

Submission of comments by Carl V. Phillips PhD
on the
Health Effects of Wind Turbines on Nearby Residents
to the Australian Senate
in the matter of
The Social and Economic Impact of Rural Wind Farms

9 February 2011

My name is Carl V. Phillips. I have been asked by Senator Steve Fielding to submit a written report in this matter and make myself available for possible oral testimony (which I am happy to do, via telecom, if we can work out the logistics). I welcome any queries on this matter from those involved in the Senate inquiry.

The following is an adaptation of an evolving report I have submitted, in various versions, to a series of regulatory hearings and trials in the United States, Canada, and New Zealand. I did that work as a consulting expert witness, retained by parties who opposed the siting of wind turbines or sought stricter regulations of them. I am an independent public health scientist in this matter, not a member of any affected community nor an activist (except in the sense of advocating for more scientific research that would allow better decision making), and I received neither input from any interested party nor any compensation for preparing the present submission.

Personal background/credentials and an overview of my scientific approach

I am an expert in epidemiology and related health sciences, as well as scientific epistemology and methodology. I earned a PhD in public policy (with an emphasis on economics-based decision making) from Harvard University, completing a dissertation on environmental policy and economics. I then completed the Robert Wood Johnson Foundation Scholars in Health Policy Research postdoctoral fellowship at the University of Michigan. Later I did a second fellowship in philosophy of science at the University of Minnesota. Before I returned to school for my PhD and began my career in public health science, I worked in consulting, primarily analyzing energy and environmental policy issues. Prior to that I earned a Master's in Public Policy from the Kennedy School of Government at Harvard, and *summa cum laude* undergraduate degrees in math and history from Ohio State University.

I spent most of my career as a professor of public health. I currently direct an independent academic-style research institute (a continuation of my university research lab) and a small consultancy. During my career as a professor, I taught at the schools of public health at

University of Minnesota, University of Texas, and University of Alberta, the evidence-based medicine program at University of Texas medical school, the University of Alberta medical school, and Harvard's Kennedy School of Government. My teaching focused on two subjects: how to make optimal public policy decisions based on scientific evidence, and how to properly analyze epidemiologic data. This subject matter, as important as it is, is generally overlooked in health science and medical education, and students frequently reported that my teaching clarified their understanding of epidemiology, science more generally, and policy decision making for the first time in their educational careers.

My research during my academic career, and continuing in my private institute, has emphasized epidemiologic methods, environmental health, science- and ethics-based policy making, the nature and quality of peer review, and tobacco harm reduction. My work on epidemiologic methods focuses on recognizing and quantifying uncertainty, recognizing and correcting for biased analyses, and translating statistical results into decision-relevant information. My initial contributions in the area of quantifying uncertainty won several awards in the early 2000s and launched a new area of inquiry in the field.

Epidemiology is the study of actual health outcomes in people, and thus is the only science that can directly inform us about actual health risks from real-world exposures. Related biological and physical sciences often provide useful information about health risks, but they are ultimately trumped by epidemiology because real-world exposures and the human body and mind are so complex that we cannot effectively predict and measure health effects except by studying people and their exposures directly. My background in epidemiology methods, scientific epistemology, and optimal policy decision-making is the background that is needed for being able to evaluate bodies of health science literature and assess their worldly implications. Most people who work in or around epidemiology learn only how to conduct particular types of studies or how to technically interpret individual study conclusions in the simplest possible way, which does not provide the tools to sort out complicated controversies. My study and research have focused on the epistemology of epidemiology: how to understand what the available evidence tells us beyond what the authors of individual studies assert.

My background in environmental economics and environmental health, with an emphasis on energy policy, provides important subject-matter background related to wind turbines, land use, and electricity generation. I have spent substantial time in the last year reviewing what has been written about the health effects of wind turbines, including research and related background science, health reports, and testimony/analysis by industry, government, and independent scientists. I have contributed my own analysis summarizing what we know and what it means (forthcoming), and initiating new field research on the topic (in progress).

Summary of my scientific opinions

It is my learned opinion that there is ample scientific evidence that wind turbines sited near residences cause serious health problems for some people living in those residences. However, we do not have enough of the systematic population-based research that is needed to estimate the portion of the population that is vulnerable to the effects, the effects of exposure variables (e.g., how risk varies depending on how far away from a residence the turbines are), and the effects of

other variables. In addition, we do not yet understand the pathways that lead from cause to effect.

It is possible to make rough estimates of the portion of the exposed population that will suffer serious health problems from turbines within one or two thousand meters of their residences, based on adverse event reporting and our limited systematic data: It is more than trivial (that is, it is at least a few percent of the exposed population) but fairly clearly a minority. This is obviously a wide range, but it does exclude the possibility that the effects apply to only a few rare individuals, let alone that they do not exist.

We would like to be able to better estimate this figure, especially as a function of intensity of exposure (distance from the turbines and other characteristics of the facility) and personal characteristics. We would especially like to know if particular types of people are at much higher risk. Moreover, we would like to have a better list of exactly which health outcomes are caused by the exposure; we have strong evidence about some, but others are more speculative. Additionally, it would be extremely useful to know which of the several candidate causal pathways (noise, shadow flicker, etc.) leads from the existence of the wind turbine to the health outcomes; without that knowledge, it is not possible to assess the options to mitigate the effects. However, the fact that we do not have such knowledge does not diminish the actual epidemiologic knowledge we have. That is, we know there is a serious problem, we just do not know exactly why it happens or how prevalent it is, let alone what steps (other than completely eliminating the exposure) could mitigate it.

If this were a pharmaceutical exposure, the manufacturer would have been required to do research to try to document the “side effects” and clear up these unknowns. Unfortunately, no entity with sufficient resources has been required to or has chosen to do that research for wind turbines, so the information can only trickle in as fast as various under-funded – often self-funded or community-funded – researchers can produce it. Still, even without sufficiently financed research and with the unknowns I have noted, one thing is quite clear from the science that has been done: There is no way that someone can claim that there is not evidence to support the claim that there are serious health effects.

To support these conclusions, my analysis follows the following logic, which is expanded upon below:

1. There is substantial epidemiologic evidence of health effects on residents from nearby wind turbines, including numerous adverse event reports (case studies), many of which are case-crossover studies, and a few population-based systematic studies.
2. When the nature of the situation and evidence is carefully considered, it becomes particularly convincing. It stands up well to criticisms and concerns that have been leveled at its validity.
3. The health effects that have been reported are serious and have important implications for people’s well-being. Attempts to dismiss these as “not real diseases” are inappropriate and contrary to widely accepted definitions of disease and health.

4. The causal relationship between exposure and disease is quite plausible and requires no great imagination. Wind turbines produce audible noise, noise at sub-audible frequencies, and shadow and light flicker, and all of these affect people's minds and bodies. The noise and light are cyclic rather than constant, which can be particularly bothersome. We are not sure how much of the outcome results from particular causal pathways, but there is no reason to doubt the accuracy of the epidemiologic evidence based on a lack of plausible pathways.

Beyond the core analysis, I include two appendices, reprinted from previous reports I have filed:

5. In the course of making the above points, I offer a general rebuttal of some of the points that are most often made in pro-industry reports (ones that are not already inherently part of the previous points, and thus already covered). Since all the pro-industry reports contain very similar claims, this approach covers most of them. I have been asked to write point-by-point responses to various reports written by industry supporters that were considered of particular interest in specific cases. The one such report that seems to be most frequently cited by industry supporters is the North American industry's main report (Colby 2009). I was also once asked to comment on two pro-industry documents published by the government of Australia, which may be of particular interest in the present forum. Thus I include what I previously wrote, as rebuttals of the claims in those documents. I will note that these were written in a scientific, not a political style, which unfamiliar readers might find a bit unusual; while political discussions call for a certain gentility and diplomacy, serious scientific argument requires calling something naïve, misguided, or wrong, if it deserves such judgment.

