Evidence, Illness, and Causation: An Epidemiological Perspective on the Russo-Williamson Thesis

Alex R. Fiorentino[1], *Olaf Dammann[1,2, 3]
1) Dept. of Public Health & Community Medicine, Tufts University Medical School, Boston, MA, U.S.A.
2) Philosophy Department, University of Johannesburg, Auckland Park 2006, South Africa
3) Dept. of Gynecology & Obstetrics, Hannover Medical School, Hannover, Germany
*Corresponding author
Olaf Dammann, MD, SM (Epi.)
Dept. of Public Health & Community Medicine,
Tufts University School of Medicine
Boston, MA 02111
Tel (+int)(617)636-3963
Fax (+int)(617)636-6635
Email olaf.dammann@tufts.edu

Abstract

According to the Russo-Williamson thesis, causal claims in the health sciences need to be supported by both difference-making and mechanistic evidence. In this article, we attempt to determine whether Evidence-based Medicine (EBM) can be improved through the consideration of mechanistic evidence. We discuss the practical composition and function of each RWT evidence type and propose that exposure-outcome evidence (previously known as difference-making evidence) provides associations that can *be explained* through a hypothesis of causation, while mechanistic evidence provides finer-grained associations and knowledge of entities that ultimately *explains* a causal hypothesis. We suggest that mechanistic evidence holds untapped potential to add value to the assessment of evidence quality in EBM and propose initial recommendations for the integration of mechanistic and exposure-outcome evidence to improve EBM by robustly leveraging available evidence in support of good medical decisions.

1. Introduction

Evidence-based medicine (EBM) has become the predominant paradigm of medical decision-making. It is described by its founders as "the conscientious, explicit, and judicious use of current best evidence in making decisions about the care of individual patients," (Sackett 1995). This definition inevitably raises the question: what constitutes the current best evidence? Surely many approaches to medical decision-making aspire to use the best evidence; EBM differentiates itself by weighing evidence in a particular way. Specifically, EBM is founded on the belief that information derived from clinical experience or from knowledge of basic biological mechanisms can often be misleading or incorrect if not supplemented by systematic observations from clinical research (Evidence-Based Medicine Working 1992). EBM's proponents therefore generally count randomized trials and systematic reviews of randomized trials as the most trustworthy sources of evidence, with non-randomized cohort studies, case-control studies, and case series occupying the middle ground, and mechanistic reasoning, clinical experience, and expert opinion forming the least trustworthy sources of evidence (if they are included at all) (Force 2008, Group 2011).

EBM's rise in popularity has been accompanied by significant criticism, not least from philosophers of science, for whom the theory of EBM touches upon many familiar topics, including the logic of evidence, causal inference, mechanism, and explanation. Some philosophers have questioned whether EBM's preferred prioritization of evidence sources is justified (Clarke, Gillies et al. 2013), others have asked whether evidence even *can* be prioritized by source (Worrall 2010), and still others have largely defended EBM's general principles while offering resolutions to specific issues and paradoxes (Howick 2011). One particular criticism, which we will explore here, suggests that EBM undervalues "non-statistical evidence of mechanisms" (Clarke, Gillies et al. 2013). The philosophical underpinning for this criticism comes from the Russo-Williamson Thesis (RWT).

The RWT, first published in 2007, proposes that in order to establish a causal claim in the health sciences, one needs two types of evidence: *difference-making evidence* that the putative cause and effect are correlated and *mechanistic evidence* that there exists some mechanism explaining the difference-making relationship (Russo and Williamson 2007, Clarke, Gillies et al. 2013). Although this is certainly a helpful guideline for thinking about causal inference, we hold that if the RWT is to be of use to medical decision-making, it requires clarification beyond previous disambiguations (Illari 2011) and criticisms (Claveau 2012). Moreover, we do not subscribe to the view that knowledge of both difference-making and mechanistic evidence is <u>necessary</u> for the establishment of causal claims in the health sciences, although it may very well be (and often is) <u>sufficient</u> for it (Broadbent 2011).

Here, we approach the RWT from the epidemiologist's point of view. Modern epidemiology is the design, implementation, and analysis of studies in humans that help identify (a) risk factors for certain diseases (so-called observational studies) and (b) interventions that reduce disease burden or prevent illness altogether (mainly, but not exclusively, randomized controlled-trials). Therefore, all of the health sciences draw upon information generated by epidemiological research. We believe that the epidemiological perspective we provide in this paper is therefore useful for examining the relationship between the concepts of the RWT and the day-to-day work of evidence gathering and utilization across the health sciences.

Our goal is to determine whether EBM can be improved through the consideration of mechanistic evidence. In order to answer this question, we first attempt to clarify the practical composition and function of each RWT evidence type. We find that exposure-outcome evidence (previously known as difference-making evidence) provides associations that can *be explained* through a hypothesis of causation, while mechanistic evidence provides finer-grained associations and knowledge of entities that ultimately *explains* a causal hypothesis. In light of these clarifications, we find that mechanistic evidence holds untapped potential to add value to the assessment of evidence quality in EBM. Finally, we suggest a path forward, providing initial recommendations for the integration of mechanistic and exposure-outcome evidence to generate an EBM that more robustly leverages the available evidence to make good medical decisions.

2. Illness causation viewed from the epidemiological angle

One might be tempted to think that in order to discuss causal inference in epidemiology, one should be clear about how epidemiologists define causation. Another possible position is that we do not need to know about the nature of a cause in order to find one. In Broadbent's words, "we can know that something exists without knowing what it is" (Broadbent 2011:59).

