The affirmative ethical arguments for promoting a policy of tobacco harm reduction

Carl V Phillips

28 July 2009

Background

Tobacco harm reduction (THR) is the substitution of low-risk nicotine products for smoking, with the expectation that users will self-administer the substitute products for a long period as an alternative to smoking rather than a short-term bridge to total nicotine abstinence (see our website, TobaccoHarmReduction.org, for more information). A policy promoting THR contrasts with the standard abstinence-only approach to nicotine use that dominates discourse on the subject, where substitute products are encouraged only for very short term use for weaning the user off of nicotine entirely. Whether we should promote THR as a matter of public policy is the subject of contentious debate. However, the normative ethical arguments that are made in that debate are usually left implicit, are unexplored as if they are self-evident, and are overshadowed by discussions of empirical results and assertions, which are often presented as if they substitute for policy analysis. As a result, the normative arguments are muddled to the point that it is difficult to identify them.

Statements made by both proponents and opponents of THR almost never clearly distinguish among: (a) affirmative arguments that favor promoting THR, (b) arguments that THR has disadvantages that offset the advantages, (c) empirical questions, usually epidemiologic, on which one or more of the arguments hinge, and (d) different notions of what a pro-THR policy would entail. In particular, what is most often presented as ethical analysis is usually just an extensive examination of an empirical claim. Normative conclusions are then asserted as if they simply follow from the empirical observations, leaving both the basis for making ethical claims and the connection between the empirical results and the normative claims unexamined.

For example, epidemiologic research reports written by political opponents of THR that suggest that there is some small risk from low-risk nicotine products never include an analysis of the normative basis for promoting THR or how their result relates to normative arguments, but nevertheless quite often conclude that they have provided an argument against promoting THR. In the few cases where an opponents of THR purport to be analyzing the arguments in favor of THR, the presented arguments are generally a straw man, intended to misconstrue the arguments in favor of promoting THR or misconstrue the concept of THR itself (e.g., Tomar (2009); Zeller et al (2009); for an annotated version of the first of these that points out the straw man arguments and other weaknesses, see Phillips et al (2009)). However, blame for opponents' mischaracterization of the arguments in support of THR must be shared with proponents who do not effectively, clearly, and completely present the arguments in favor of particular policies.

Authors of analyses and commentaries arguing in favor of THR typically seem to assume there is consensus about how to identify a good policy, as well as about what specific policy is being discussed. This is probably adequate in most contexts: The experience of those of us who promote THR is that when a description of the current situation and the existing scientific evidence is presented to someone who does not have overriding political beliefs about substance use in general or nicotine/tobacco in particular, they immediately think that promotion of THR is a good and proper idea. Intuitive notions of what constitutes a measure of the good and proper, and what would constitute the promotion of such, are adequate for these educational efforts. However, a serious policy analysis requires making these intuitive notions explicit. This becomes particularly important when there is contentious political debate that obscures the core points. It is not, after all, possible to assess the goodness of a policy without first establishing what constitutes goodness and exactly what the policy is. Scientific empirical claims are not normative arguments; they can only inform the policy question after establishing the policy and ethical framework.

To the extent that authors do implicitly identify what they think is the measure of a good policy, it is usually the reduction of disease or increasing of longevity. As argued below, the case for THR can clearly be made based on this, but it is the least compelling of the main arguments from an ethical perspective, and the most sensitive to differing empirical claims. But even to the extent that a goal is stated or implied, there is little attempt to identify exactly what specific policy alternatives are supported by the goal and accompanying empirical claims.

Perhaps the most complete previous attempt to balance the arguments and emphasize those that are most compelling can be found in our website (Phillips et al., 2009). However, since that information is presented in the form of educational information pitched to the consumer, and organized in support of this goal, the fact that it is also an attempt to present clear and comprehensive policy-ethics analysis has not been widely recognized. To address this deficit, this paper presents a more formal and signposted version of the main affirmative ethical arguments in favor of THR.

Structure of this analysis

To be able to focus on the affirmative arguments for THR, we have presented them here in as much isolation as possible. No attempt is made here to establish the empirical claims that underpin the arguments since it would be impossible to do justice to such a complex topic as an aside in an ethical analysis (just as it is impossible to do justice to the ethical analysis as an aside in an empirical paper). Instead, relevant empirical claims are explicitly posited in the first section. Readers will have to look elsewhere to assess the validity of the claims (summary analyses and sources for the relevant scientific information can be found at our website's FAQ section (Phillips et al. 2009), as well as Phillips et al. (in press) and Rodu & Godshall (2006)). For empirical assumptions, specific citations are offered for the convenience of the reader when the claim is not part of the general background knowledge that can be found in most analyses of THR, but no attempt is made to argue the claim. Sometime analytic studies that establish a claim are

cited. However, when an empirical claim cannot be legitimately supported by citing a small number of references (e.g., claims that depend on synthesizing a large body of empirical research reports), no citations are presented; the common practice of presenting one or two references for such claims merely give readers the illusion that the claim is being properly argued rather than just posited, and that is avoided here. When an ethical argument depends substantially on an empirical point where there is genuine scientific doubt, this is noted and the implications of alternative empirical assumptions are analyzed.