Affirmative evidence that individuals are suffering disease from wind turbines

The greater part of the evidence about individuals living near wind turbines takes the form of adverse event reports, also known as case reports. Many of these are individual stories that are self-published. Others have been collected more systematically, such as the WindVOiCe collection from Ontario (Krogh et al. 2011), the scholarly book by Pierpont (2009), and in Harry (2007). Adverse event reporting is a cornerstone of identifying emerging health risks, since it is obviously impossible to constantly study every possible health risk in a more systematic way, waiting for health outcomes to appear. Pharmaceutical regulators rely heavily on clearinghouses they create for adverse event reports from drugs. The WindVOiCe report collection is an example of this same well-accepted kind of active-recruiting data collection system.

These adverse event reports provide useful information in several ways:

First, the reported health problems are similar across reports and are plausibly related to each other and the exposure. If people were complaining about a collection of seemingly unrelated ailments, it might suggest that they were simply blaming the turbines for coincidental problems. However, this is not the case. Instead, these outcomes consistently include a combination of the same list of related problems, including sleep disorders, general distress, and mood disorders, headache, fatigue, vestibular (balance) problems, and tinnitus. These problems exist at the border of the psychological and physical, and can all be caused by plausible effects of wind turbine exposure: stress reactions and/or vestibular disturbance. There are also a few reports of

hypertension, though since this is difficult for individuals to monitor themselves it would be unlikely to appear in most adverse event reports.

Second, the sheer volume of adverse event reports suggests that the problems are not restricted to a few rare highly-susceptible individuals. It is impossible to make a very confident estimate based only on adverse event report collections, though pharmaceutical regulators attempt it, as best they can, as a matter of standard practice. We do not know what portion of the people experiencing the adverse events choose to volunteer the information in a form that is accessible to researchers and regulators (although this is typically very low); moreover, we do not even know how many people are exposed. That said, the rate of volunteered complaints suggests that the prevalence of serious health problems is well above 1% of the exposed population. When communities or turbine operators actively solicit reports from a population of known size, several percent of the population in the monitored area (typically a fairly wide radius from the nearest turbine) make complaints.

Third, several of the case studies provide *case-crossover study* data, albeit without exactly the optimal data collection. The case-crossover study (Maclure 1991) is one of the most effective methods for assessing the transitory effect of a transitory exposure (which means in this context: occurring fairly soon after being exposed to turbine noise and disappearing fairly soon once the exposure is removed). It is a method for transforming the limited information that may come from a single case study to a controlled study with many observations. This study design is one familiar to all of us in our everyday lives: Impose the exposure at a time when the outcome of interest is absent and see if the outcome occurs; withdraw the exposure and see if the outcome disappears.

For example, if you think that a particular food gives you stomach pain, note whether that pain occurs after eating the food but not at other times; avoid eating the food for a while and see if the pains are absent; and if you really want to pay the price of finding out, try the food again and see if the pain returns. It is obvious how to translate the resulting observations into a causal conclusion about the food and the pain. For an even simpler illustration, our usual study design to figure out if a particular switch turns on a particular light is to flip the switch (often three times for some reason) – nothing more complicated is needed. When it is possible to collect case-crossover data, it usually provides among the best possible epidemiologic information. Its advantages include more data (each exposure change serves as an observation, whereas most other study designs produce only one observation per person) and individuals serving as their own comparison population.

For the case of turbines, many of the reports (personal testimonials and collected case series) recount the onset of distress beginning shortly after the activation of the turbines – that is, when the person crossed over from being unexposed to being exposed. Moreover, some residents have reported intermittent reductions in their health problems under certain conditions (when the turbines stop turning, or other effects of wind direction or speed), further supporting the conclusion. Finally, some of those who sought relief through physical (soundproofing) or geographic (moving their home) methods crossed back to unexposed and reported the results. While the cause-and-effect pattern might not be so obvious as the light switch experiment

(because the appearance and disappearance of many of the effects would not be immediate), this is compelling evidence.

Many of those collecting data and individuals reporting on their experiences intuitively understand the value of crossover data. It does not appear that anyone has instituted a formal method for such data collection (the term “crossover” does not appear to have been used in discussions of wind turbines and health before I first addressed the point), and so did not collect all the possible data about timing and crossing back in the optimal form. But there is still convincing data and I have encouraged the collecting of this information more formally in the future, and I have a plan to systematize the data that exists, though it is difficult to get funding to support new research.

In addition to individual case-based data, there is a small collection of studies that use data gathered in a systematic way. There is a small collection of epidemiologic studies of people exposed to wind turbine noise in Europe by Pedersen and colleagues (2004, 2007, 2009, 2010). These studies suggest that some substantial portion of exposed individuals experience harms, some of which constitute health problems by any modern definitions of health. The studies have various limitations, but they provide a quantification of a nontrivial number of cases. Among the collections of cases that have been reported by advocacy groups or consultants working on previous regulatory cases, many are systematic studies though they may have been mischaracterized as non-systematic case series. It appears that no one has done a systematic review of such studies, so it is not possible to generalize about them. But as one example, Phipps (2007) presented a systematic study that provides further evidence of health effects, and I am aware of no examples that provide contrary evidence.

Most recently, Nissenbaum et al. (2011) surveyed residents living near turbines about most of the above-mentioned health conditions and compared them to similar people living further away. At the time of this writing the details of their research – which I have seen and commented on – are still confidential, but some of it has been reported in public forums so I am comfortable making some broad statements about it. The data appear to support most of the widely-stated hypotheses about the health effects of nearby turbines, and provided no contrary evidence. This study added a systematically-collected, population-based study with formally measured health effects to the types of evidence that already existed, and suggests that the portion of the population suffering the effects is much higher than the conservative estimates I present above. This is not to say that this one formal study provides definitive evidence, nor is it better evidence than we already had just because it was more systematic. But it does suggest that the estimates based on adverse event reporting are extremely conservative (as they often are, since most adverse events are never reported). Moreover, this study shows that different types of evidence, gathered in different ways, tend to further support what the adverse event reports show.

Finally, observations about behavior and expenses endured provide further evidence of a causal relationship between turbines and health problems, as well as offering a measure of the magnitude of some of the problems. The reported adverse event reports suggest that a nontrivial number of residents who experienced severe problems concluded that the turbines were damaging their health with sufficient confidence to move their residence or retrofit the structure to try to block the noise. These are expensive actions that would not be taken by people who

were suffering only minor problems or who had not made every effort to make sure the cause of their disease was indeed the turbines. Some case studies recount residents attempting to sell their properties but not finding buyers at a price they would accept, suggesting that potential buyers anticipated suffering health problems if they moved near the facility.

It might be surprising to see such observations being used as epidemiologic data, but scientific inference is not a matter of following a recipe for conducting the same type of study for any question. In this case, there is the opportunity to infer health information from data about real estate sales.

It is theoretically possible that everyone involved (residents, all potential buyers) is so misled about the causes of their health problems that they would spend their lives or waste thousands of dollars in error, and others would fail to take advantage of their error and not buy the heavily discounted houses. However, economists recognize that when there is data like this (called “revealed preference”), it is usually the most compelling evidence available.

In summary, though no one has been willing to pay for some of the studies that many of us would like to see done, we have evidence from population-based surveys, adverse event reports, individual experiments (case-crossover studies), and economic revealed preference that all points toward the same conclusion: Wind turbine facilities are causing serious health effects at a nontrivial rate.

Epistemic observations about the affirmative evidence

There are several epistemic points, points about the nature and meaning of the affirmative evidence, that should be addressed. Some of these are responses that are commonly made by the industry and its consultants to try to dismiss the evidence. Others are reasons why this body of evidence is more compelling than some observers might realize.

The breadth and nature of the existing data make it more compelling than the sum of its parts. As with most complicated science, no individual piece of evidence is compelling. Similarly, gathering more similar information in the same way is often not compelling. In this case, we have different types of data, collected in different ways. Some of it is experimental (a case-crossover study is an experiment on oneself), some of it is systematic, and most of it is based on pre- and post-exposure experiences of particular individuals.