However, if asked which ontological position they take when thinking about causation, mechanistic or difference-making, most epidemiologists will probably agree that their notions of illness causation processes seem to fit current philosophical definitions of mechanism quite well (Glennan 1996, Machamer, Darden et al. 2000, Bechtel and Abrahamsen 2005, Illari and Williamson 2012). One particularly attractive conception of mechanism is Garson's "functional sense of mechanism" (Garson 2013). According to this view, mechanisms serve functions and illness causation can be conceptualized as the aberrant function that is due to disrupted mechanisms. Using biomedical terminology, this is the process of moving from the physiological to the pathophysiological state by means of pathomechanisms. The fact that "pathomechanism" is frequently used in biomedicine suggests that Garson's view that "there are no mechanisms for pathology; pathologies result from disrupting mechanisms for functions" (Garson 2013:317) might not be shared by at least some biomedical scientists.

Unfortunately, pathomechanisms cannot be directly observed by epidemiologists. Indeed, the intra-individual component of the illness causation process (pathogenesis) requires biological studies in model organisms. Epidemiologists can examine how certain antecedent variables act and interact in predicting health outcomes in populations (etiology). The etiology of illness, therefore, can be viewed as the natural history of disease from exposure to risk factors until the illness becomes clinically detectable. As such, disease etiology includes the pathogenesis as the disease mechanism that occupies the black box between exposures and outcomes, though epidemiologists can discuss pathogenesis itself only in rather limited ways. Nevertheless, specialized forms of epidemiologic inquiry such as molecular and genetic epidemiology can make quite interesting contributions to the interdisciplinary discussion of pathomechanisms.

Still, even without exact pathogenetic (mechanistic) knowledge, information about difference-making evidence (we prefer "exposure-outcome evidence", v.i.) can be very useful. This approach has been called the "black box stance", which holds that "epidemiologists need not concern themselves with the discovery of mechanisms, but can directly attack causal questions without worrying about the mechanisms underlying the hypotheses they generate" (Broadbent 2011:60). Historical case studies of successful prevention of childbed fever by hand disinfection and of lung disease by smoking cessation support the notion that "laboring to uncover mechanisms may well prove to be a waste of time and money, from a public health point of view" (ibid).

Most activities in the health sciences are targeted at prevention and health promotion in public health and about treatment of illness by intervention in medicine. Our view of etiology and pathogenesis outlined above suggests that although epidemiologists might be interested in pathomechanisms, they can live without them while still making a contribution to public health by generating etiologic knowledge. Biological bench scientists are better equipped to focus on pathomechanisms that can be targeted with pharmacological interventions.

3. Redefining Difference-Making Evidence for the Health Sciences

Difference-making evidence in support of a given causal claim in the health sciences has been defined as evidence that a cause "makes a difference to" an effect (Russo and Williamson 2007) and "evidence that the effect does indeed vary with the postulated cause" (Illari 2011). The trouble with the term "difference-making" is that it closely resembles the notion of causation, simply because of the similarity between the verbs "making" and "causing". If the statement "A makes a difference to B" is taken as equivalent to the statement "A causes B," then evidence of difference-making is indistinguishable from evidence of causation generally. Obviously, this is not the meaning that Russo and Williamson intended, because it would contradict their notion that evidence of difference-making does not suffice to support causal claims.

For epidemiologists, evidence of difference-making is the result of a two-step discovery process in which evidence of (statistical) association between exposure and outcome is gathered first in order to establish that it is possible that the exposure causes the outcome. This evidence is then supplemented with evidence for possible mechanisms by which the exposure can cause the outcome. Together, both kinds of evidence are sufficient to support a causal claim. Thus, from the epidemiological perspective, evidence for difference-making comprises two kinds of evidence that are very close to the two kinds of evidence required by the RWT. We, therefore, see two differences between the RWT and the epidemiological perspective. First, for epidemiologists, finding evidence for difference-making requires both evidence of association and evidence of a mechanism, which is the same as satisfying both requirements of the RWT, one of which is called "difference-making", which is a misnomer from our epidemiological perspective. Second, epidemiologists, and arguably other health scientists, consider evidence for difference-making sufficient for causal inference, while Russo and Williamson seem to seem to consider their two kinds of evidence necessary for it.

If we are to adopt a more conservative definition of difference-making evidence that does not presuppose causation, then difference-making evidence begins to resemble a concept familiar to epidemiologists everywhere: statistical association. An association, reported in the epidemiologic literature using relative risk (RR), odds ratio (OR), attributable risk (AR), etc., implies that two variables move together, but not necessarily that movement of one variable causes the movement of the other. Indeed, we agree with Alex Broadbent that "the mathematical definition of (the numerical values of RRs, ORs, and the like) do not tell us how to understand the causal import of a measure of association" (Broadbent 2013:30). Causation is, in fact, only one of several possible interpretations of an association. An observed association between *A* and *B* might have arisen

because *A* caused *B*, but it might also be that *B* caused *A*, or *A* and *B* shared a common cause, or some bias within the study caused *A* and *B* to appear associated, or even that *A* and *B* moved together by chance. Without further information, we do not know which of these explanations of the association is correct.

The idea that difference-making evidence refers to associations is not ours. Claveau notes that "difference-making evidence […] amounts to the statistical association between the two relata" (Claveau 2012:809). Russo and Illari confirm that "evidence of difference-making usually consists of correlations, associations, and dependencies in statistical analyses" (Illari and Russo 2014:55), but add that it "can also comprise counterfactuals (based, for example, on established medical knowledge of case reports (Clarke, 2011, sec. 4.5)". Indeed, Clarke suggests that in some cases of medical inquiry into disease causation difference-making is non-statistical (or non-probabilistic), because in some such scenarios only very few individuals with a certain disorder exist and the association between causative characteristics (e.g., genotype) and phenotype is due to a quasi-deterministic relationship between genotype and phenotype where nearly all individuals with a genetic abnormality will have the clinical disease features of interest and where it is impossible to have the phenotype without having the genetic abnormality (Clarke 2011:99). Still, we think that even in situations where "association" is used to denote the co-occurrence of characteristics in individuals and not in a statistical sense, the main evidence-maker is the fact that such co-occurrence is identified in multiple individuals, not just one.

