invalidation of evidence of one type does not invalidate any evidence of the other types. All evidence I refer to comes from observations in epidemiology or from observations in lab experiments or clinical/population trials. No causal attribution is required. The causal notion is based entirely on the degree to which it is possible to integrate the available error-independent *associative* evidence in a coherent explanatory framework.

In the next section I argue that in order to serve as a justification for health interventions, EEs should be coherent systems of multiple types of evidence that are strong individual predictors, mutually coherent in their predictions, and not jointly weaker in their predictive success than individual predictors are.

### **Explanatory-Predictive Coherence**

What kind of characteristic must scientific evidence have to get us from the belief that an intervention will be helpful to the knowledge that it is highly likely to be? I will not discuss in detail what constitutes knowledge. I simply suggest that an evidential system that integrates explanatory and predictive components is likely to become a belief-justifier, hence knowledge-maker, if it is consistently coherent and jointly more successful than its components. In essence, I

argue that an explanatory-predictive coherentist approach will render a justified intervention *well* justified.

Explanatory coherentism claims that a well-justified belief requires multiple items of evidence that cohere with each other. When the items cohere, they can be offered together in support of hypotheses (Poston 2014). Unlike foundationalism, coherentism does not require a foundation upon which all parts of the system are built, but only mutual support by all parts of the explanatory system. I have recently argued that Hill's qualitative approach to causal inference in epidemiology is an interesting example of an explanatory coherentist framework (Dammann 2018b). In that paper, I suggest that explanatory coherence might be one potential way to integrate causal and mechanistic evidence into comprehensive EE. In this paper, I go one step further and argue that if all types of evidence considered in one such set are coherent, adding predictive evidence to the mix will yield sets of explanatory-predictive evidence that justify interventions designed to improve health.

*Coherence* refers to situations in which pieces of information are viewed in light of each other and examined for mutual and reciprocal support. Information becomes evidence by comparison. If new information appears to be aligned with previously accepted information, and if the pieces of information reinforce each other, the transformation from information to evidence is supported (Dammann 2018a).

I distinguish between *explanatory* evidence, such as data that help us explain cogently, and *predictive* evidence, such as facts about successful prediction. Explanatory evidence is a set of facts that we use to tell the story of why and how a phenomenon comes about, while predictive evidence is a set of facts that establish our ability to correctly forecast when, why, and how future instances of the phenomenon will occur. In medicine, explanatory evidence comes from causal and mechanistic research, which tell us, respectively, why and how illness occurs. Explanatory evidence may include epidemiological association measures (type A evidence) and laboratory data (type B), while predictive evidence may come from the accurate prediction of an association (type C) and of intervention success (type D). Depending on the situation, it may be the observed data that count as "evidence," or it may be the accurate prediction.

Some hold that explanation and prediction are logically identical, only that the thrust of explanations is rooted in the past while the power of predictions will be clear only in the future. However, Rescher (1958) has argued against this notion, using the concept of evidence to explicate the difference. He points out that the difference between evidence from the past and the potential presence of future evidence render explanation and prediction epistemologically asymmetrical. More recent work by Spirtes, Glymour, and Scheines (2000) has explored interesting relationships between causation and prediction, the discussion of which is beyond the scope of this paper.

Multiple authors have analyzed what kinds of evidence it takes to support *causal notions* in the health sciences. Some have stressed the need for both probabilistic and mechanistic evidence (Russo and Williamson 2007), while others emphasize the added benefit of parallel (confirmatory) evidence (Howick, Glasziou, and Aronson 2009). Yet others appear to disagree whether interventional evidence is necessary to support causal claims in medicine (Campaner and Cerri

2020; Gillies 2019). Moreover, Campaner and Cerri (2020) suggest distinguishing between interventional evidence (evidence *obtained by* intervention) and evidence for interventions (evidence that *enables* interventions). However, my analysis in this paper is not about causal inference, but about justification of intervention. An in-depth study of the relationships between these topics seems warranted but is not provided here.

Now, let's dissect JES in terms of the evidence needed to support its *component relations*, association [A], biology [B], confirmation [C], and difference [D] (see Table 2).