Some commentators have attempted to dismiss the entire body of evidence because most of it is not based on one of the two or three epidemiologic study types that they claim to understand. With the Nissenbaum study and other ongoing work, this rhetoric will become difficult to sustain. But it still seems likely that dismissing all of the evidence will remain a preferred tactic for some time. The present analysis should help show why these nihilistic claims are not defensible and further demonstrate the fallacy of the “no evidence” claim. A huge portion of all knowledge, including formal scientific inference, is based on data that is not from studies designed according to certain standard approaches. Collecting more information never hurts, but we have to reach the best conclusion we have based on what we know. It should be obvious that the statement “does not tell us everything we want to know” does not mean “has no information

content”. Claiming that there is no evidence even though there are reports of individuals suffering is akin to claiming that there is no evidence that people get injured as a result of text-messaging while engaged in other activities because, even though the pathway is obvious and there are numerous accidents occurring from some activities, there is often not a “real study” that allows us to make various quantitative estimates. It is perfectly reasonable to try to make a case that our existing knowledge does not provide convincing evidence of a claim, but when someone simply tries to convince the reader that most of that evidence does not even exist or has no content, it suggests that they do not feel like they can make that case.

The affirmative evidence, particularly in the context of how to best make sense of this particular exposure-disease combination, presents a compelling case that there are substantial health effects. It is difficult to see how anyone could take seriously the assertion that there is no evidence to support the case. Even the most dubious observer would have to conclude that negating the existing evidence requires either further study (assuming it produces results that contradict the hypothesis) or compelling arguments about why each of the many sources of evidence is so flawed as to be completely uninformative. Denying the existence of evidence or casually stating, without support, reasons why some of the data might be misleading are simply not adequate.

It is helpful that that this exposure-disease combination is relatively easy to study, including by individuals recounting their own experiences, in contrast with more typical claims about health effects from noxious facilities. In cases of environmental pollution there is often a fear of slow-developing diseases (especially cancer, for which we cannot even define a time of incidence – i.e., when the disease actually started – only of diagnosis) that occur seemingly at random because they have many causes, usually far in the past, that it is impossible to sort out a specific cause for a particular case. In such cases, when local residents claim “I got cancer because of the effluent from this factory” the standard response is that it was inevitable that some people near the factory would get some cancer someday, and so it is impossible to make that causal conclusion. Indeed, to make any such conclusion it is generally necessary to systematically collect enough data on enough exposed cases, as well as on non-cases and an unexposed comparison group, so that statistical comparisons can be made. (The caveat “generally” is meant to recognize the fact that if twenty cases of the same rare cancer were reported among a few hundred exposed individuals we would not actually need to know more than this and common knowledge to draw the causal conclusion.) Contrast this epistemic situation with the case of a traumatic injury from a car crash: If following a crash a passenger in the car has a laceration on his head that he did not have a few minutes earlier, we would not hesitate to say, based on that information alone, that the crash caused the injury. Why? Because head lacerations do not slowly develop from unknown causes, appearing years later (like cancer); instead they are almost always diagnosed within seconds after a causal event occurs. Moreover, we can almost see the causal pathway in the form of the crash causing rapid deceleration which caused an impact between head and something in the car, and it is such impacts that cause trauma.

The case of the health outcomes that appear to be caused by wind turbines lies somewhere between the cancer and crash examples, but is rather closer to the latter. Unlike for the trauma case, we cannot fully envision the causal pathway. But the particular health problems and general distress that has been observed are not phenomena that, like cancer, often suddenly occur

without any observable proximate cause; if the problem is new, the cause is almost certainly new. In addition, the sensory impact of nearby wind turbines is readily observable, like the car crash, and unlike chemical exposure. The reasonable expectation of a proximate explanation and ability to observe the turbine noise as the apparent cause make this case more like the car crash than the cancer. Some of the authors writing reports for the wind energy industry come from consulting groups that write a lot of reports for industry that argue, often correctly, “you can never know what caused those disease cases because you do not know exactly when it was caused and there are many plausible causes other than our client’s product/facility.” But they have tried to import that exact claim to the present case where the observation is simply incorrect.

Interestingly, some authors who dispute the claim that the operation of nearby turbines caused observed health problems have proposed other causal pathways from the turbine installation to the problems. They have claimed that those with health complaints just do not like the turbines for other reasons, or have been talked into believing that there is a health risk, and so are distressed because of that, or perhaps they even concoct health claims because they have other complaints about the land use. A more elaborate version that appeared in at least one report is that those suffering health problems after local turbines began operation would have just lost a battle against the siting of the turbines and would be suffering from exhaustion or a sense of defeat that would cause onset of the disease. Whatever the merits of these suggestions (and it should be emphasized that they are just hypotheses, without any supporting evidence), those making the suggestions are implicitly acknowledging that the nature of these health problems means it is likely that onset was caused by a recent event. Moreover, their alternative explanations for the health effects tend to concede that the installation of the turbines is indeed the event that is causing the problems, and they are merely trying to insinuate that the specific pathway somehow means that the health problems do not “count”.

Though it is not quite so easy to observe the proximate cause of distress and the other psychological and physiological manifestations associated with wind turbines as it is to observe the crash as the cause of head trauma, a subject’s own observations about his own case are still scientifically informative. This contrasts with most types of cancer, wherein neither the victim nor any clinician or scientist can offer a legitimate conclusion about causation, other than in the form of far-from-certain probabilities derived from statistical comparisons. Except for the very few cancers where we know the causes of almost every case (i.e., cervical cancer), a claim that “this exposure caused my (or this patient’s) cancer” is never justified. But if someone claims “this noise is driving me crazy and keeps me from sleeping” we have good reason to believe him. For a more subtle exposure, like a relatively low decibel periodic noise, the conclusion is less certain than it would be for a loud party next door, but the individual’s assessment still has substantial value. This is true even apart from the crossover data that an individual will naturally accumulate, so when combined with crossover data (either from actively moving away from the area or just the inevitable periods of low wind activity) and common intuition about how to reason based on crossover evidence, the individuals’ assessments are even more compelling.

In understanding these points, it is important to not confuse incidence (the onset of a particular condition) with prevalence (having the condition). A scientifically incorrect response to some of the above points is “but at any given time, lots of people have these health problems, so why

should we conclude they were caused by the turbines?” But this describes prevalent cases, while adverse event reports and the systematic data gathering focus on problems with incidence shortly after the start of operation of the turbines. Going back to the car crash comparison, lots of people have headaches at any given time, but someone who acquired a headache at the time of a crash probably suffered an injury from the crash. When coincidence is common (e.g., someone has cancer and once lived near an industrial facility), assessing whether there are extra cases that are apparently caused by the exposure requires the complicated statistics used in some epidemiology studies. But the number of coincidental incident cases is very low for the car crash or turbine examples, making population-based statistics less necessary to merely establish that something occurs (though they are still needed to quantify the risk).

Thus, unlike the case of trying to detect an elevated incidence rate above some baseline level of a disease that has distant and uncertain causes – which is generally impossible absent formal studies that are specifically designed to do just that – the natural observations in this case are quite compelling.

A few specific points made by industry proponents about the nature of the evidence are worth addressing:

Subjectivity of the measures

Industry consultants have correctly observed that most of the observed health effects are subjective (meaning that they cannot be measured except by the person that is experiencing them, as is the case for mood or headache). But this does not make the results less real. Many important diseases are diagnosed and, indeed, defined based entirely on subjective experience, everything from suicidal depression to a minor headache. To dismiss subjective experiences would be to dismiss the vast majority of what people genuinely care about in the world, as well as many fields of science and medicine. It is certainly true that the reported outcomes are much more difficult to measure than many other health outcomes. That, of course, is not a reason to ignore what we do know. We have to make the best of whatever form our data takes rather than declaring the data to be less than one might want and misconstruing that disappointment as complete ignorance.

Often when commenting on the issue of subjectivity, industry proponents claim that *all* of the effects are subjective, but this is not true: For example, inability to sleep can be measured objectively, and loss of concentration can manifest in reduced productivity, which is an objective measure. Many commentators may be confused, thinking “subjective” means “psychologically mediated”.