One might argue that study design and analysis techniques provide additional information that, together with the association, suggests causation. For example, an epidemiologic study might be designed to ensure that the observed exposure predates the observed outcome; it might adjust for known confounders (common causes); and it might include statistical tests that indicate the likelihood of the association arising by chance. These aspects of study design and analysis narrow the parameters of the association, and in so doing reduce the likelihood of some explanations of the association and raise the likelihood of others. They do not, however, change the fact that the thing being directly evidenced is an association.

The issues discussed above lead us to suggest abandoning the term "evidence of difference-making" in favor of "*exposure-outcome evidence,*" which we define as follows:

Exposure-outcome evidence is evidence from observational (cohort, case-control) or interventional studies (randomized controlled trials) that an exposure (the putative cause) and an outcome (the putative effect) are associated in populations and/or co-occur in individuals more frequently than expected by chance, for which causation is one biologically possible and scientifically plausible explanation.

We are aware that this brings us back to a probabilistic, not-necessarily-causal interpretation of such evidence, somewhat reversing Illari's and Russo/Williamson's change of terminology from "probabilistic evidence" (Russo and Williamson 2007) to "evidence of difference-making" (Russo and Williamson 2010). However, we find this to be the most accurate way to describe exactly what is being evidenced. Further, this term emphasizes that the variables involved in the association exist at the organizational level of exposure and outcome, which will prove important for distinguishing exposure-outcome evidence from mechanistic evidence.

4. Evaluating Exposure-Outcome Evidence

As we have attempted to illustrate, not all perceived exposure-outcome evidence actually arises from an exposure causing an outcome. Causation is one possible explanation for such evidence, but confounding, chance, and bias represent competing explanations. In order to determine the causal import of a piece of exposure-outcome evidence, we must weigh the strength of each of the above explanations and infer causation only when it is the best explanation, a process fittingly deemed "inference to the best explanation" (Lipton 1991). If we wish to use such causal inference to inform a future decision, as in EBM, we must further assess whether the causation we infer to have occurred in an observed group of people is likely to recur in some other group of people concerned in our future decision.

This leads to a pair of practical questions: 1) How do we assess the strength of the causal explanation for an observed association? 2) How do we assess whether a causal inference from one population is applicable to a distinct population of interest? These questions have been of great interest to epidemiologists generally, and were perhaps most famously addressed by Sir Austin Bradford Hill, who in 1965 proposed nine "aspects" of an association that one should "especially consider before deciding that the most likely interpretation of it is causation" (Hill 1965). Drawing upon Hill's work, as well as characteristics in support of a causal claim previously published by one of us (Dammann and Leviton 2007), we suggest that any assessment of exposure-outcome evidence should examine at least the following five characteristics:

1. The *magnitude* of the association, i.e., the degree to which the variables move together, as measured by the numerical value of the relative risk, odds ratio, attributable risk, or other measure of association;
2. its *statistical significance*, i.e. the likelihood that the observed association was not due to chance, as measured by a confidence interval or *p*-value;
3. its *temporality*, i.e. the ability of the study design to demonstrate that the exposure preceded the proposed outcome;
4. the extent to which the study design and analysis provides evidence that the observed association is not the result of known *confounders or other biases*;
5. the *relevance* of the observed association, i.e. the similarity of the population, conditions, and variables of the observed association to the population, conditions, and variables concerned in the causal claim.

Though the above characteristics help to inform causal inference from exposure-outcome evidence, the key concept of the RWT is that causal inference should be informed by *both* exposure-outcome and mechanistic evidence. In keeping with the RWT, we suggest that the strength of a causal claim should depend not only upon how well the causation hypothesis *explains* existing exposure-outcome evidence, but also upon how well the causation hypothesis can *be explained by* a plausible mechanism. We explore this topic below.

5. Rethinking Mechanistic Evidence for the Health Sciences

Mechanistic evidence is perhaps an even trickier concept to pin down than exposure-outcome evidence. One commonly cited definition of mechanism is that of Machamer, Darden, and Craver: "Mechanisms are entities and activities organized such that they are productive of regular changes from start or set-up to finish or termination conditions" (Machamer, Darden et al. 2000). Stuart Glennan's is: "a complex system that produces [a] behavior by the interaction of a number of parts, where the interactions between parts can be characterized by direct, invariant, change-relating generalizations" (Glennan 2002).

Both of these definitions, like the definition of difference-making evidence, rely on causal language: Machamer et al. state that mechanistic entities and activities "are productive of regular changes," which we take to mean that they *cause* regular changes. Glennan's idea of a mechanism as a "complex system that produces [a] behavior" seems indistinguishable from a complex system that *causes* a behavior. As Paul Thagard states, "Knowledge of mechanisms is itself knowledge of causally related events that must have somehow been previously acquired" ((Thagard 1999)p109).

Because the concept of mechanism is intertwined with the notion of causality, mechanistic evidence faces the same major challenge as exposure-outcome evidence: causation cannot be observed and must be inferred from indirect evidence. What, then, does this indirect evidence of mechanisms look like? We propose that mechanistic evidence in the health sciences consists of at least three distinct sub-types:

- ***Entity-based mechanistic evidence*** provides information about the existence and structural organization of entities hypothesized to be involved in a causal relationship.