#### [A] Association Beyond Chance

The target health phenomenon is real, and we have an idea *why* it occurs by demonstrating a strong *association* between putative determinants and the phenomenon. We need to demonstrate that “there is something to intervene on.” A cause (*c*) is anything that contributes to the occurrence of disease (*d*) as an initiator of a mechanism (*m*) which in turn contributes to the occurrence of (*d*) as a pathomechanism. Thus, we need to demonstrate that

*c* is associated with *d*;  
*c* is associated with *m*; and *m* is associated with *d*.

These associations need to be statistical (and ideally statistically significant, free of bias, etc.), because we are interested in the co-occurrence of *c* and *d*, of *c* and *m*, and of *m* and *d* beyond what would be expected by chance.

#### [B] Biology

We further need to demonstrate that we have an idea of *how* the phenomenon occurs and what the underlying *biology* might be (“it might work like this”):

*c* can biologically plausibly initiate *m*; and  
*m* can biologically plausibly be a pathomechanism for *d*.

#### [C] Confirmation

We also have to demonstrate that the phenomenon “occurs as expected” in real life, not just in the research lab, based on our knowledge from [A] and [B]. This is *confirmatory* evidence as well as evidence of successful prediction. We have to show that we can predict the presence of *m* based on knowledge about the presence of *c*. Moreover, we have to be able to predict the occurrence of *d*, given *c* and *m*. Thus, we need to demonstrate that knowledge about the presence of *c* or *m* enables the prediction of *d*. We need to demonstrate that

*m* occurs more frequently in the presence of *c* than in the absence of *c*;  
*d* occurs more frequently in the presence of *c* than in the absence of *c*; and *d* occurs more frequently in the presence of *m* than in the absence of *m*.

**Table 2. Criteria for proposed explanatory-predictive coherence model**

Field System Approach	Epidemiology		Bench
	Human		Nonhuman
	OS	IT	E
<b>EXPLANATION</b>			
<b>A. Association</b> ~ “There is something!”			
<i>c</i> is associated with <i>d</i>			
<i>c</i> is associated with <i>m</i>			
<i>m</i> is associated with <i>d</i>			
<b>B. Biology</b> ~ “It might work like this.”			
<i>c</i> can biomedically plausibly initiate <i>m</i>			
<i>m</i> can biomedically plausibly be a pathomechanism for <i>d</i>			
<b>PREDICTION</b>			
<b>C. Confirmation</b> ~ “It happens—as expected!”			
<i>m</i> occurs more frequently in the presence of <i>c</i> *			
<i>d</i> occurs more frequently in the presence of <i>c</i> *			
<i>d</i> occurs more frequently in the presence of <i>m</i> *			
<b>D. Difference</b> ~ “We can make a difference.”			
Eliminating or reducing <i>c</i> is associated with elimination or reduced occurrence of <i>d</i>		PT	
Interfering with <i>m</i> is associated with elimination or reduced occurrence of <i>d</i>		CT	
<i>Abbreviations: c: cause; m: pathomechanism; d: disease; OS: observational studies; IT: intervention trial; PT: prevention trial; CT: clinical trial; E: experiment.</i> <i>*More frequently than expected based on its occurrence in its absence.</i>			

Of note, I am not proposing that the other types of evidence in the JES would not benefit from confirmation as well. Experiments need confirmation by similar experiments in another laboratory, and clinical and community trials need confirmation by similar trials in different and larger populations. Even confirmatory type C evidence becomes stronger if confirmed independently. What is special about the confirmation of the initial type A evidence

(epidemiological observations) by type C (predictive) evidence is that the mere existence of a potential association needs to be ascertained.

#### [D] Difference

The litmus test for establishing the  $c > m > d$  process is successful prediction of intervention success. We must demonstrate that by intervening in controlled human research environments, such as clinical or population trials (showing efficacy), we can make a *difference*. We need to demonstrate that eliminating or reducing the occurrence of  $c$  eliminates or reduces  $m$ , and thereby  $d$ . Of note, I do not think that randomization of an intervention trial is the main justifier of health interventions. Instead, this is provided by the explanatory-predictive coherence of the error-independent components of a JES.