The observations about subjectivity seem to be intended to downplay the importance of the experience. But the real challenge of many subjective effects is that context affects how much suffering they cause, and so only evidence about this specific experience is of much value: A minor trauma might cause a pain we just shake off while playing sports but the pain from such a trauma might be completely unpalatable if it occurs while sitting quietly at dinner, and if that pain were inflicted by someone walking by and hitting us on the head we might consider it even more painful and a criminal act. Many studies of the effects of noise on people take place in the workplace or short-term artificial settings where someone might find the stimulus tolerable and

typical, whereas they would find it distressing if experienced constantly and at home. The challenge is not just that the experiences are different based on circumstance, but that none of these experiences is more legitimate than another. If the subjective experience of the resident at home is different from the worker on the job or someone visiting the home to assess the effects, it needs to be dealt with as it is. Observing that sometimes people are able to ignore particular impacts without health effects might suggest intervention methods for reducing the impact, but the observation does not in itself reduce the effects.

Heterogeneous effects

Some observers appear to be confused by the heterogeneous effects – e.g., some people report experiencing debilitating symptoms from their exposure to turbines, while others with greater exposure (as measured physically) have no significant problems. But this is not at all unusual, and similar patterns can be observed for most any exposure-disease combination. For example, many heavy smokers never get cancer or suffer any other major disease that is often caused by smoking, but that does not lead us to doubt that smoking, even light smoking, gives some people lung cancer.

This relates to the above points about contextual effects. Some observers have had brief exposures to the noise and experienced no adverse effects, and perhaps concluded that the exposure would never bother them. This obviously does not constitute evidence that no one ever suffers from the effects; the individual in question might be immune while others are not. Moreover, he might be wrong about what would happen if he were exposed longer, since health problems caused by noise exposure tend to be cumulative, as is typical for other exposures that produce stress reactions. Stress-causing exposures (e.g., social harassment, pain, sleep deprivation, physical restraint) which may seem trivial for an hour can become torture after a week. Even exposures for eight hours per day (like workplace exposures) may have quite different effects than exposures that last all day and overnight. Some exposures that people intentionally seek controlled versions of for an hour or a workday-length period (hot weather, loud music, exhausting exercise) cause stress reactions and health problems with unremitting, uncontrolled long-term exposure. Something that is beneficial in one controlled form is harmful in other forms. (It might seem like a strawman argument to make that last observation, but I mention it because I have repeatedly observed industry proponents argue that low-frequency sound is used as a therapeutic tool under some circumstance and therefore there is no way it can be harmful. Apparently they are unfamiliar with the various uses of knives.)

When effects are heterogeneous across people and across intensities of exposure, we can often find measurable characteristics that are associated with whether someone experiences the effect. A great deal of epidemiology is devoted to measuring these associations. Yet some of those who deny that wind turbines affect health seem to find such associations surprising or want to imply that the heterogeneity of effects suggests that the effects are less “real”. Those authors may be making the mistake – common among people who do not understand complicated sciences like epidemiology – of thinking that if an outcome has one cause (personal characteristic) then another factor (noise) is not really the cause. In fact, the proper way to think of it (though it takes some getting used to) is that *both* the noise and the personal characteristics caused the disease (as did a multitude of other factors); if either one of them was absent then the disease would not have occurred. In ethical or policy discourse (as opposed to scientific analysis) we

often reduce our list to causes that someone actively brought about (i.e., the causes someone is culpable for, not the ones that simply are). So, for example, a murder is not excused, and is not considered to not be the cause of the death, if his victim would have survived had he been stronger or closer to a hospital – those non-act-based personal and geographic characteristics *also* caused the death, but the murderer is still a cause and thus is guilty.

What constitutes evidence?

Some commentators take what might be called a legalistic approach, which I address here to point out why it is not science. In court cases there are rules about what cannot be used as evidence and what categories of evidence are considered to dominate others, and such rules might be seen as necessary to bring order to an adversarial system. Truth seeking via science, however, is messy and does not function with such rules.

On one side, no appeal to authority can outweigh actual analysis of the evidence or, indeed, in itself constitute evidence. There are several government and NGO reports and many regulatory decisions that have been made that favor the installation of wind turbines near residences. But these are only as good as the evidence they are based on and the quality of the analysis. Simply noting that someone already reached a particular conclusion matters for something like legal precedent, where the previous conclusion has weight in itself, but is not a basis for scientific analysis. Indeed, when it comes to the science this difference seems to be recognized by courts – e.g., at least in some jurisdictions a scientific expert witness cannot simply base his opinion on what some authority has declared to be true. Scientific analysts should, of course, draw upon the evidence and arguments used by previous authors. But it appears that most of those writing reports supporting further building of wind turbine facilities have put inappropriate faith in the conclusions of previous authors without actually looking at their evidence or argument. In short, a statement like “this siting complies with WHO recommendations” or “with the local law” tells us nothing about what the health outcomes might be unless it is also shown that that WHO or the local regulators knew what the health outcomes would be (which they did not).

In addition, science evidence is not excluded because it does not pass some test. Some authors have claimed that there is “no evidence” because there is nothing that looks at exactly what we want to know, with the study carried out in the best possible way, published in a particular way. But this argument, much liked by liability defendants, can be used to deny all of our knowledge about *anything*. We have the evidence that we have and need to use it.

Perhaps some rule about “only consider evidence published in certain journals” is reasonable for a debating league, since games played for fun always have arbitrary rules, and may be defensible in a stylized battle like a liability trial (though even then it is a formulaic substitute for really doing the science), but it is clearly not appropriate for policy decisions which are supposed to be made in the best interest of the public. Scientific experts should be aiding the process of making the best sense we can of the evidence we have, not trying to pretend we have no evidence. Of course, non-experts often have no choice but to rely on rules of thumb (e.g., believe the WHO, focus on journal articles, ignore anything else due to inability to judge its worth). So if policy makers without relevant scientific background have to make a decision on their own, rules of thumb would have to do. But for the case of wind turbines, experts – who by definition are those who are capable of independently assessing whether evidence has value – are available.

Some of those who seek to deny the evidence tried to argue that the many adverse event reports should be discounted because they are not in a peer reviewed publication (though some of them actually are; Pierpont 2009 was peer reviewed). Apparently they somehow believe that peer review can improve the accuracy of someone's report about their experience. Non-scientists often think that the peer review process offers some magical verification of analyses and data. But in health science, reviewers almost never assess these (they do not get access to the data, let alone to how it was collected; they see nothing more about the calculation methods than the reader sees, and thus cannot even check the work; see Heavner et al., 2009). Thus it should be obvious that even when adverse event reports are collected into journal articles (as they inevitably will be), the peer review process will do nothing to change the accuracy of the reports.

Similarly, some commentators seem to not understand (or pretend to not understand) that epidemiologic knowledge can be gained from many sources and in many forms, denying the bulk of the evidence not because it is not peer reviewed, but because it does not come from a particular type of study. First-semester epidemiology students may only learn how to make sense of two types of studies, but experts in the field can do much better. It is fair to say that we wish we had particular forms of data, since some studies could tell us more than others. But failure to have the perfect data obviously does not mean we have no data. We simply need to be careful about only drawing the conclusions we can. This means that we can currently be confident that a nontrivial number of exposed people suffer serious health problems, but we cannot be confident that any particular mitigation measure or offset distance rule is sufficient to protect them.

The next two sections present further reasons why we should take this epidemiologic evidence seriously.

The observed outcomes are real and very important for public health

Some commentators have tried to dismiss the reported health outcome because they are primarily psychological or because they do not officially constitute a real disease. Both of these claims are groundless and, indeed, insulting to millions of people who suffer important health problems that are similar to those being discussed here. It seems reasonable to suspect that most of the apparent health effects of wind turbines are psychologically mediated, though they may have little to do with thought processes, bypassing the conscious mind and affecting mood, behavior, or even the body. Psychological conditions and those with manifestations on both sides of the psych-physical border – a category that includes stress, depression, and many other ailments – probably account for the loss of more quality-adjusted life years than purely physical diseases, at least in the West and possibly even worldwide. Most accepted definitions of individual or public health include psychological health as part of the consideration, and usually refer to an overall state of well being rather than just an absence of a particular diagnosed physical ailment.