- ***Association-based mechanistic evidence*** demonstrates associations between entities that, if causally interpreted, suggest a causal pathway connecting exposure to outcome.

- ***Activity-based mechanistic evidence*** provides direct evidence of the activities that entities participate in, without requiring the inference of these activities from observed associations.

Entity-based evidence is vital for the discovery of how entities and their components might be involved in causal processes. One example is the application of knowledge based on Rosalind Franklin's x-ray crystallography images of DNA (Franklin and Gosling 1953) by Watson and Crick, who were famously able to leverage structural knowledge of DNA to make several accurate predictions about the mechanisms of DNA replication (Watson and Crick 1953). For certain causal processes, namely those in which known laws of physics and chemistry are expected to govern system behavior in simple, predictable ways, entity-based evidence may even be sufficient for inferring mechanisms. In the complex systems of the health sciences, however, such predictable scenarios are quite rare. Neither the overall effects of a drug upon a patient nor the totality of the mechanisms that generate its effects can be predicted with satisfactory confidence from its chemical structure and knowledge of physics. Accordingly, Watson and Crick cautioned in 1953 that their proposed mechanisms of DNA replication "must be regarded as speculative." For the vast majority of health sciences mechanisms, entity-based evidence must be accompanied by association-based evidence in order to demonstrate with any confidence that a given mechanism exists.

For an example of association-based mechanistic evidence, one needs look no further than the line graphs, bar graphs, dot plots, and other graphical data visualizations present in virtually any medical or biological journal article. For example, let us examine the content of the article, "Meteorin-like is a hormone that regulates immune-adipose interactions to increase beige fat thermogenesis," published in the journal *Cell* in 2014 (Rao, Long et al. 2014). This article reports the identification of a novel hormone, Metrnl, and finds that Metrnl, among other functions, "stimulates energy expenditure and improves glucose tolerance." The identification of this hormone represents entity-based evidence: it describes a previously unknown molecule and its spatial distribution as a circulating factor within the blood. The remainder of the article consists largely of association-based evidence aimed at demonstrating the (causal) processes in which Metrnl is involved.

The article contains 7 figures (excluding supplemental material), which are divided into 50 sub-figures (1A, 1B, 1C, etc.). Of these 50 sub-figures, 44 contain a bar graph, a line graph, or a dot plot. Each of these 44 graphs pools data from repeated statistical trials to demonstrate that a difference in one variable is associated with a difference in another variable. Figure 5F, for example, demonstrates that when mice are experimentally induced to express Metrnl (vs. a control gene), the Metrnl-expressing mice tend to express the inflammatory molecules interleukin (IL)4 and IL13 at higher levels. Viewed independently, this association could be interpreted as exposure-outcome evidence for the causal claim that increased Metrnl expression causes increased IL4/IL13 expression. This exact evidence, however, could also be applied as association-based mechanistic evidence in support of the broader causal claim that increased Metrnl expression stimulates energy expenditure. In this context, the putative causal relationship between Metrnl and IL4/IL13 represents one step in a mechanistic pathway that explains how Metrnl might cause increased energy expenditure. Each of the steps in this putative pathway is supported through observed associations (Figure 1).

We believe that activity-based mechanistic evidence – evidence in which activities are observed directly – is relatively uncommon in the health sciences. Its prevalence may be growing, however due to the use of technologies that enable continuous monitoring of biological processes. One recent example of activity-based mechanistic evidence comes from the article, "Stage Specific Assessment of *Candida albicans* Phagocytosis by Macrophages Identifies Cell Wall Composition and Morphogenesis as Key Determinants," published in *PLoS Pathogens (Lewis, Bain et al. 2012)*. Supporting video S3 from this article uses video microscopy to show a macrophage moving toward and subsequently engulfing a *C. albicans* yeast cell. This video demonstrates not only the spatial and temporal organization of the macrophage and yeast cell (as would be observed in entity-based evidence) but also directly demonstrates what these entities are doing over time. Information about these activities need not be inferred (which would be association-based evidence); instead, activities comprise an inherent part of the evidence itself. Interestingly, the authors also report a variety of associations involving the activities observed in the video. For example, they report that when they repeated the experiment using yeast with mutated glycosylation genes, "visual inspection of the videos suggested enhanced macrophage migration towards *C. albicans* glycosylation mutants…compared to wildtype control." This demonstrates that characteristics of activities themselves (such as macrophage migration) can serve as variables in biological associations.

Entity-based and association-based mechanistic evidence are both crucial to the discovery and validation of the mechanisms that govern causal relationships. Entity-based mechanistic evidence identifies players that might be involved in mediating a causal relationship and describes the structure and organization of these entities, thereby suggesting hypothetical mechanisms. Association-based mechanistic evidence supports (or refutes) hypothesized mechanisms through the observation of associations between entities, which may permit causal inference at the level of individual mechanistic steps. Association-based mechanistic evidence is in this way identical to exposure-outcome evidence except that it concerns a single causal mechanistic step within a broader causal claim, rather than the broader causal claim itself. This distinction is determined purely by the particular causal claim that is of interest to the claimant. Activity-based evidence, where available, can also support the existence of hypothesized mechanisms.