To establish type D evidence, we need to demonstrate that eliminating or reducing  $c$  is associated with elimination or reduced occurrence of  $m$ ; and interfering with  $m$  is associated with elimination or reduced occurrence of  $d$ .

Effective manipulation is part of components [B] and [D] in JES. Only if high-quality evidence of all four types is provided — in other words, if each one of the four types of evidence is plausibly supported by the other three, and none of the four contradicts any one of the other three — is the explanatory-predictive evidence coherent and the initiation of population health interventions fully justified. Of course, some interventions may be justified even if the JES is still incomplete, for example, if an intervention is desperately needed and preliminary evidence seems to suffice.

Consider the following brief example. Showing that (1) coffee consumption in pregnancy (risk factor  $r1$ ) is associated with low birthweight, and that (2) caffeine exposure in pregnant laboratory animals reduces the birthweight of mouse pups (by mechanism  $m1$ ) is not yet sufficient to fully justify population campaigns to reduce coffee consumption during pregnancy. Why not? It could very well be that another risk factor ( $r2$ ) for low birthweight, such as smoking, is strongly associated with coffee consumption *and* has a much stronger effect on birthweight via another mechanism ( $m2$ ) than does caffeine (via  $m1$ ). For example, if pregnant women tend to hide smoking behavior during pregnancy from health researchers, but tell the truth about their coffee intake, it would seem as if  $m1$  explains a stronger effect of  $r1$  on birthweight than it really has, because the strength of the association is really determined by both  $m1$  and  $m2$ . In such case, we would expect to be able to predict low birthweight based on knowledge about  $r1$  and  $m1$ , without even knowing about  $r2$  and  $m2$ , whose effect is screened off by (hidden behind)  $r1$  and  $m1$ .

However, any prediction based solely on such explanatory evidence of what would happen if we intervened on  $r1$  would probably be exaggerated, because such intervention would affect only  $m1$  and not  $m2$  (unless intervening on  $r1$  somehow also modifies exposure to  $r2$ ). We would probably be able to make a correct prediction about the occurrence of low birthweight based on [A] and [B], but we would likely be disappointed at step [D] because our intervention on (weak)  $r1$  might reduce the risk for low birthweight only minimally, while an intervention on (strong)  $r2$  would be much more effective. In this scenario, predictions [C] and [D] would yield incoherent

results. My argument is that only if they are coherent—in other words, only if [C] and [D] cohere with each other *and* with [A] and [B]—do the predictions fully justify intervention.

It is important to keep in mind that explanatory and predictive evidence are not the same. Although both kinds of evidence are based on the same data (for example, observational data from epidemiological studies), the mere observation of an association in a population is explanatory and, once observed again in another population, is confirmatory, but still explanatory as well. What makes such (subsequent) observation predictive evidence is not the repetition of an observation in a separate study (in what Hill called *consistency*), but the success of a prediction: explanatory evidence is data such as risk ratios or odds ratios and their confidence intervals, while predictive evidence is a simple yes/no answer to the question “Is our prediction accurate?”

Furthermore, the explanatory power—the oomph, thrust, vigor—of type A evidence lies in its capability to generate a causal hypothesis that explains the observed data. The explanatory power of type C evidence as standalone evidence would be the same as type A evidence, but given previous type A evidence, type C evidence also has additional confirmatory power that type A doesn't have. This confirmatory power derives from both the repeated observation—what Howick and colleagues (2009) call “parallel evidence”—*and* the success of the prediction we made based on type A evidence. In this sense, type A and C evidence are very different in the way they can be interpreted: type A evidence is explained by a causal hypothesis, while type C evidence is explained by both a causal hypothesis and by a hypothesis of accurate prediction, which turns out to be true.