Legal and regulatory arenas often put greater emphasis on damage to someone's physical health, probably because it is easiest to measure. But unless physical damage causes unrelenting pain or the loss of the ability to communicate, it generally rates as much less important in terms of quality of life than psychological problems that affect mood or functionality. Indeed, research

into quality of life suggests if people were allowed to experience both and choose between them, most would rather lose the use of an arm than to suffer constant insomnia or anger, or the inability to concentrate.

It is sometimes tempting for people who do not suffer from addictions, depression, or other psychological or physical-and-psychological diseases to have the attitude that sufferers do not have a real disease and should just snap out of it. Versions of this in the case of wind turbines include such rhetoric as “it does not bother me” or “the decibel level is no greater than that of the wind blowing through leaves” so you are just imagining it is a problem. There is substantial literature that documents this dismissive attitude toward psychological and sensory stressors, and argues that it makes no more sense than telling someone with appendicitis that he is just being negative and can just “get over it”. For example, it is difficult for most of us to imagine why someone would have the urge to drink himself into dysfunction every day and not even apparently enjoy it, and so it is tempting to think, “all he has to do is realize that this is a bad thing to do and he will quit”. But it should be obvious that this is not sufficient and his affliction, however difficult it might be for most of us to empathize with, is quite real. Those who are assessing the health effects of wind turbines should consider this and realize that dismissing the problems is similar to dismissing any other problem that is “in someone’s head”.

Some of the attempts to dismiss the importance of the observed health problems are mere semantics that are intended to minimize people’s experiences via labeling. The term “annoyance” has been adopted as jargon by researchers to refer to certain psychological problems resulting from noise, and has the implication (presumably intended in some cases and not in others) that the causal pathway from noise to disease involves the type of psychological experience that is typically referred to as “annoyance” in natural language. However, it is clear that there is nothing in the use of the jargon that implies that the harm is minor, and indeed the term is used in contexts in which it includes life-ruining distress. But some commentators have tried to equate this annoyance with “mere annoyance”, a common language term used to denote a very minor problem. But of course adopted jargon does not generally have full natural language meaning. Just as “insult” can mean cause of cancer, “annoyance” in this case includes serious physical and psychological health problems.

Some commentators have tried to dismiss the observed health effects by declaring them to not really be a disease. Since the term “wind turbine syndrome” was coined by one researcher, other commentators have tried to imply that the absence of this term from official catalogs of diseases must mean that there is no disease. Obviously neither the presence nor absence of an accepted title, nor even disagreement about whether what we are seeing is a sufficiently unitary phenomenon that should be titled, has a bearing on whether people are suffering health effects. Disease definitions are created, altered, and dropped over time based on both science and politics, without changing whether and what people are suffering. Consider that there is no official disease designation for “injuries resulting from a driver of a subcompact car being hit from the left side by an SUV”, but there is certainly a collection of injuries that typically occur from that, and the absence of a label or recognition as an official syndrome will not improve that driver’s condition.

By the same token, some commentators have tried to claim that the reported effects are “just a list of symptoms” rather than a disease. This overlooks the fact that the use of the word “symptom” in this context is short for “symptom of a disease”, and the rather more practical observation that “symptoms” can be roughly defined as “the manifestations of a disease that sufferers experience, and thus care about”. Obviously sufferers’ experience is not affected by whether someone arbitrarily designates their condition to be “a disease”, let alone whether someone else denies the existence of that particular disease or wants to call it instead “a collection of symptoms”. Since those suffering the health problems are unaffected by the labels, policy makers should be equally unaffected.

Another semantic claim is that the observed suffering is not a *health* problem, but is some other form of unpleasantness. There is no accepted definition of which problems are indeed health, and certainly no bright line that demarcates the boundary health problems, and so there is genuine ambiguity. Most everyone would agree that a momentary noise that awakens someone for a few minutes imposes a real cost, but the experience does not constitute a health problem, while a constant noise that consistently prevents sleep for weeks or months, leading to depression, headaches, and an inability to function, has created serious health problem. The point in between that marks the border between the experience being “disease” or “health problem” and being some other type of cost is arbitrary and not defined. Of course, many of the problems that have been recorded – headaches, sleep disorders, etc. – are unambiguously considered to be health problems. Moreover, for some people the recorded suffering is so severe that even if someone were to argue that it is not officially a health problem, they still could not possible argue that it does not matter.

Politically, it is sometimes the case that harms are considered more important if they can be designated as diseases. For example, instead of recognizing that the aesthetic and minor immediate health effects of involuntary exposure to second-hand smoke can justify bans in indoor public spaces, bans have only been justified based on claims about life-threatening disease risk. Not only is there no bright line between suffering certain types of diseases and otherwise being forced to suffer a less pleasant life due to ailments, but the same sciences (epidemiology, econometrics) that measure the former inevitably also measure the latter. From an economic or policy ethics point of view, there is no meaningful difference: A major cost inflicted upon someone’s psychological well-being matters, whatever it is called. Anyone who is attempting to argue that the harms do not represent an official disease and that this should affect how we should treat them in decision making should be asked to declare explicitly (as they are admitting implicitly) that there really are serious effects that are important to those suffering them, but they simply should not be called diseases. Only having admitted this can they then argue that these effects should therefore be ignored in policy decisions because they are not a disease, which is different from arguing that they do not exist or are small.

The plausibility of the causal pathways

When conducting epidemiologic assessments, it is useful to establish a plausible causal pathway (i.e., that there is an apparent mechanism via which an exposure could cause a impact on a person which *could* cause the outcome in question). Such plausibility is not always necessary, and in many cases major epidemiologic discoveries were made before any mechanism was

recognized. But often when an association is observed without identifying any plausible mechanism before looking at the data, it turns out to be a statistical accident or a bias in the analysis. Of course, recognizing a plausible pathway does not prove that there will be a health outcome; assessing that depends on the epidemiology. By the same token, it is important to realize that just because we cannot figure out exactly how an exposure is causing a health effect obviously does not mean there is no health effect.

A huge portion of what turbine proponents have written about the health effects – most of the studies of turbine acoustics, reviews of research about the effects of noise and flickering, and such – can be summarized by saying “hmm, we just cannot figure out why this exposure causes serious harm to people’s health”. Some researchers have directly countered this, providing evidence-based hypotheses about why. But the key response is simply that someone’s inability to figure out why a well-documented phenomenon is happening hardly constitutes evidence that it is not happening. This would be true even if no scientist had any idea why the phenomenon was occurring, and it is certainly the case when the inability to understand any reason why the phenomena might occur appears confined to those who would prefer that the phenomena did not exist.

Many plausible candidate pathways

It is clear that the physical effects – noise traveling through the air, noise/vibrations traveling through the ground, and flickering shadows and light – do reach the bodies of local residents, and noise and flickering do cause health problems under some circumstances. This observation alone, given that effects of noise and cyclical stimuli are so varied, unpredictable, and often downright strange, is sufficient to make health effects plausible. Additionally, various studies show that lower level sound and vibrations affect the body via the ear (hearing and vestibular systems), skull, skin, viscera, and other body parts. Some authors theorize that the health effects from wind turbines seem to result from impact on the vestibular system, while others are more inclined to suspect other sensory effects. Either of these could have secondary effects due to stress reactions that could cause many other harms.

Wind turbine proponents sometimes claim that particular theories of how noise and light affect people say that there is not a health problem. The proper interpretation of such theories is that the models would not have *predicted* there would be a health problem before it was observed. Prediction is useful, but obviously its value ends once we have an actual observation. The fact that last week’s weather forecast predicted rain yesterday is not a helpful observation about what the weather was yesterday. It might be a useful observation about weather forecasting, though – i.e., the inability of the models to explain why there are health effects means that we do not fully understand the pathways. But of course, ignorance (about exactly what is happening) cannot constitute evidence (that nothing is happening).