Similarly, presentation of and responses to claims of THR's possible negative consequences (that might weigh against the affirmative arguments) are omitted. Negative consequences that directly subtract from the objective function being presented are necessarily included (e.g., possible negative health consequences when the objective is improved health), but orthogonal negatives that might weigh against the affirmative arguments (e.g., claims that all drug use is *per se* immoral, and thus facilitating it is inherently bad) are catalogued and analyzed in a separate paper.

Policies to promote THR can take many forms, and not all are supported by all affirmative arguments. Thus, this analysis begins by positing some basic empirical points that are necessary background, and then identifies a hierarchy of potential policies. It then presents the three distinct affirmative arguments: the deontological argument based on individual rights and autonomy; the consequentialist argument based on maximizing individual (and, thus, social) welfare; and the consequentialist arguments based on an objective function of maximizing individual or public health.

Empirical Background

The central empirical claim that is necessary to posit for any pro-THR argument is that THR is possible. That is, available alternative nicotine products are substantially less harmful than smoking cigarettes, and they have enough appeal that many smokers would switch to them if they were aware of the comparative risk (or if they were pushed to switch, even ignorant of the comparative risk). Such products include smokeless tobacco (a.k.a., snuff, snus, chewing tobacco), electronic cigarettes, and pharmaceutical products (see Royal College of Physicians (2007), Rodu & Godshall (2006), and Ballin (2007) for details about products), though specifying exactly what products qualify is not necessary for present purposes. To avoid complication, the products fitting this description will simply be referred to collectively as "THR products"; some might be better than others by some measures, but that is not addressed here.

The exact magnitudes of "substantially less harmful" and "many smokers would switch" have some implications for the arguments below and are addressed in context. THR proponents usually concede that most smokers find smoking to be a superior experience compared to most or all THR products, but argue that many find (or would find if they tried it) at least one product to be somewhere between tolerable and almost as good. Opponents have not substantially disputed this claim. As background (though it is not critical to any of the arguments), it is worth noting that the various THR products are equivalent in terms of health risks (i.e., there is no evidence of any differences among

them, and good evidence that they are quite similar), so THR proponents encourage consumers to choose whichever product they find best.

The typical comparative risk quantification offered by proponents of THR is that THR products are roughly 99% less harmful than smoking (opponents typically do not offer a quantification). This claim is widely repeated and accepted, though we should disclose a bias for the 99% figure; it appears to always trace to one of our calculations (Phillips et al. 2006), since we are not aware of any other published calculations that present that figure. The other often repeated estimate, 98% less harmful, can be traced to our calculation or others that predate it (e.g., Rodu, 1995). Importantly, a much smaller reduction is also sufficient for any of the arguments below; if the risk reduction from THR products is merely a 90% reduction compared to smoking – and no one seems to seriously claim it is any higher than that – all the claims are still supported. Most would be supported if the figure were merely 50%.

It is useful to observe that the 99% estimate means that for the average smoker, smoking for one more month creates more risk of mortality than switching to a THR product and never quitting (see Phillips 2009 working paper for the calculations). As context, the portion of established smokers who quit each month is well under 1%.

It is also necessary background to know that most people (including smokers, but also opinion leaders who influence what smokers believe) are not aware that THR products are so low risk. Were this otherwise, there would be no reason to focus on policy options that involved informing smokers about THR products and their low risk. Identifying the reason for this widespread ignorance is not necessary for presenting the affirmative arguments themselves. However, the question of which policy is warranted is influenced by the following observation: Proponents of THR have accused most of the major actors in the anti-tobacco arena in most Western countries – government agencies and quasi-governmental actors (major anti-tobacco organizations that are closely affiliated with the government and substantially financed by taxes) – of intentionally misleading consumers into believing that THR products cause risks similar to those from smoking, and so there is no low-risk option.

For the argument about maximizing health (though not the other two arguments), it is necessary to posit that people will continue to use nicotine, and that this will be primarily in the form of smoking unless harm reduction efforts succeed. That is, there is no serious prospect that nicotine use will be eliminated by voluntary or forced changes in behavior.