While explanatory evidence may be pointing in the right direction and may, in some cases, even yield success once implemented, it is the added predictive evidence and the resulting coherence of items [C] and [D] that justifies health interventions. In other words, providing *both* explanatory and predictive evidence in support of a proposition to intervene offers better evidence than does a cogent explanation alone. But my argument goes one step further, proposing that successful prediction joins explanatory evidence in transforming merely cogent evidence into evidence that justifies interventions.

This point may be a departure from previous approaches that tend to separate explanation from confirmation. By integrating the explanation and confirmation, the overall explanatory power of the joint explanatory system of evidence is strengthened. At the point where all members of JES cohere, the *predictive* power of [C, D] adds to the *explanatory* power of [A, B] to form JES, whose joint explanatory power is greater than that of [A,B].

What makes me think that the joint explanatory power of JES goes beyond that of [A,B]? Consider, again, the notion of *evidential error independence*. As explained above, this notion holds that the joint evidential power of two pieces of evidence is greater if they are error independent than if they are not. Error independence means that if one of the two pieces of evidence turns out to be faulty, the other one does not automatically fail as well. For example, if one witness saw the gardener arrive at the house of the victim on the evening of the murder and also saw the gardener having a heated argument with the victim on the back porch of the victim's house an hour later, the explanatory power of these two pieces of evidence can usually be considered jointly stronger than either one alone. However, they might not be error independent,

since anything that limits the credibility of the witness (her having been inebriated that night as witnessed by a local bartender, or her being mentally ill and often confused if she forgets to take her medication, or her being the disgruntled ex-wife of the gardener) will limit the credibility of both pieces of evidence offered by her, not just one.

The notion that [A], [B], [C], and [D] are defined by their explanatory versus predictive characteristics and by their being gathered by observation versus intervention (Table 1) makes them error-independent, at least along these two dimensions. Thus, error independence increases the evidential vigor of JES = [A,B,C,D] beyond that provided by just [A,B].

In sum, in order to promote a successful health intervention, one needs (1) explanatory evidence [A,B], (2) strong individually successful predictors [C] and [D], which are (3) mutually coherent in their predictions and thus form a coherent jointly predictive set [C,D] that is (4) jointly better in its predictive success compared to individual predictions [C] and [D]. At this moment, the predictive evidential set [C,D] is deemed coherent with the explanatory evidential set [A,B], and the joint explanatory-predictive coherence of JES justifies the proposed intervention. The presence of *incoherent* evidence would suggest, first, that starting an intervention might not be a good idea and, second, that the validity of the available incoherent evidence should be scrutinized.

## Evidence Mapping

Available evidence supports mapping explanatory-predictive coherence claims. The result is an evidence map that systematically summarizes research data (Miake-Lye et al. 2016). Table 3 is a proposed rubric for evidence mapping to support the claim that a health intervention is justified. Specific selection criteria and quality control methods, as well as effective ways to visually represent the evidence, are still to be developed. The following description of what kind of documentation should be available in support of explanatory-predictive coherence claims may serve as a humble beginning, and I invite proposals how to improve the framework.

### Section A: Association

The first prerequisite for an intervention is having an observable target problem. This is often the identification of a health condition we want to alter, prevent, treat, and so forth. In part A of the map, I suggest documenting observational evidence from human individuals and populations in support of the notion that the problem at hand is real (A.1, *occurrence*). Simple prevalence and incidence data provide such information. In the story of AIDS, it was the publication of a case series of homosexual men with Kaposi's sarcoma and a rare form of pneumonia that helped identify AIDS as a population health problem (CDC 1982).

If we don't know *what* to change, it is very hard to come up with a good idea what to do to bring about change. In order to reduce the burden posed by the problem, we need something to intervene upon. Thus, we need to have some idea of what might be the causes of the condition of interest, or risk factors that we think of as initiators of the process that culminates in the condition. This is achieved by detailed epidemiological studies. In the history of AIDS research, scientists found not only a strong association between HIV infection and the subsequent

development of AIDS but acquired data “from prospective epidemiological studies that document an absolute requirement for HIV infection for the development of AIDS” (Blattner, Gallo, and Temin 1988). Evidence that certain candidate causes have been identified is documented in part A.2 of Table 1, *association*.