Moreover, despite the claims that the models would not have predicted the problem, there are many studies that might lead someone to expect that turbines would cause health problems. These include evidence of impacts of low-frequency noise (e.g., Qibai 2005) and the fact that measures of acceptable noise level are typically “penalized” (i.e., are treated as if they are higher energy (decibels) than they actually are) if they are cyclic or seem out of place for the environment (as turbine noises are), as well as the long list of problems people sometimes

experience from noise, and the very strange collection of effects that noise has been observed to have on people (e.g., vibroacoustic disease as defined by Alves-Pereira and Castelo Branco, 2004). A common conclusion of many studies of noise is that exposure to noise that is not immediately harmful can, over a period of days or months, affect the body in many ways.

As noted above, the common effects of noise are the kind that people can analyze fairly effectively themselves. They can observe that the noise from the turbines seems to be bothering them, and can surmise that what they are noticing may be causing their diseases. It is well established that local residents can hear sound from turning turbines and see flickering when the direction of the sun and blades are at particular angles. Apparently without exception, the acoustic predictions and empirical observations show that the physical impacts reach nearby houses, including those that are several kilometers away from any turbine (a fact that is not changed when some of these physical researchers assert – without basis – that it does so only at magnitudes that cannot cause health effects).

Proposed alternative pathways

Those who seek to argue that there are no health effects from wind turbines sometimes claim that people are suffering health problems simply because they dislike the turbines. They do not seem to realize that what they are doing is proposing alternative causal pathways, not denying that turbines are causing health problems.

In at least one study (Pedersen and Waye 2004) a correlation was found between health problems and a negative opinion about the facilities. Some commentators have suggested that this means that the health problems are therefore less real, or perhaps even concocted due to other motives for disliking the facilities. It is theoretically possible that people who dislike the turbines for reasons that have nothing to do with their noise or light impacts could be lying about health problems, though it is difficult to understand what they do not like, if not the noise and light. Perhaps, the claim might go, they do not like the noise and light so much that they claim they are suffering health effects. But, again, why? At best they are wasting time, since experience shows there is little chance that complaints will cause an installed facility to be shut down, and at worst they are convincing potential buyers of their property to stay away.

Absent out-and-out lying, we still have the observation that health problems and opinions about the facilities are, at least sometimes, correlated. We should obviously expect to see this when all data is collected after the turbines are operating (which includes all the adverse event reports and most of the other studies): Anyone suffering health problems that they perceive to be caused by the turbines is going to have a negative opinion, and is likely to better recall any negative opinion they had before the facility was built. The causal pathway is quite plausible: the impact of the turbine causes health problems, which then causes the sufferer to dislike the facility. Even if disposition data is collected before the turbines start operating, there is still a good chance of causation running from health concerns to disposition. People who recognize, from experience or other self awareness, that they are more likely to suffer health effects from noise pollution are among those who will most strongly object to the siting and have negative feelings about it. Indeed, it seems safe to predict that a larger than average portion of the population with those feelings will be near new facility sites, since local residents have chosen to live in quiet rural areas. It is certainly the case that the average resident will be more sensitive than people who

self-select into noisy occupations (i.e., the people who are the subject of most studies of the effects of noise).

That said, it is not implausible that dislike of the facilities triggers or exacerbates health problems. Indeed, increase in distress caused by a stimulus due to frustration with having it imposed upon you, and not being able to do anything about it, is inevitable. But this is simply part of the causal pathway. That the pathway partially passes through local residents' dislike of the facilities is hardly an argument that the facilities are less damaging, though industry supporters sometimes suggest as much. The most charitable interpretation of the claim is basically, "people so hate having the audible, visual, and other effects of these facilities imposed on them so much that it ruins their health."

I am not sure whether this represents an addition to the industry's standard arguments, but I observed in one recent regulatory proceeding the industry's consultants focusing on the argument that local residents have an irrational fear of the turbines, like the fears that have been documented about radiation or chemical pollution, and that they can just be educated out of this. However, there is no basis for claiming that people have the same attitude toward simple noisy mechanical objects that they have toward mysterious invisible threats like radiation. Equally important, even if they do, there is absolutely no evidence that people can be "educated" into not being bothered by nearby turbines. If the industry can do this, they should be doing so in existing facilities, of course, rather than simply claiming it *could* be done as a basis for siting new facilities. But there is no evidence that they can or that they genuinely believe that they can. Unless it can be shown that this education reduces people's health problems, then the theory about the causal pathway is purely theoretical and represents a distinction without practical difference.

Epistemic conclusions about pathways

With the exception of damage to hearing, noise causes problems via mechanisms that we do not fully understand. Similarly, flickering lights cause well-understood health problems under a few circumstances, but the reasons that they bother people in other circumstances are simply not known. The conclusion that we should reach from this is not "we have no evidence that turbine noise would cause health problems" but rather "we know so little about how noises cause health problems that we will have to defer to the epidemiology in assessing whether they are actually occurring", as well as "we should be quite cautious about imposing novel noises on people because they could easily have unpredicted effects." The scientific reasoning is not that if we cannot figure out how the physical impact causes a disease then there is not disease, but rather if there is disease then it would be useful to figure out how the exposure is causing it.

It may be that many people outside of health science think that we know exactly how most diseases are caused by the exposures that cause them. In reality, most causal relationships have some black box in the pathway, especially when both the exposure and the outcome are complicated, as they are in this case. For example, we are only starting to learn why smoking causes as much heart disease as it does. We know some of the pathways and have guesses about others, but I suspect most people would be surprised that we are only just figuring out some key bits of the pathway. That, of course, does nothing to change our knowledge that smoking causes heart disease or to change our knowledge of how much heart disease it causes, which we

estimate based on the epidemiology. Indeed, the most useful knowledge that comes from discovering these pathways is that it might enable us to reduce the impacts of smoking.

The implication of last point is worth reiterating: Not understanding the causal pathways does not keep us from knowing that there is a problem. However, it makes it almost impossible to claim “if we just do X there will be no problem.” In other words, when industry proponents point out that we cannot figure out why there is a problem, they are really arguing that we do not yet have any way of knowing what regulations are sufficient to eliminate the problem.

As a final point about understanding causation, it is possible to provide an intuitive reality check for readers seeking to go beyond the abstract level. Chances are that sometime in your life you have experienced sensory distress from trying to do mental work near rumbling machinery or a ticking clock in an otherwise silent room, or trying to sleep in a hotel room where a LED was blinking on a television or smoke detector, or driving where the sun is flickering through the trees, or perhaps trying to relax quietly but being bothered by the muted sound of your neighbors arguing, their baby crying, or the bass on their stereo. All of these and countless similar low-energy sensory inputs sometimes bother us for reasons we cannot fully understand. Imagine, then, if whichever one of these you have experienced bothering your sleep, concentration, or mood invaded your home frequently and you could not get used to it. This is not to say that what is experienced by residents near wind turbines is exactly like one of these processes (the ticking clock would need to go “whoop...whoop...whoop” and the blinking light would need to be a passing shadow darkening the entire room many times per minute). It is not the case that most people are bothered by stimuli like these most of the time. The point is that you can recall one of these bothering you, and so should be able to imagine the impact of turbines. These stimuli would not damage your hearing or sight, trigger an epileptic seizure, or break your bones, but one of them might make you miserable and physically or mentally dysfunctional until you could get away from it. That is a serious health effect.

APPENDIX

My analysis of Colby et al (2009), as included in testimony from July 2010

Most of what I have to say about the scientific and pseudo-scientific claims made in scientific reports is addressed above, addressing the substance of the claims. However, there are a few comments that are worth making outside of the context of specific scientific arguments. [...] Because these might tend to influence policy decisions, even apart from their specific arguments that are addressed above, I believe it is important to further illustrate the failings of these reports as legitimate health science analysis. While many analysts insist on only writing competing monologues that address opposing arguments only obliquely, if scientists do not directly and explicitly confront scientific errors, non-scientific readers generally have a difficult time determining which of the competing claims is accurate. Thus, it is part of the duty of those offering scientific advice for policy making to explicitly explain why claims that they disagree with lack validity.