6. Evaluating Mechanistic Evidence

As with exposure-outcome evidence, not all mechanistic evidence provides equal support for causal inference. As Jeremy Howick notes, many "bogus" mechanistic arguments have been used to justify

harmful or useless interventions throughout the history of medicine, including relatively recent cases such as the use of antiarrhythmic drugs to treat ventricular arrhythmia, which may have caused thousands of avoidable deaths ((Howick 2011) Chapter 10). To avoid such harmful decisions, it is critical that we distinguish the mechanistic evidence that provides strong evidence of causality from that which provides weak evidence, or no evidence. We propose that this assessment entails two key questions: (1) Can we reasonably infer that each individual step within the proposed mechanistic pathway represents a real causal relationship? (2) Do these individual mechanistic steps sum to form a continuous and reasonable mechanistic pathway?

In inferring that a causal mechanistic step does indeed occur, we are inferring that the inputs to the mechanism cause the outputs. The evidence we can use to assess this inference consists of entity-based evidence, association-based evidence, and, where possible, activity-based evidence. We will consider the evaluation of these evidence types sequentially.

Entity-Based Evidence
Assessment of entity-based mechanistic evidence should in theory be relatively straightforward, since it is limited to information about the existence and static organization of entities, and is not concerned with processes involving change over time. As a starting place, we suggest assessment of entity-based evidence should revolve around the following questions:

1. ***What are the entities of the proposed mechanism?*** These might include hormones, receptors, enzymes, entire organs, and other physical entities, as well as social and psychological constructs, depending on the causal claim of interest.

2. ***How convincing is the evidence that these entities exist within the system mediating the causal claim of interest?*** For entities discovered in model systems and organisms, additional evidence may be necessary to demonstrate that the entities are also present within the human body systems of interest and within the disease states of interest.

3. ***How convincing is the evidence for the structure and spatial arrangement of these entities, if these characteristics are important to the function of the proposed mechanistic pathway?***

Association-Based Evidence
Since association-based mechanistic evidence and exposure-outcome evidence both rely on observed associations to make causal inferences about the variables concerned in the association, these two evidence types can be evaluated in much the same way. The strength of the association-based evidence for each step in a mechanistic pathway will thus depend, as in section 3, on the direction, magnitude, statistical significance, temporality, and relevance of the observed association, as well as the degree of confidence that it is not due to confounding or other biases.

Activity-Based Evidence
As we have noted, activity-based evidence is likely to be less common in the health sciences than the first two types of mechanistic evidence. When it is available, we propose it can be evaluated roughly according to the following considerations:

1. ***Is the activity reproducible?*** This includes the consideration of the statistical significance of associations, since it addresses the likelihood that the activity was observed simply by chance.

9

2. ***Is it reasonable to attribute the perceived activity to the available evidence?*** It is desirable to avoid unjustified assumptions about the nature of an observed activity. For example, if we are concluding that the aforementioned video provides evidence of the activity of a macrophage engulfing a yeast cell, we should evaluate how certain we are that the macrophage is engulfing the yeast cell, rather than simply moving between the yeast cell and the microscope lens.

3. ***How convincing is the evidence that this activity takes place within the system mediating the causal claim of interest?***

In addition to the strength of its individual pieces of evidence, the overall strength of a proposed mechanistic pathway from cause to effect depends upon the *completeness* of the pathway, i.e. whether the existing association-based evidence forms a continuous inferential chain from exposure to outcome without any "obvious gaps" (see Howick (Howick 2011), p144). If we have observed an association between exposure *E* and proposed mediator *M*, but no association between *M* and the outcome *O*, then we do not know whether our proposed mechanism could actually connect *E* to *O* (Figure 2). This partial mechanistic evidence is clearly much weaker than it would be if we had also observed an *M-O* association.

Depending how successfully individual mechanistic associations and an overall proposed mechanistic pathway address the above issues, they may constitute no evidence, weak evidence, or strong evidence for a given causal claim.

7. Debunking Some (Tempting) Assumptions about Evidence Types

Assumption #1
Mechanistic evidence comes from laboratory research, while exposure-outcome evidence comes from epidemiologic research.

As noted above, a single association will often serve as exposure-outcome evidence for one causal claim and association-based mechanistic evidence for a broader causal claim. For example, an observed association between high dietary LDL cholesterol and the formation of atherosclerotic plaques would serve as exposure-outcome evidence with respect to the claim that high dietary LDL causes atherosclerosis, but as mechanistic evidence with respect to the claim that high dietary LDL causes premature death. If any given association can constitute either exposure-outcome evidence or mechanistic evidence, it follows that epidemiologic research and laboratory research are both capable of generating either evidence type, albeit for different causal claims.

Where laboratory and epidemiologic research differ is in the organizational level at which they are comfortable collecting information. Laboratory research generally examines variables at the level of the person and below (including the organ, the cell, the molecule, etc.) – a territory we call the *pathogenetic level* of disease (Figure 3). Epidemiology, with the exception of molecular epidemiology, generally examines variables at the level of the person and above (including the group, the society, etc (Stallones 1980). Disease causation, however, can involve any organizational level, so the territory we define as the *etiologic level* of disease spans the entire spectrum of levels.

Assumption #2
Mechanistic evidence deals exclusively with variables at a lower organizational level than the hypothesized cause and effect variables it connects.

The very word "mechanism" seems to carry with it an image of mechanical parts, smaller than the whole, which relate the cause to the effect. As a result, it is tempting to think that mechanistic evidence in the health sciences lives only at lower levels of organization than the causal claim of interest. Indeed, of the three biological meanings of "mechanism" described by Daniel Nicholson, two ("mechanicism" and "machine mechanism") deal only with components or internal workings (Nicholson 2012).

In the context of medicine, however, the biological, psychological, and sociological processes that lead to health or disease do not function as isolated machines. These processes may be influenced by their context as well as their components, and so the entities involved in disease causation may be both large and small. As Illari and Williamson note, "mechanistic explanations use parts organized to produce the phenomenon, and may also look to containing systems to set the phenomenon" (Illari and Williamson 2011:827). This idea is also discussed at length in the social sciences literature; social psychologist J. Richard Hackman writes: "…moving up one level of analysis can add at least as much explanatory power, and sometimes more, as moving down a level" (Hackman 2003).