## Section B: Biology

The second section is to be populated by detailed information about the mechanics of the process that leads from cause to effect. For obvious reasons, pathomechanisms cannot be studied directly in humans. Thus, targeted laboratory experiments have to be designed and performed to provide type [B] evidence. By definition, such experiments come with two prominent drawbacks. First, the phenomena observed in such experiments are not occurring naturally, but only after some specific manipulation of the system parameters. Part of this caveat is that exactly what makes experiments credible as evidence for biological mechanisms (design, manipulation, controlled environment, randomization) makes them unnatural in the sense of not being a reflection of what happens to a person while she is getting sick. Second, the biological phenomena observed in the experimental setting are frequently from work done *in vitro* (“in glass,” test tube research) or *in vivo* (“in life,” animal research). Neither system is identical with the “human system.”

The focus on animal models comes from the intent to render experimental models as relevant to the (human) target condition as possible but introduces new layers of ethical and comparability concerns. One striking example is the wide-ranging incongruence between human and murine inflammatory responses that might have important consequences for the use of experimental mouse data to support claims about human inflammatory disease (Seok et al. 2013). References that elucidate the mechanism that explains the association between the purported cause(s) and the illness under investigation are documented in section B of our evidence map: type B.1 being evidence from *in vitro*, and type B.2 from *in vivo* studies. Either type can come without or with evidence from blocking experiments, while such evidence always strengthens experimental evidence.

Note that for the prevention of illness, knowledge of mechanisms (type [B] evidence) is sometimes not necessary. In the history of AIDS research, this concept was alluded to by Blattner, Gallo, and Temin in 1988, when they wrote that “Epidemiological data show that transmission of HIV results in AIDS and blocking HIV transmission prevents the occurrence of AIDS.” By “blocking HIV transmission,” the authors refer to public health interventions that interrupted HIV transmission via blood transfusion, which resulted in an reduction of the occurrence of AIDS among individuals who had received transfusions containing HIV (type [D] evidence).

**Table 3** *Proposed multi-evidence map for documentation of evidence needed to form*

*Target condition potential cause(s)*

Section	Evidence Type	Hypothesis supported	Evidence
<b>A</b>	<b>Observation</b>		
A.1	Occurrence	The problem occurs naturally in humans	
A.2	Association	One or more candidate causes have been identified in well-performed observational studies	
<b>B</b>	<b>Experiment</b>		
B.1.a	In vitro, effect	An in vitro laboratory model using the candidate cause(s) from A.2 induces characteristics of the target condition	
B.1.b	In vitro, effect blocked	Blocking the candidate causal mechanism in above model prevents characteristics of the target condition from occurring	
B.2.a	In vivo, effect	An in vivo laboratory model using the candidate cause(s) from A.2 induces characteristics of the target condition	
B.2.b	In vivo, effect blocked	Blocking the candidate causal mechanism in above model prevents characteristics of the target condition from occurring	
<b>C</b>	<b>Confirmation</b>		
C.1.a	Similar population/ scenario, same type of observation	The problem continues to occur as predicted by knowledge from type A evidence	
C.1.b	Similar population/ scenario, different type of observation	(as above)	
C.2.a	Different population/ scenario, different type of observation	(as above)	

**Table 3 (cont.)**

Section	Evidence Type	Hypothesis supported	Evidence
C.2.b	Different population/ scenario, same type of observation	(as above)	
<b>D</b>	<b>Difference</b>		
D.1	Clinical trial	An intervention targeting X turns out to be efficacious in clinical trial(s) (randomized or not)	
D.2	Population trial	An intervention targeting X turns out to be effective in population/community trial(s) (randomized or not)	
	<b>Additional Evidence</b>		
C*	Effectiveness	Population-level changes based on vital data	
C*	Effectiveness	Post-marketing research	
D*	Positive meta-analyses	E.G., Cochrane analyses suggest solid and significant efficacy	

Section C: Confirmation

Repeated observation of the occurrence of the phenomenon under investigation, both in terms of incidence and prevalence of the illness and in terms of observed recurrent association between purported causes and the illness, serve as confirmation that we have a potential target for intervention (documented in section C). The term *confirmation* here refers to the confirmation of the natural occurrence of the illness, not the confirmation of research studies or experiments as is standard operating procedure in science. The strength of the confirmatory evidence increases if we have such confirmation from multiple different populations and from multiple studies with different designs.