One category of argument directly hinges on (and another makes little sense without) the assumption that a substantial fraction of smokers get some benefit (not necessarily net benefit) from smoking, and that they would still get some or all of that benefit from THR products. It is generally assumed that these benefits come primary from consuming nicotine. Some THR proponents (and others who have studied nicotine) argue that established benefits from its use include aiding focus, alertness, relaxation, weight control, and a relief from various psychological pathologies, as well as less specific contributions to pleasure. For the present purposes, it is not necessary to posit the exact

form of the benefits, nor is it necessary to assume any particular magnitude of benefit. Note that this assumption specifically excludes the claim that the benefits from smoking are substantially due to aspects of smoking that would not be provided by any THR product; it is often claimed that *some* of the benefits of smoking fit this description (which is not inconsistent with any of the arguments below), but there is no empirical support for the claim that *most* of the benefit comes from these aspects.

While the assumption that people get some benefits from a behavior that they engage in at great cost might seem blindingly obvious to readers who are not familiar with the discourse over tobacco, it is useful to point out that this is not generally recognized within the politicized discourse. To explain the behavior, THR opponents often appear to work from a premise (generally without trying to defend it, almost never explicitly stating it) that smokers get no benefit from smoking and continue to smoke only because of some involuntary compulsion. Thus, there seems to be an implicit claim (though, again, no explicit stating and defense of it) that smokers would experience a welfare gain (or at least, no loss) if they were prevented from doing what they are currently choosing to do, even apart from eliminating the health costs. This possibility is sufficiently implausible – a violation of everything we know about human nature and economics – that it is ignored here; it will be addressed as a possible counter-argument in the next paper.

It appears that former smokers find it easier to quit THR products after switching to them, than they found it to switch from smoking to abstinence. This is based on minimal evidence, however, so it seems unwise to make any arguments that depend on it. However, there is ample evidence that THR products are, at worst, not substantially more difficult to quit to abstinence than are cigarettes. This observation is useful, but not critical, for one of the points below.

It is sometimes hypothesized that implementation of THR would cause more people to smoke. Positing the degree to which this is true is necessary for one of the arguments below. There is no empirical evidence to suggest that promoting THR has ever increased smoking prevalence, but the possibility is discussed in context below. Often opponents of THR demonstrate confusion about the relevant point, arguing that use of THR products (in the absence of active promotion of THR) has been shown to cause users to take up smoking later, but careful analysis of the evidence shows that it does not support even this claim (Timberlake et al., 2009; Bates et al., 2003; Phillips, 2004).

<u>Spectrum of pro-THR policies</u>

Potential pro-THR policies vary widely, and each argument supports certain policies over others. Confusion about what policy is being discussed muddles many discussions in the literature and in the political arena.

It is a bit artificial to force the list onto a single spectrum since orthogonal elements of the policies could theoretically be separated, and thus the policies could have completely independent degrees of honest provision of information, availability of products, and so on. However, many combinations (such as nudging people to adopt THR via pricing while not informing them of its benefits) seem unsupportable by any ethical argument,

and even absurd, and so their omission is not a problem. The simplification to a single dimension with artificial bright lines is useful for presentation and does not seem to interfere with the analysis, since every policy that is identified as being supported by a particular ethical argument is included.

1. At its strongest, a policy of promoting THR would force, or nearly force, smokers to switch to THR products. This would include the provisions of policy 3 along with substantial pushing to switch. This has never been seriously proposed, but as noted below, at least one set of ethical views and empirical beliefs would call for it.

2. An active but more libertarian approach is to provide incentives that nudge (Thaler & Sunstein, 2008) smokers toward THR products in addition to the provisions of policy 3. This is widely advocated, particularly in the form of making sure the purchase price is lower for the low-risk products, through taxation or other policies.

3. The strongest policy that does not involve directing people to a particular behavior is for the government and its surrogates to actively provide accurate information about the benefits of THR and ensure the easy availability of THR products, but nothing more.

4. A weaker policy would be for government and quasi-governmental actors to not take pro-THR steps, but to refrain from promulgating misleading information that discourages THR. Proponents and private opponents of THR would be free to make their case.

5. Weaker still, to the point that it could barely be considered a pro-THR policy, would be to merely allow pro-THR information to compete in the marketplace of ideas and to make sure that THR products were legally available as competitors to cigarettes in the marketplace; government and its allies could continue to take all other anti-THR actions. Readers unfamiliar with the political reality might think that such conditions would be the most anti-THR policy anyone would seriously consider, but it should not be taken for granted. For example, our own experience: Anti-THR activists (often backed by public funds and even including government actors), have repeatedly tried to censor and shut down our research and education group, which is arguably the only educational organization in Canada devoted to promoting THR (Phillips, 2007; Phillips, 2008; Libin, 2007; Balfour et al., 2009). Such attempts at censorship of accurate information have also occurred to a lesser degree in the United States and elsewhere. Governments have banned or made inaccessible the most promising THR products (sometimes snus, sometimes electronic cigarettes, and sometimes both) in many places where smoking is common and legal, including the European Union, Canada, Australia, New Zealand, and elsewhere.