Therefore, compared with the exposure and outcome variables, the mediators of a causal relationship may be variables at a lower level, higher, or identical organizational level, or (most likely) some combination thereof. For example, the putative causal pathway that links taking an antidepressant medication (a person-level variable) to an improved self-described level of happiness (a person-level variable) may well involve mediating variables including serotonin reuptake rates (a cell-level variable) and changes in family dynamic (a group-level variable). This aligns most closely with Nicholson's third meaning of mechanism, "causal mechanism," which avoids any reference to internal workings or components: "A step-by-step explanation of the mode of operation of a causal process that gives rise to a phenomenon of interest" (Nicholson 2012).

8. The RWT and EBM
In light of the above clarifications of the RWT, we now return to the topic of evidence-based medicine and the question: can the RWT be used to improve EBM? As a starting place, let us reconsider the critique offered by Clarke et al.: that EBM undervalues "non-statistical evidence of mechanisms" (Clarke, Gillies et al. 2013). The phrasing of this critique is inconsistent with our view of health sciences evidence, since we believe that both exposure-outcome relationships and mechanisms in medicine and public health are often inferred from statistical associations between variables. To overcome this issue, we will rephrase the critique to say that EBM undervalues mechanistic evidence as defined in section 4 of this paper. In other words, EBM undervalues the sort of (mechanistic) evidence that *explains* how a given causal relationship might come about, and focuses too heavily on the sort of (exposure-outcome) evidence that *is best explained by* a causal relationship.

We believe this critique is justified with respect to the existing EBM evidence hierarchies. Our reasoning, previously expressed in one form or another by several authors on this topic (Illari 2011, Clarke, Gillies et al. 2013), is that mechanistic evidence and exposure evidence tend to have different weaknesses, so by combining both evidence types, we reduce the probability of any single weakness leading us into a spurious causal inference.

Phyllis Illari explains this error independence by stating that the major weakness of difference-making (exposure-outcome) evidence is confounding, while the major weakness of mechanistic evidence is "masking" – a scenario in which other, unseen mechanistic pathways between exposure

and outcome could in reality overwhelm the operation of the observed mechanistic pathway (Illari 2011). In her view, "each [evidence type] addresses the major weakness of the other evidence."

We generally agree with Illari's assessment, but would suggest a slightly more general interpretation. In our view, the major weakness of exposure-outcome evidence is not simply confounding, but is the possibility that any non-causal explanation of the evidence (including confounding, chance, and bias) is the correct explanation. We call this weakness the *non-causal possibility*. In fact, the non-causal possibility is not only a weakness of exposure-outcome evidence, but also applies to association-based mechanistic evidence, which might similarly arise through causation, confounding, chance, or bias. Mechanistic evidence therefore has two weaknesses. The first is the non-causal possibility for each of its proposed mechanistic steps, and the second is the possibility that unseen variables and mechanisms may make the net relationship between exposure and outcome different from what we predict based on our mechanistic evidence. These unseen mechanisms could exist along our proposed pathway, or along alternate pathways between exposure and outcome (in which case we have masking). We call this general weakness the *missing mechanism possibility*.

These adjustments do not undermine Illari's assertion that mechanistic and exposure-outcome evidence address each other's weaknesses. By adding strong mechanistic evidence to existing exposure-outcome evidence, we reduce the non-causal possibility, since the mechanistic evidence explains how the exposure-outcome evidence could have arisen through causation. By adding strong exposure-outcome evidence to existing mechanistic evidence, we reduce the missing mechanism possibility, since the exposure-outcome evidence demonstrates the net association between the variables. Therefore, by considering both evidence types, we can reduce our chances of making a spurious causal inference.

9. Mechanistic Evidence in Practice

Although we find mechanistic evidence to be undervalued in EBM hierarchies, it is not necessarily the case that mechanistic evidence is entirely ignored in the medical decisions that are guided by EBM. Though it is relegated to the bottom of many evidence pyramids, in the current era of research, mechanistic evidence is often a precursor of the human experimental studies that occupy the top of the pyramid. This is perhaps most obvious in the development of new pharmacologic interventions, in which mechanistic evidence about disease causation drives drug discovery by suggesting mechanisms upon which a hypothetical drug could act.

For example, the first three drugs approved by the U.S. Food and Drug Administration in 2014 were dapagliflozin, tasimelteon, and elosulfase alfa (Administration 2014). Dapagliflozin (for treatment of type 2 diabetes) was discovered through a search for compounds that selectively inhibit SGLT2, the glucose transporter that largely drives reabsorption of glucose in the kidneys (Meng, Ellsworth et al. 2008). Tasimelteon (for treatment of sleep-wake disorder in blind individuals) was discovered as a novel agonist of human melatonin 1a and 1b receptors (Vachharajani, Yeleswaram et al. 2003). Elosulfase alfa (for treatment of Morquio A syndrome, an inherited metabolic disorder) is a form of the human enzyme N-acetylgalactosamine-6-sulfatase (the enzyme that is deficient in individuals with Morquio A syndrome) that is manufactured through recombinant DNA technology (Sanford and Lo 2014). For all three of these interventions, existing evidence about the mechanisms of disease suggested the mechanism of a hypothetical drug, and this mechanistic reasoning drove the development of the drug before any exposure-outcome evidence could be collected.