Section D: Difference

As discussed above, the scientific experiment is considered the litmus test for causal claims. What the lab experiment is for the explanatory axis, the intervention trial is for the predictive axis. By

*prediction* I do not refer to the simple fore-casting of events, but the *accurate* prediction of intervention success. Results from a successful intervention would be entered in section D.1 of the evidence map, and results from a successful population trial in section D.2.

As previously discussed, in order to justify interventions, at least three conditions should be fulfilled by the evidence collected in the evidence map: first, high-quality evidence of all four types should be provided; second, each one of the four types of evidence should be plausibly supported by the other three; and third, none of the four types of evidence should contradict any of the others.

The first requirement asks for high-quality evidence in all four areas. Here, the current scientific standard should guide judgment. In sections A and C, sufficient study size as well as control for confounding and other biases will be important. In section B, experiments should be controlled and blinded, and statistical procedures should be appropriate. In section D, published guidelines to rate the quality of the available evidence should be used, for example those from the GRADE initiative (Balshem et al. 2011).

The second condition is that the evidence from all four sections should be mutually supportive in logical and biological ways. For example, the initial observational evidence that HIV is associated with AIDS (Laurence et al. 1984) (section A) was supported by confirmation (section C) of an association between HIV and AIDS in multiple or even many different populations (Quinn, Zacarias, and St. John 1989), as well as by laboratory findings (section B) that a “virus-like infectious agent” transfected with DNA of an AIDS patient caused fatal wasting disease in monkeys, that zidovudine is an effective anti-HIV agent (Mitsuya et al. 1985), and by subsequent evidence that such intervention is effective in humans (section D) (Fischl et al. 1990).

Finally, any contradiction between the four types of evidence would reduce the coherence of the system. Of course, noncontradiction in itself isn't helpful, particular noncontradiction by the absence of evidence of a different type. For example, an observed association between a purported cause and some disease [A] would not be contradicted by the mere absence of supporting experimental evidence [B], but by the failure to cause the illness in the lab using the purported cause as the experimental intervention.

## **Conclusion**

In this paper, I have attempted to bring together the concepts of etiological explanation and explanatory-predictive coherence by offering a proposed outline for a multi-evidence map, in which evidence of association, biology, confirmation and difference making can be documented. Moreover, I have proposed that high-quality evidence that is mutually supportive and noncontradictory might be a good starting point for judging the degree of explanatory-predictive coherence. More work is needed to outline formal ways to extract evidence from the literature and organize it using evidence mapping to justify interventions in biomedicine and population health.