6. An odd policy combination that could not be considered pro-THR, but is useful for analytic completeness, would be to ensure access to accurate information, but prohibit or severely interfere with THR product availability.

Obviously these brief policy descriptions leave out many details that would need to be addressed for a practical analysis. For example, access and product availability do not

imply that children should have free access to purchase these products. The right to provide accurate information does not necessarily mean manufacturers would be allowed unfettered commercial speech to make accurate health claims. Details about such limitations, and others, would have to be determined.

Argument: People have a right to know the truth and make their own decisions about THR

The preeminent principle of Western health ethics for more than half a century has been that of informed autonomy: People have a right to make their own health-affecting decisions and to be provided with the information necessary to make those decisions (Beauchamp & Childress, 2009). Assigning a right to one individual is synonymous with imposing obligations on others, and in the present case for the rights to be meaningful the obligation must apply to governments and other powerful organizations that influence education and policy regarding nicotine products. It is difficult to identify any specific interpretation of this ethic in the present context other than the following: People have a right to choose THR products, whether they are using them in lieu of smoking or not, and they should be accurately informed about how risky they are, particularly in comparison to smoking. It follows immediately that policy 5 or something higher in the above list would be warranted as necessary to preserve autonomy. The right to accurate information further implies that public actors should be forbidden from providing disinformation, ruling out policy 5 in favor of 4 or 3. The principle of informed autonomy seems to be silent regarding policy 2. Policy 1 would seem violate this right, since an informed individual who still chooses to smoke instead of switching should not be prohibited from doing so.

Actively pushing particular information for making health-affecting decisions is not necessarily the obligation of government or quasi-governmental organizations, even in light of the right to informed autonomy. People make dozens of such decisions per day (how to travel, what products to buy, how much to eat, etc.), and for many the relevant health information is never actively delivered. However, government and other organizations have actively adopted the role of advising people about the risks from smoking and other nicotine use, and thus have accepted the obligation to provide accurate information autonomous decisions. Moreover, if the empirical claim is accepted that these actors have actively misled people into false beliefs about the possibility for THR, their obligation extends to actively correcting their previous violations of people's rights. Presumably such attempts to mislead people were motivated, at least in part, by the goal of getting people to do the "right" thing "for their own good". But since intentionally misleading people with false information into making a particular health-affecting decision, whatever the motivation, is a clear violation of the informed autonomy right, there is an obligation to correct the effects of such behavior from the past. This suggests that policy 3, rather than just 4, is demanded in the majority of societies where the government and its surrogates have adopted anti-smoking policies, and especially where they have misled people about THR.

In sum, the strong form of the affirmative argument in favor of THR that is based on the right to informed autonomy is as follows: Individuals (including but not limited to

smokers) have a right to choose to use THR products if they want, and to know what the actual risks of those products are. This forbids interfering with the availability of products or providing disinformation to manipulate people's choices. Moreover, because there has been so much disinformation (which would likely continue from some sources) resulting in widespread ignorance about the actual level of risk, government and its allies have an affirmative obligation to push accurate health risk information (policy 3). A case can be made that there should be some encouragement for smokers to adopt THR beyond simply promulgating accurate information, to overcome the inertia created years of disinformation and smoking -- in effect, to mimic the situation that would exist if informed autonomy had always been respected. But a case can also be made that this is just an alternative manifestation of the "for their own good" violation of autonomy, so it is not clear whether policy 2 is supported by this ethical principle.

Is drug use an exception to the informed autonomy principle? It should be recognized that actual practice shows that the right to autonomy over one's health is not considered absolute in most Western societies. Possession or use of some drugs is prohibited due to the effect they have on users' health. (At least, that is the ostensible reason for the prohibition. Many critics disagree, pointing out that the policies are not health-maximizing and the rhetoric often invokes deontological claims, and thus health is usually a stalking horse for other motivations.) This makes the exercise of some health-affecting choices impossible if the law is obeyed or enforced. While some people object to such prohibitions in all cases, and call for a nearly complete autonomy, there is sufficient support for the existing policy regime that an ethic of complete autonomy in matters of drug use should not be assumed. What are the implications for the pro-THR argument if this weaker version of informed autonomy (that allows prohibition of the choice to use particular psychoactive chemicals) is accepted as the underlying ethic?