Of course, many medical interventions do not originate from mechanistic reasoning. Further, even for those interventions based on mechanistic evidence, the strength of the mechanistic evidence might vary quite significantly. Therefore, we believe the quality of EBM's decisions would be enhanced if it were to consider mechanistic evidence explicitly, and alongside exposure-outcome evidence, in making decisions based on causal inference.

10. Conclusion: Moving EBM Forward

We have sought here to clarify the Russo-Williamson Thesis by interpreting its evidence types through the pragmatic lens of epidemiology. From this perspective, we find that both mechanistic evidence and exposure-outcome evidence are largely comprised of observed associations between variables, and this leads to further clarifications about the meanings of these evidence types and how each should be evaluated. In the end, however, our primary interest as epidemiologists is the improvement of health outcomes. Therefore, we ask whether, in light of our interpretation, the ideas of the RWT can be leveraged to improve the medical decisions flowing out of Evidence-Based Medicine.

As we have described above, the complementarity of each evidence type's weaknesses means that consideration of both evidence types reduces the chance of making a spurious causal inference. From our perspective, this is a good thing for EBM, as it potentially reduces the chance of making a medical decision that is harmful or not helpful to the patient. Therefore, we support the notion that EBM should more explicitly consider available mechanistic evidence as complementary to (rather than far inferior to) the type of exposure-outcome evidence gained from a randomized controlled trial examining patient outcomes with and without the intervention of interest.

Of course, this sort of vague recommendation is of little direct use to EBM, which is a highly structured and formalized system. Therefore, we offer below a few specific recommendations for the consideration of both mechanistic and exposure-outcome evidence within the world of EBM. These recommendations will require refinement and formalization by others within the EBM field, so our recommendations here should serve only as a starting place for further work toward the practical improvement of EBM.

- ***EBM should systematically review the literature for mechanistic evidence relevant to a causal claim.***

EBM relies heavily on systematic reviews, which "seek to collate all evidence that fits pre-specified eligibility criteria in order to address a specific research question" [Cochrane Handbook]. Such systematic reviews increasingly serve as a major point-of-care reference for clinicians. However, the "pre-specified eligibility criteria" alluded to above generally ensure that only exposure-outcome type evidence is included in systematic reviews. We suggest that systematic reviews should be strengthened by the consideration of mechanistic evidence. Developing methods to compile all relevant mechanistic evidence for a given causal question will undoubtedly pose a challenge, but it will be necessary to overcome this challenge if we wish to make medical decisions that are informed by the best available evidence, not simply the most easily searchable evidence. One possible strategy would involve two stages of searching. First, reviewers would identify all reasonable mechanistic pathways that have been proposed for the causal claim of interest in the literature. Some of this information could undoubtedly be gleaned from the same randomized controlled trials already being reviewed, as many authors attempt to justify an observed association by suggesting a causal mechanism. The second stage of searching would involve dissecting each proposed pathway into its separate mechanistic steps, and searching out the available association-based evidence to support the causal claim represented by each mechanistic

step. This would undoubtedly involve many of the same searching techniques already used in EBM to investigate broader causal claims.

- ***EBM should evaluate how strongly mechanistic evidence supports causal inference.***

The purpose of considering mechanistic evidence alongside exposure-outcome evidence is to distinguish those causal claims that are well supported by mechanistic evidence (and thus less likely to be spurious) from those that are poorly supported by mechanistic evidence (and thus more likely to be spurious). Therefore, in order to *use* mechanistic evidence meaningfully, we will need a reliable system of *evaluating* mechanistic evidence. We have laid out some basic ideas for evaluation in section 5. Generally speaking, we believe any evaluation of a proposed mechanistic pathway should ask at least three questions: Do the entities involved in the proposed pathway exist? Based on the evidence, is each step in the pathway best explained as causation? Are there any evidential gaps in the pathway?

- ***Once exposure-outcome and mechanistic evidence have been evaluated individually, the total strength of the causal claim should be evaluated based on both evidence types.***

Perhaps the most challenging task in the effort to integrate mechanistic evidence into EBM is the actual integration step, in which we hold exposure-outcome and mechanistic evidence together and weigh the causal claim accordingly. It is difficult to imagine a formula that encompasses the confidence we should feel in a causal claim based on the combination of these two different evidences. One solution, however, emanates from the particular weaknesses of the evidence types. Once we have assembled all of the exposure-outcome evidence and all of the mechanistic evidence, we can ask to what extent we can rule out each weakness: how confident are we that the association observed in the exposure-outcome association reflects causation, as opposed to chance, confounding, or bias? How confident are we that the causal relationship predicted by the proposed mechanism is the same as the net causal relationship between exposure and outcome? In certain cases, strong exposure-outcome evidence may narrow the non-causal possibility enough to justify causal inference in the absence of mechanistic evidence. In other cases of well-understood and reliably demonstrated mechanisms, mechanistic evidence alone may justify action. In general, however, we expect that consideration of both evidence types will improve medical decision-making and ultimately save the lives of patients.

References
Administration, U. S. F. a. D. (2014). "New Drugs at FDA: CDER's New Molecular Entities and New Therapeutic Biological Products of 2014."

Bechtel, W. and A. Abrahamsen (2005). "Explanation: a mechanist alternative." Stud Hist Philos Biol Biomed Sci **36**(2): 421-441.

Broadbent, A. (2011). Inferring causation in epidemiology: mechanisms, black boxes, and contrasts. Causality in the sciences. P. M. Illari, F. Russo and J. Williamson, Oxford University Press**:** 45-69.

Broadbent, A. (2013). Philosophy of epidemiology, palgrave macmillan.

Clarke, B. (2011). Causality in medicine with particular reference to the viral causation of cancers, UCL.