It is notable that the Colby et al. panel did not include any population health researchers, even though the question they claim the report addresses is one of population health. Their expertise seems to be limited to the relevant physical sciences and clinical medicine. This explains the dominance of physics-based analyses in the report, discussions that are interesting and informative in some ways, but have very little bearing on question of actual health effects. From the perspective of a population health analysis, those entire sections can be summarized by the following: Turbines make noise; noise often affects people's bodies and health, though in the present case – if it can be established that there are health effects – we have no idea exactly what physical pathway leads from the turbines to the effect. Beyond that, these sections appear to be little more than general background about the technology. More cynically, they might be seen as impressive-sounding filler that might lead the causal reader to think there is a lot more substance to the report than there really is.

Some clinicians are trained and qualified population health scientists, but there is no evidence that those who participated in the Colby panel have such expertise. Unfortunately, it is often the case that physicians who may be quite skilled in their clinical abilities do not recognize that population science requires an entirely different set of skills. (Moreover, in the present case there is something quite curious to a population scientist who has years of experience observing physicians making population science conclusions without understanding the science: Usually physicians can be relied upon to say that if they have seen a case of a disease then the disease exists, and then they often err by over-concluding (about prevalence and especially cause) based on nonsystematic observations. In this case, however, they seem quite anxious to claim that we have seen the cases but they do not really exist, a very un-physician-like behavior.)

Section 4.5 of the Colby report exemplifies the lack of expertise in population health science. It is difficult to make this clear without seeming petty, but this section reads like it was written by someone who took a single class on how to understand epidemiology, and half understood the material. Like most readers of epidemiology, they present statistical significance as the measure of the accuracy of a study. Setting aside the complexities and common errors associated with this concept, they make a huge error: When they invoke it they are talking about a case series where there is no statistic that even could be statistically significant (or fail to be); the concept does not even apply. They go on to characterize “statistically significant” (in this context where it does

not actually mean anything at all) as the opposite of the vague concept “simple coincidence” when it actually refers to a fairly technical test of how likely a particular observation is due to sampling error. But “coincidence” mostly invokes the situation where cases that occurred among the exposed population but not because of the exposure (as I use it above), a concept that is quite different from statistical significance or the lack thereof.

(For those who are interested, what the authors seemed to be trying to understand and explain is the contrast between associations (which cannot actually be calculated from the particular data) that have a low degree of statistical robustness – because they quite plausibly could have resulted from random sampling error – versus more robust results that we would not expect to see as a result of that random error alone (a particular technical definition of which is “statistical significance”).)

The question of whether “there is a causal relationship between the exposure and the disease” (p. 4-12) is not a matter of whether there are control subjects, as they characterize. Also, they use the word “uncontrolled” to refer to lack of comparison populations, though this is a strange term to use for this (when used in epidemiology, it almost always means the analysis does not include covariates to try to control for confounding or is used to differentiate a purely observational study from a trial/experiment when the investigator controls the exposure). They describe case-control studies in the (admittedly common) naive way – that the comparison is between people with and without the disease, while that is merely the way the data are collected and like most other epidemiologic study designs the comparison is between those with and without the exposure. But they also offer the strange characterization that such studies always match subjects on other variables that might affect the probability of having a disease, which is one option for doing such a study but not the only one. They claim to describe a cohort study (which follows a population to look for new events) but actually use language that better describes a cross-sectional study (a fundamentally different design which takes a snapshot in time looking for existing ongoing diseases). They imply that the only types of studies that exist are case-control and (the mis-described) cohort, ignoring cross-sectional, the usually favored randomized trial, the critically important case-crossover, and others.

These errors paint a picture of authors who are dramatically overstepping their expertise and hoping that no reader will ever have the expertise to notice and a forum like the present report in which to expose it. But even more important than these failures to understand epidemiologic methods, they fail to understand how to draw scientific conclusions in epidemiology.

In Colby et al.’s conclusions (section 5), even after citing many pieces of evidence that suggest turbines are having health effects, they repeat their “no evidence” claim. They apparently are basing this on the observation that “there is nothing unique about the sounds and vibrations emitted by wind turbines”, combined with the claim that “the body of accumulated knowledge about sound and health is substantial”. Their logic (they do not explain, so I must infer) seems to be “this is just like other noises” and “we already know everything there is to know about those other noises and they do not cause health problems”. The first of these is utter nonsense. While their physics studies may not be able to identify what the relevant differences are, anyone who understands epidemiology knows that similar exposures sometimes have quite different effects. No other exposure is going to be exactly like the noise from wind turbines. Indeed, the

“substantial” body of literature that they cite is not really all that impressive, and covers in depth only a few of the many forms of exposure to noise that people experience. In effect they are saying “we would not have predicted, based upon the limited analysis we can do using analogy and extrapolation, that health effects would have been observed, so we are going to insist that they really were not actually observed”.

Perhaps this can be attributed to just sloppy presentation of summary points. But deeper flaws in their scientific reasoning can be found. They claim that some quota of studies proves a causal relationship while fewer tell us nothing. Their example is that “multiple case-control studies were necessary before the link between smoking and lung cancer could be proved” (p. 4-12), when in reality the first English language studies (which appeared almost simultaneously) are generally regarded as being quite sufficient for reaching the conclusion. (Moreover, the example is a very poor choice since – as an expert in epidemiology would know – those were the studies at the center of establishing the validity of the case-control study in the first place, so any uncertainty was more about the study design, not its results.) Sometimes a single study is quite convincing, while other times a collection of studies leaves a lot of room for doubt. This is contrary to their assertion that “only after multiple independent-controlled studies show consistent results is the association likely to be broadly accepted”.

That last line is wrong at several levels. The first demonstrates their lack of understanding about what we are actually trying to infer (as well as the nature of epidemiology), since the question is not whether the association is accepted, but *causation*. The association is apparent in each individual study, or not. The phrase “independent-controlled” might appear to be jargon from the field, but it actually has no obvious meaning, while “broadly accepted” is a measure of public opinion, not scientific inference. Importantly, no studies ever prove causation – that is not how science ever works. So when Colby et al. denigrate case series data as not being able to “prove that an exposure is really harmful” (p. 4-12) they are saying nothing of substance. What they say that is of substance is that case series “can do no more than suggest hypotheses for further research” which is nonsense. To offer just one clarifying example, continuing on from a previous example, early case studies of tongue cancer cases in young people were sufficient to show that it was not being caused by tobacco use or drinking (on which most geriatric oral cancer was blamed), and later case series of oral cancer that tested cancer tissue for HPV provided very convincing evidence that that virus was often the cause. Whether a particular study provides useful information about a question is not a simplistic function of the study type. In this case, since the most important question is “does it appear that turbines may be causing diseases”, the case series is entirely adequate.

Indeed, this failure to understand what they are analyzing is worth emphasizing. Colby et al. write a (mangled) discourse on what is supposedly required to establish, “prove”, or make “broadly accepted” a causal conclusion, when the conclusions of their report are basically that there is no evidence that there is any problem and no reason to do further study. In other words, they lay out (what they claim is) the burden to prove a hypothesis of a particular exposure-disease relationship is true, but then try to use the (claimed) failure as the basis for saying that the hypothesis is false. It is possible that this is calculated misdirection, though my reading of their many failures of scientific reasoning suggest that this elementary error may well be inadvertent. It is really difficult to believe that people who wrote that section have any

understanding of epidemiology. And since epidemiology is the core science for understanding human health effects, it seem rather odd that this report is characterized as “an expert panel review” of “health effects”.

My comments on Australian government (2010a, 2010b) documents, as included in testimony from September 2010

The Australian Government (2010a) recently published Wind Turbines and Health, A Rapid Review of the Evidence. At the risk of seeming glib, perhaps they should not have conducted their review quite so rapidly. The first paragraph concedes that the report relies upon the highly flawed Colby et al. report. Where the Australian report simply repeats one of the claims from Colby et al., I refer to the appendix that contains my original criticisms of Colby et al. and those appearances in the above text. This is both for efficiency and because the Australian report seems to have garbled some of the points and scattered them throughout the document, and so the points can be better addressed as they appeared in the original. This channeling of Colby et al. includes referring to a “nocebo” effect, dismissing the problems by manipulating the word “annoyance”, implying that the heterogeneity of effects across people and that effects are associated with individual characteristics somehow makes them less real, and suggesting that the correlation of health problems and general attitudes toward the facilities makes the effects less real.