If prohibition is considered ethical for some drugs, then it cannot simply be ruled out for THR products based on an ethic of autonomy. However, in free societies and open governments, there is no ethical basis for the government or its surrogates to mislead people about health risks, let alone to interfere with others who are providing accurate information. When government or quasi-governmental actors mislead the public to justify warfare, infringement of civil liberties, taxes, or other policies, it is decried by those concerned with ethical governance (not always at the time – narrow political interests often prevent ethical analysis when the issue or personalities are still in play – but almost universally with the distance of history). Thus, it could be argued that the principle of informed autonomy does not necessarily support any policy higher on the list than policy 6.

However, it is difficult to formulate a justification for prohibiting use of THR products by arguing that using them is a health-affecting decision that is so universally and tremendously terrible that no one should be allowed to make it, for several reasons. Such claims are generally justified only for the most extreme cases where it is claimed that the drug in question makes someone dysfunctional (depriving the community of a substantial portion of their potential productivity), dangerous to those around them, inclined to criminal behavior, or doomed to a life that is nasty and short. None of these are the case

for THR products, which (based on the posited empirical claims) actually increase productivity and have a trivial effect on life expectancy. Moreover, any argument that people should not be free to choose THR products must be at least as compelling for disallowing access to cigarettes, which remain legal and widely available. While a case could be made that prohibitions against drug use should extend as far as the banning of tobacco smoking, no one has seriously proposed it, so there cannot be much popular belief that this is an ethical position, and thus even less for THR product prohibition.

Thus, the drug prohibition exception to the right to autonomy could never justify providing false information or interfering with others' freedom to provide accurate information about THR. Moreover, the observed applications of this exception are products whose effects are completely different from THR products, not even extending to smoking. Thus, there seems to be little basis from retreating from the demand that policy 3 (or perhaps policy 2) should be implemented.

Argument: Promoting THR will improve welfare

It should not be surprising that an argument from the perspective of informed autonomy and an argument based on maximizing welfare reach the same conclusion. Economic theory shows that under certain conditions, which are never perfectly met but are often approximated, increasing information and autonomy improves welfare. Those conditions include sufficiently low transaction costs (i.e., the welfare loss from the costs of learning and choosing is low compared to the advantage of making the personally optimal choice), informed welfare-maximizing choices can be made by individuals on their own (in particular, they can understand the information), and the choice does not produce extensive negative externalities. The latter of these conditions is simple: Unlike smoking, THR products do not produce any obvious negative externalities, so the effect on externalities is entirely beneficial, in the form of reducing second-hand smoke exposure. The other two conditions are addressed below.

When those conditions are substantially met, a simple analysis of welfare maximization shows that adding options or providing accurate information, is always (weakly) better for everyone, since any inferior option or irrelevant information is just ignored, while any superior option or decision-useful information increases welfare. This results in what is known as a "Pareto improvement" in economics – a change that makes no one worse off and at least one person better off; from any welfarist perspective (i.e., when the goodness of an action is measured based on its impact on people's welfare), such a social welfare improvement with no "distributional" issues (i.e., a need to weigh one person's loss against another's gain) is indisputably good unless it crowds out an alternative improvement that is even better. Thus, the simple analysis says that every policy up to policy 3 (policy 2 is addressed below) is justified, since they are Pareto-improving and it is difficult to imagine how they could crowd out any other welfare improvement.

This analysis is not dependent on details of empirical claims, as is the "maximize health" objective, discussed below. Everyone who voluntarily adopts THR products must be improving his welfare. Every smoker who switches represents a benefit primarily due to health risk reduction. But also, every nonsmoker who takes up THR products or smoker

who was destined to quit but switches instead also represents a welfare gain. While there may be some minor reduction in life-expectancy, those individuals still had abstinence as an option but chose the THR product instead, so using THR products must have provided net greater utility.

This introduces an observation that is undoubtedly true but strangely controversial: Honestly informing people about the low risks from THR products (which most people now incorrectly believe are roughly as harmful as smoking) would inevitably lead some individuals to use the products who would have otherwise been abstinent. These are nonsmokers or would-be quitters who find nicotine (or some other aspect of product use) to be insufficiently beneficial to justify the great risk from smoking, but sufficiently beneficial to justify a tiny fraction of that risk. There are often disputes about this prediction, as obvious as it is. Anytime the cost of something is lowered, we expect consumption to increase**.** Put another way, we know that many people quite rationally quit or avoid smoking because of the health risk, in spite of its benefits, so presumably many of them would quite rationally choose to use a substitute for smoking that had much of the benefits and a much lower cost (i.e., health risk). Indeed, smokers have a history of switching to products that they thought (usually mistakenly) were lower risk (e.g., Zeller et al. 2009). While improving the welfare of nonsmokers is not usually considered a goal of THR, letting people know about a potentially appealing option must be either neutral or positive for their welfare; this contrasts with the perspective when only longevity is considered, where such effects are considered negative.