Clarke, B., D. Gillies, P. Illari, F. Russo and J. Williamson (2013). "The evidence that evidence-based medicine omits." Prev Med **57**(6): 745-747.

Claveau, F. (2012). "The Russo-Williamson Theses in the social sciences: causal inference drawing on two types of evidence." Stud Hist Philos Biol Biomed Sci **43**(4): 806-813.

Dammann, O. and A. Leviton (2007). "Perinatal brain damage causation." Dev Neurosci **29**(4-5): 280-288.

Evidence-Based Medicine Working, G. (1992). "Evidence-based medicine. A new approach to teaching the practice of medicine." JAMA **268**(17): 2420-2425.

Force, U. P. S. T. (2008). "USPSTF Procedure Manual." **AHRQ Publication No. 08-05118-EF**.

Franklin, R. E. and R. G. Gosling (1953). "Molecular configuration in sodium thymonucleate." Nature **171**(4356): 740-741.

Garson, J. (2013). "The Functional Sense of Mechanism." Philosophy of Science **80**(3): 317-333.

Glennan, S. (2002). "Rethinking mechanistic explanation." Philosophy of Science **69**(S3): S342-S353

Glennan, S. S. (1996). "Mechanisms and the nature of causation." Erkenntnis **44**: 49-71.

Group, O. L. o. E. W. (2011). "The Oxford 2011 Levels of Evidence." Oxford Centre for Evidence-Based Medicine.

Hackman, J. R. (2003). "Learning more by crossing levels: evidence from airplanes, hospitals, and orchestras." Journal of Organizational Behavior **24**(8): 905-922.

Hill, A. B. (1965). "The environment and disease: association or causation?" Proc R Soc Med **58**: 295-300.

Howick, J. (2011). The philosophy of evidence-based medicine. Chichester, West Sussex, UK, Wiley-Blackwell, BMJ Books.

Illari, P. and F. Russo (2014). Causality : philosophical theory meets scientific practice. Oxford, Oxford University Press.

Illari, P. and J. Williamson (2012). "What is a mechanism? Thinking about mechanisms *across* the sciences." Eur J Phil Sci **2**: 119-135.

Illari, P. M. (2011). "Mechanistic evidence: Disambiguating the Russo–Williamson thesis." International Studies in the Philosophy of Science **25**: 139–157.

Illari, P. M. and J. Williamson (2011). Mechanisms are real and local. Causality in the Sciences. P. M. Illari, F. Russo and J. Williamson. Oxford, OUP**:** 818-844.

Lewis, L. E., J. M. Bain, C. Lowes, C. Gillespie, F. M. Rudkin, N. A. Gow and L. P. Erwig (2012). "Stage specific assessment of Candida albicans phagocytosis by macrophages identifies cell wall composition and morphogenesis as key determinants." PLoS Pathog **8**(3): e1002578.

Lipton, P. (1991). Inference to the best explanation. London ; New York, Routledge.

Machamer, P. K., L. Darden and C. F. Craver (2000). "Thinking About Mechanisms." Philos Sci **67**: 1-25.

Meng, W., B. A. Ellsworth, A. A. Nirschl, P. J. McCann, M. Patel, R. N. Girotra, G. Wu, P. M. Sher, E. P. Morrison, S. A. Biller, R. Zahler, P. P. Deshpande, A. Pullockaran, D. L. Hagan, N. Morgan, J. R. Taylor, M. T. Obermeier, W. G. Humphreys, A. Khanna, L. Discenza, J. G. Robertson, A. Wang, S. Han, J. R. Wetterau, E. B. Janovitz, O. P. Flint, J. M. Whaley and W. N. Washburn (2008). "Discovery of dapagliflozin: a potent, selective renal sodium-dependent glucose cotransporter 2 (SGLT2) inhibitor for the treatment of type 2 diabetes." J Med Chem **51**(5): 1145-1149.

Nicholson, D. J. (2012). "The concept of mechanism in biology." Stud Hist Philos Biol Biomed Sci **43**(1): 152-163.

Rao, R. R., J. Z. Long, J. P. White, K. J. Svensson, J. Lou, I. Lokurkar, M. P. Jedrychowski, J. L. Ruas, C. D. Wrann, J. C. Lo, D. M. Camera, J. Lachey, S. Gygi, J. Seehra, J. A. Hawley and B. M. Spiegelman (2014). "Meteorin-like is a hormone that regulates immune-adipose interactions to increase beige fat thermogenesis." Cell **157**(6): 1279-1291.

Russo, F. and J. Williamson (2007). "Interpreting Causality in the Health Sciences." International Studies in the Philosophy of Science **21**(2): 157-170.

Russo, F. and J. Williamson (2010). Generic versus Single-case Causality: the Case of Autopsy.

Sackett, D. (1995). "Evidence-based medicine." Lancet **346**(8983): 1171.

Sanford, M. and J. H. Lo (2014). "Elosulfase alfa: First Global Approval." Drugs **74**(6): 713-718.

Stallones, R. A. (1980). "To advance epidemiology." Annu Rev Public Health **1**: 69-82.

Thagard, P. (1999). How scientists explain disease. Princeton, Princeton University Press.

Vachharajani, N. N., K. Yeleswaram and D. W. Boulton (2003). "Preclinical pharmacokinetics and metabolism of BMS-214778, a novel melatonin receptor agonist." J Pharm Sci **92**(4): 760-772.

Watson, J. D. and F. H. Crick (1953). "Genetical implications of the structure of deoxyribonucleic acid." Nature **171**(4361): 964-967.

Worrall, J. (2010). "Evidence: philosophy of science meets medicine." J Eval Clin Pract **16**(2): 356-362.