The report grossly oversimplifies the scientific discussion by saying “there are two opposing viewpoints regarding wind turbines and their potential effect on human health”. The anonymous authors maybe have been trained as journalists and were following the journalistic tradition of trying to figure out how to turn any complicated disagreement into something as simple as a football match. In the present situation there are two opposing viewpoints only in the sense that there is a political faction that insists that there is absolutely no risk and that nothing should be done to further explore it and there are those who disagree with this dismissal. Within the latter category, however, there is far wider distribution of belief in what, exactly, the current evidence shows than the attempt to portray it as a football match implies. Indeed, the latter category includes those who do not believe that the health effects are great enough to warrant new policies but who nonetheless recognize the absurdity of the claim that there is absolutely no risk and that we should avoid learning more as well as those who are convinced that the effects are so great that we should impose an immediate construction moratorium until we can better understand and mitigate the risks.

Moreover, after asserting that there are two sides, the report then goes on to make clear which team it is cheering for.

Addressing the effects of noise, the report quotes the assertions from several reports that basically say “there are no effects”. What it claims to present for the “opposing view” is an overly simplistic and misleading representation of Pierpont’s book, trying to imply that she is alone in her opinions and dismissing her analysis by asserting that her conclusions “have been heavily criticised by acoustic specialists”, though how exactly acoustic specialists could criticize her epidemiologic study is not entirely clear. Following that, and with no further analysis, the authors conclude “based on current evidence, it can be concluded that wind turbines do not pose a threat to health if planning guidelines are followed”. What analytic method the authors used

for reaching that conclusion is left to the readers' imagination. This requires quite a bit of imagination since the authors do not appear to have actually reviewed the evidence beyond gathering a haphazard collection, and there is no evidence in the report that they even read more than the abstracts and conclusion statements of what they gathered, let alone critically assessed them.

The analysis of the effects of flickering is basically limited to attacking the strawman argument that no one makes, that the flickering causes epileptic seizures. The section on mitigating the effects (notwithstanding their conclusion that there really are not any effects) basically just quotes a few people who say we should follow the guidelines, and calls for teaching people it is really all in their head.

The report also includes some strange gaffes. The authors claim that Pederson and Persson Waye (2007) argue that noise from wind turbines does not cause physiological effects such as anxiety, tinnitus or hearing loss; whatever those authors might believe about this point, that article does not include any such argument. Though the report is about human health, the authors spend more words talking about the effects of electromagnetic radiation on communications equipment than they do about its effect on health (the latter is another strawman). The conclusion section consists of irrelevant statements about the health effects of energy sources in general, followed by an assertion: "There are no direct pathological effects from wind farms" (they fail to explain what they mean by the caveat "direct" or why they chose the rather odd word "pathological"), which they immediately contradict by concluding the sentence with "and that any potential impact on humans can be minimised by following existing planning guidelines" (how is it possible to minimize something that does not exist?).

The Australian Government (2010b) also published an accompanying public statement. This is basically a three page advertising pamphlet centered around the absurd claim, "there is no published scientific evidence to support adverse effects of wind turbines on health". The document is basically cheerleading, with no substantive analysis. Strangely, in this document they also mis-cite Pedersen (2007), but this time as a source for the claim that there is popular demand for green energy. They also uncritically quote the line from Colby et al. about there being no evidence of physiologic effects, the one that is contradicted by the very next sentence as I discuss above. The balance of the document is basically the standard industry assertions distilled into an advertising flyer (e.g., the noise is too low energy to cause damage; the flickering cannot cause epileptic seizures and so must be harmless) that are all addressed elsewhere in my report. Since there is nothing novel or analytic in this document, there is really little more I can say about it.

In sum, the Australian government's reports appear to have been thrown together by people who are not experts on health science, spent only a few days doing the research and writing, and relied on the industry to provide them with most of their conclusions. There is simply no reason to take any of their conclusions seriously, and there is no actual analysis, just repetition of assertions, so it would be impossible to take their analysis seriously.

References

Australian Government, National Health and Medical Research Council. Wind Turbines and Health; A Rapid Review of the Evidence. July 2010a. Available at <http://www.nhmrc.gov.au/publications/synopses/new0048.htm>

Australian Government. NHMRC Public Statement: Wind Turbines and Health. July 2010b. Available at <http://www.nhmrc.gov.au/publications/synopses/new0048.htm>

Castelo Branco NAA, Alves-Pereira M. Vibroacoustic disease. *Noise & Health*, 6(23):3-20, 2004.

Chief Medical Officer of Health [of Ontario]. The Potential Health Impact of Wind Turbines. May 2010.

Colby WD et al. Wind Turbine Sound and Health Effects; An Expert Panel Review. Prepared for: American Wind Energy Association and Canadian Wind Energy Association, December 2009.

Krogh C et al. Wind Vigilance for Ontario Communities; A Self-reporting Survey: Adverse Health Effects, Industrial Wind Turbines (IWT) and the Need for Vigilance Monitoring. 2011. Available at windvigilance.com.

Harry, A. (2007). Wind turbines, noise, and health. February 2007. http://www.wind-watch.org/documents/wp-content/uploads/wtnoise_health_2007_a_harry.pdf

Heavner K, Phillips CV, Rodu B. Peer review in epidemiology cannot accomplish its ostensible goals due to incomplete reporting and unverifiable analyses. Proceedings of the International Symposium on Peer Reviewing, at the 3rd International Conference on Knowledge Generation, Communication and Management, 2009. <http://www.iiis.org/CDs2008/CD2009SCI/ispr2009/PapersPdf/V524QH.pdf>

Maclure M. The Case-Crossover Design: A Method for Studying Transient Effects on the Risk of Acute Events. *American Journal of Epidemiology* 133(2):144-153, 1991.

Michael A. Nissenbaum MA, Aramini J, Hanning C. Personal communication based on a draft of the currently confidential “Adverse Health Effects Related to Industrial Wind Turbines (IWTs) - a Retrospective, Cross- Sectional Epidemiological Study”. 2011.

Pedersen E, Waye KP. Perception and annoyance due to wind turbine noise; a dose-response relationship. *Journal of the Acoustical Society of America*. 116(6):3460-3470, 2004.

Pedersen E, van den Berg F, Bakker R, Bouma J. Response to noise from modern wind farms in The Netherlands. *J Acoust Soc Am*, 126(2):634-43, 2009.

Pedersen E, Persson WK. Wind turbine noise, annoyance and self-reported

health and well-being in different living environments. *Occup Environ Med* , 64(7):480-6, 2007.

Pedersen E, van den Berg F, Bakker R, Bouma J. Can road traffic mask sound from wind turbines? Response to wind turbine sound at different levels of road traffic sound. *Energy Policy* 38(5):2520-2527, 2010.

Phillips CV. The economics of ‘more research is needed’. *International Journal of Epidemiology* 2001;30:771-776.

Phillips CV. Analysis of the Epidemiology and Related Evidence on Health Effects of Wind Turbines on Local Residents. July 2010. Available at <http://www.windaction.org/documents/28175> (which implies no affiliation with that group – they just have a good archive)

Phipps R. In the Matter of Moturimu Wind Farm Application. Evidence to the Joint Commissioners, Palmerston North. March 2007. <http://www.ohariupreservationsociety.org.nz/hipps-moturimutestimony.pdf>

Pierpont N *Wind turbine syndrome: a report on a natural experiment*. K-Selected Books. 2009.

Qibai CY, Shi H. An Investigation on the Physiological and Psychological Effects of Infrasound on Persons. *Noise Notes* 4(2):25-30. June 2005. <http://www.windturbinesyndrome.com/news/wp-content/uploads/2010/12/An-investigations-on-the-physiological-and-psychological-effects-of-infrasound-on-persons.pdf>