To expand beyond the simple assumptions, it should be recognized that the "more options and information are better" principle is not always true. The possibility of a welfare-improving change in consumption can be outweighed by the costs of learning and choosing, and sometimes the learning is not successful, resulting in poor choices. However, in the case of one of the most important consumption decisions someone will ever make (smoking, THR product use, or abstinence), the cost of learning is clearly justified. The information is not difficult to understand, and can be made extremely simple without becoming inaccurate. This is not analogous to the cases of choosing among different insurance plans (where the differences among risks and costs are quite large, but arcane) or fifty ever-changing varieties of toothpaste (where the differences among them are too small to be worth the effort to make an assessment). Indeed, the more accurate the information, the easier the choice -- e.g., if consumers were told that all THR products pose approximately the same risk then they would recognize that the choice among them was no harder than picking the toothpaste that tastes best or is cheapest. Consumers with greater sophistication could, of course, spend tens or hundreds of hours optimizing their decision (as they might when buying a car, a choice with similar implications for health risks and pleasure), given that it is important and accurate information is available. Thus, the conditions for the information and choice being necessarily beneficial (no worse than neutral) seem to be met.

Another possible concern is that a particular option will lead someone down a path they will later regret, and so should not be permitted in the first place. As discussed above, it is difficult to see why such a consideration – which might justify disallowing the decision

to use crystal meth or drop out of grade school – could ever apply to a low-risk alternative to a common legal product. It is theoretically possible that nonsmoking adopters of THR products could suffer a net loss from their habit even though they thought they were benefitting, but there is no evidence to suggest this is the case. While it is true that people sometimes make bad choices, centuries of myriad evidence show that individuals are still more likely to optimize their consumption tradeoffs than are governments or other authorities trying impose better choices, except in the most technically arcane cases like which surgical procedure to choose. Since the health risk or expenditure from THR product use would be of similar magnitude of other health-affecting decisions, to consumption of dessert, recreational transport, french fries, sports, or beer, they are clearly not of a magnitude that should trigger paternalistic interception of adults' own decisions. (Note: It is inevitable that our future selves will wish they could change many of the choices our present selves make to enjoy life rather than investing in our future, whether it be taking leisure time or using drugs. This inevitable regret is sometimes mistaken for evidence that we did not maximize our lifetime welfare, perhaps even that people should be prevented from making particular decisions, but it actually has almost no information value.)

Promoting welfare maximization may or may not justify a bit of gentle pushing by those of us who have greater expertise and have thought more about an issue than the average consumer. If consumers always process all available information and make calculated decisions then those of us concerned with improving public health would want to only provide information and options. For most decisions, however, it is unrealistic to expect such attention, and demanding it for minor decisions would just be a waste of people's time. Thus, for minor decisions there is a strong case to "nudge" people, allowing them to overrule our suggestions if they choose to, but pointing them in what we think is the right direction (a case for doing so can be found in Thaler & Sunstein (2008)). It is difficult to see the harm in guessing that the average smoker would have much greater welfare if he switched to a THR product, though he may not realize it yet; his expected health would be much better and his benefits of consumption would only decrease a bit. Moreover, he could switch back if this turns out to not be the case. So instead of just educating him, we might nudge him toward considering THR (e.g., providing free THR product samples with cigarettes, keeping THR products cheaper than cigarettes). On the other hand, since smoking is such a major decision, consumers have strong incentives to be well-read and think carefully about their decisions regarding it. If they all do this, then too strong a push could risk lowering welfare by altering thoughtful decisions, ruling out policy 1. The case for nudging (policy 2) seems stronger than the suggestion it might be costly (arguing for policy 3), but the argument is not definitive.

Is there any basis for questioning the conclusion that welfarism supports promoting THR? It is difficult to identify any weakness in this argument. Both consumers who take advantage of the new opportunity to consume nicotine with low risk and those who quit smoking are benefiting from the knowledge that THR products are low risk and the availability of those products. None of the conditions that often render free choice and free consumer markets less beneficial than theory suggests seem to apply (externalities are reduced, the choice is understandable and simple).

There are rare extreme individual choices where we refuse to count as a positive the individual welfare benefits, and thus improving welfare is not necessarily good. It may be that some fierce opponents of nicotine/tobacco actually believe this applies to using THR products. But such exceptions are generally reserved for cases with extreme negative externalities (e.g., the benefits experienced by a child molester are not seen as lowering the net social costs) or that trigger our strongest instincts of disgust (e.g., benefits from freely-chosen incest and suicide are often not counted on the positive side of the ledger). There is no basis for suggesting that benefits from using mild, productivity-enhancing psychoactive products in ways that do not interfere with operating machinery or cause social disturbance fit these exceptions. Arguments that THR products have negatives that should be weighed against the welfare benefits is taken up in the subsequent paper, but there is no accepted basis for arguing that the welfare benefits are not legitimate.