## References

- Balshem, H., et al. 2011. "GRADE Guidelines: 3. Rating the Quality of Evidence." *J Clin Epidemiol* 64 (4): 401–6.
- Blattner, W., R. C. Gallo, and H. M. Temin. 1988. "HIV Causes AIDS." *Science* 241 (4865): 515–16.
- Campaner, R., and M. C. Galavotti. 2012. "Evidence and the Assessment of Causal Relations in the Health Sciences." *Int Stud Philos Sci* 26 (1): 27–45.
- Campaner, R., and M. Cerri. 2020. "Manipulative Evidence and Medical Interventions: Some Qualifications." *Hist Phil Life Sci* 42, article 15.
- Centers for Disease Control and Prevention (CDC). 1982. "A Cluster of Kaposi's Sarcoma and Pneumocystis carinii Pneumonia Among Homosexual Male Residents of Los Angeles and Orange Counties, California." *MMWR Morb Mortal Wkly Rep* 31 (23): 305–7.
- Dammann, O. 2017. "The Etiological Stance: Explaining Illness Occurrence." *Perspect Biol Med* 60 (2): 151–65.
- Dammann, O. 2018a. "Data, Information, Evidence, and Knowledge: A Proposal for Health Informatics and Data Science." *Online J Public Health Inform* 10 (3): e224.
- Dammann, O. 2018b. "Hill's Heuristics and Explanatory Coherentism in Epidemiology." *Am J Epidemiol* 187 (1): 1–6.
- Dammann, O. 2020. *Etiological Explanations: Illness Causation Theory*. Boca Raton: CRC Press.
- Fischl, M. A., et al. 1990. "The Safety and Efficacy of Zidovudine (AZT) in the Treatment of Subjects with Mildly Symptomatic Human Immunodeficiency Virus Type 1 (HIV) Infection: A Double-Blind, Placebo-Controlled Trial. The AIDS Clinical Trials Group." *Ann Intern Med* 112 (10): 727–37.
- Gillies, D. 2019. *Causality, Probability, and Medicine*. London: Routledge.
- Glass, T. A., et al. 2013. "Causal Inference in Public Health." *Annu Rev Public Health* 34: 61–75.
- Hernán, M. A. 2004. "A Definition of Causal Effect for Epidemiological Research." *J Epidemiol Community Health* 58 (4): 265–71.
- Hill, A. B. 1965. "The Environment and Disease: Association or Causation?" *Proc R Soc Med* 58: 295–300

Howick, J., P. Glasziou, and J. K. Aronson. 2009. "The Evolution of Evidence Hierarchies: What Can Bradford Hill's 'Guidelines For Causation' Contribute?" *J R Soc Med* 102: 186–94.

Laurence, J., et al. 1984. "Lymphadenopathy-Associated Viral Antibody in AIDS: Immune Correlations and Definition of a Carrier State." *N Engl J Med* 311 (20): 1269–73.

MacMahon, B., and T. F. Pugh. 1970. *Epidemiology: Principles and Methods*. Boston: Little, Brown.

Miake-Lye, I. M., et al. 2016. "What Is an Evidence Map? A Systematic Review of Published Evidence Maps and Their Definitions, Methods, and Products." *Syst Rev* 5: 28.

Mitsuya, H., et al. 1985. "3'-Azido-3'-Deoxythymidine (BW A509U): An Antiviral Agent That Inhibits the Infectivity and Cytopathic Effect of Human T-Lymphotropic Virus Type III/Lymphadenopathy-Associated Virus in Vitro." *Proc Natl Acad Sci USA* 82 (20): 7096–7100.

Paneth, N., and M. Joyner. 2021. "The Use of Observational Research to Inform Clinical Practice." *J Clin Invest* 131 (2): e146392. DOI: 10.1172/JCI146392.

Poston, T. 2014. *Reason and Explanation: A Defense of Explanatory Coherentism*. London: Palgrave Macmillan.

Quinn, T. C., F. R. Zacarias, and R. K. St. John. 1989. "HIV and HTLV-I Infections in the Americas: A Regional Perspective." *Medicine (Baltimore)* 68 (4): 189–209.

Rescher, N. 1958. "On Prediction and Explanation." *Br J Philos Sci* 8 (32): 281–90.

Rubin, D. B. 1974. "Estimating Causal Effects of Treatments in Randomized and Non-randomized Studies." *J Educ Psychology* 66 (5): 688–701.

Russo, F., and J. Williamson. 2007. "Interpreting Causality in the Health Sciences." *Int Stud Philos Sci* 21 (2): 157–70.

Salmon, W. C. 1984. *Scientific Explanation and the Causal Structure of the World*. Princeton: Princeton University Press.

Seok, J., et al. 2013. "Genomic Responses in Mouse Models Poorly Mimic Human Inflammatory Diseases." *Proc Natl Acad Sci USA* 110 (9): 3507–12.

Spirtes, P., C. Glymour, and R. Scheines. 2000. *Causation, Prediction, and Search*. Cambridge: MIT Press.

Stegenga, J. 2009. "Robustness, Discordance, and Relevance." *Philos Sci* 76: 650–61.

Wright, L. 1976. *Teleological Explanations: An Etiological Analysis of Goals and Functions*. Berkeley: University of California Press.