Argument: Promoting THR will improve health outcomes

The argument that promoting THR will improve public health (generally implicitly defined as average population longevity) is less defensible as an ethical goal and far more complicated to make than either of the above arguments. Nevertheless, most implicit and many explicit ethical arguments for or against THR seem to assume that some form of this consequentialist argument is the only basis for making normative claims about THR. Thus there is great practical value in providing this argument as an additional ethical basis for supporting promotion of THR.

Defending maximization of health outcomes as a valid ethical goal is a sketchy proposition at best. It is far weaker still when only physical health is considered, as is usually the case in discussions of smoking and nicotine use, in an area where major psychological health costs and benefits are at stake. Maximizing welfare, rather than just one of the many contributors to welfare (health), is a much more sensible description of human behavior (thus, the construction of the concept of utility or welfare) and is accepted in policy decision-making methods (e.g., cost-benefit analysis). The welfare objective function is necessarily more difficult to quantify than any single component of it, which presents a problem for some analyses. But in the present case the welfare argument is actually easier to make, since it can be based on the observation that free choices must be welfare improving, while the longevity-based argument requires further calculations.

The apparent reason for what is probably best described as the "maximize longevity" objective function is the context in which discussion of THR usually takes place. Assessing which policies or practices improve longevity is a perfectly valid positive analysis in public health science. But a political faction within public health, often known as "health promotion", tends to construe this as the worldly goal. The social goal of maximizing population longevity, what might be called the "health promotion pseudo-ethic", is not a defensible ethical position based on either revealed preference/belief or ethical theory. Every individual makes many choices that lower their life expectancy in pursuit of other components of welfare, and so people do not accept this objective; no

serious analysis of ethics proposes that maximizing longevity is an appropriate social ethical goal; and serious attempts at normative analysis in health (such as cost-effectiveness analyses in medicine) consider resource costs (i.e., how much we are willing to sacrifice other human wants) and quality of life, not just longevity. But there is a faction in public health that tends toward believing this pseudo-ethic and the closely associated "we will choose what is best for you" political ethic. (The implications of those alternative ethics themselves are addressed in the subsequent paper.) The public health movement evolved partially from fringe groups who were trying to purify people's bodies and souls, and for most of human history public health dealt with major threats where there was little discord between health and welfare. As a result, some residue of this attitude remained even as public health became more scientific and mainstream, and (in the richest countries) moved beyond major health issues like basic nutrition and sanitation to focus on behaviors and other relatively minor risks. In particular, the goal of maximizing physical health, even to the point of inflicting great damage to liberty or welfare, is often implicitly invoked as if it were generally accepted.

Fortunately, for present purposes, it is not necessary to argue the merits of this goal compared to the above ethical arguments, since the conclusions based on either are the same. The "maximize health" objective function results in clear affirmative argument for pro-THR policies, and indeed, is the only basis for supporting the strongest possible pro-THR policies

(Readers unfamiliar with objective functions or consequentialist analysis may find the following useful in understanding the above points: It is only possible to maximize one objective at a time, so if longevity is being maximized then other wants (physical comfort and pleasure, psychological functionality, companionship, creative accomplishment, *joie de vivre*) are necessarily not maximized. Saying "we should maximize health and happiness" is nonsense, since it is impossible to maximize both unless they are perfectly collinear (i.e., equivalent). The concept of "welfare" or "utility" is a scalar function that combines the various factors that contribute to individual well being. Strictly speaking, maximizing a particular objective means that we are willing to give up any amount of anything else to increase that objective a small amount – e.g., if the objective is really to maximize longevity, we would be willing to make someone utterly miserable so that he could live a day longer, which clearly no one would consider an improvement. More practically we can recognize that all but the most fanatical and out-of-touch health promotionistas would not demand unlimited loss of other benefits to prevent physical disease, but instead are demanding a highly disproportionate weighting of physical health in the objective function. But since promoting THR increases both longevity and welfare, any convex combination of the two is also increased.)

Arguments about improving health outcomes can be divided between the analysis of an individual smoker and analysis of population aggregate outcomes. The first of these is simple: As posited, use of a THR product is far less harmful than smoking, so a smoker who can be persuaded to switch has much better expected health outcomes. This observation supports each of the above-listed policy options more strongly than it supports the one below it, and thus policy 1 tends to be supported: If smoking is not

going to simply be eliminated by fiat, then the goal of maximizing health justifies informing smokers about the benefits of switching and justifies maximally aggressive pushing of smokers to adopt THR.

A common challenge to this conclusion is that some smokers who are destined to quit nicotine entirely might switch instead, and thus their total risk would increase. This claim is sensitive to the quantifications in the empirical assumptions. If switching to a THR product substantially displaced abstinence and was far less risk-reducing than becoming abstinent, then it would be plausible that switching decreased longevity for the average smoker who chose to switch. (Welfare would still be increased if that smoker understood the risks and preferred the THR option.) But since, as posited above, the average smoker would have to quit within the next month or so to suffer less risk from his continued smoking than from a lifetime of using a THR product, at any given moment there are very few smokers who would suffer a net expected increase in physical health risk by taking up THR. Moreover, since such individuals are on the verge of quitting, and presumably are aware of that, they would have no incentive to switch merely to reduce their risks. Finally, switchers could still quit entirely if that was their preferred final outcome, and many presumably would. Even if the empirical quantifications are inaccurate by a factor of two or even ten, the number of smokers who suffer a loss in life expectancy due to THR is very small, as is the magnitude of their net loss, and so the net effect across all smokers is clearly positive.

Nevertheless, a few individual smokers will be made worse off by switching both for the reason noted in the previous paragraph and the following: Average lower risk does not mean every individual benefits. If smoking substantially hastens the death of 1 in 3 users, while THR products substantially hasten the death of 1/300th of users, then it might be that some users who would have survived smoking will succumb to the THR product. This is probably less than the 1/450 that would result if these probabilities were independent, since most people who were susceptible to the risk from the THR product would probably be even more so to smoking, but it is presumably nonzero. This is not a problem given a realistic interpretation of the objective function as being in expected value terms. Notwithstanding that many naïve commentators seem to think that there is an accepted health ethic of "do no harm" that prohibits individual acts that might, in retrospect, be responsible for a net increase in harm, no one seriously proposes that the goal of "maximize health" implies never doing anything that might possibly do more harm than good for a particular individual. Such an ethic would demand something close to inaction: no vaccinations, few medicines, no non-emergency surgery, and not much else beyond basic hygiene. Thus, the observation that a very few smokers might actually be healthier if denied THR does not diminish the affirmative argument, which like virtually everything else in health practice is based on the expected value of the action.

The same conclusion is reached if the analysis is done at the population level, taking into consideration nonsmokers whose risk could theoretically increase. In contrast with the welfare analysis, nonsmokers who learn about the low risk from THR products and choose to start using them count as negatives from the perspective of physical health risks. However, given the numbers posited in the empirical analysis, such negative

effects will clearly be overwhelmed by the positive results. In societies where smoking is popular and widely accepted, and there is little worry about the health impact, prevalence typically reaches about 60% of the adult population; we can use this as an upper-bound estimate for the portion of the population that likes using nicotine (i.e., gets net benefits considering everything except the health costs). Even if use of THR products increased nicotine use to this level, it would represent less than 40% of the population switching from abstinence to THR products. At a 99% risk reduction compared to smoking, only 1/2 of 1% of the population would need to switch from smoking to THR products to more than compensate for this maximum conceivable effect.

Far greater negative health effects would occur if a THR policy caused more people to smoke. In that case, the longevity tradeoff would not be 1-for-99, but 1-for-1. A common claim is that promoting THR could cause a "gateway effect", whereby people who would never have smoked start using a THR product and then switch to smoking. It turns out that such arguments that are made by THR opponents are actually that the existence of THR products, not the promotion of THR, might lead to more smoking via this path. Even if the evidence is interpreted as showing that THR product use causes some people to become smokers (as opposed to recognizing it merely shows that some people like nicotine more than others), it comes from communities where it is not widely known that THR products are low risk, and thus users believe that they might as well smoke. If the products were causing would-be nonsmokers to become smokers, it would be an argument for promoting THR, to inform those THR product users that switching is very bad for their health. It is difficult to identify any mechanism by which informing people about the low risk of THR products would encourage someone who would not have smoked to switch from those products to smoking.

There is a mechanism for THR promotion increasing smoking that, unlike the gateway story, does have some plausibility. There may exist would-be smokers who want to smoke for a limited period but are dissuaded from trying because they fear that they will find it too unpleasant to quit when the time comes to do so. By providing them with the promise of a better cessation method, they might be inclined to start smoking (see Bolton et al. (2006) for a discussion and examples of how "remedy messages" like these might increase consumers' intentions to engage in risky behavior). Of course, this concern applies equally or more strongly to other methods for helping smokers become abstinent, ranging from counseling to pharmaceuticals. (This presumably explains why this concern is basically never mentioned by anti-THR activists even though it has far more *prima facie* legitimacy than many of their claims: Those activists tend to be politically or financially committed to other smoking cessation methods and have a history of avoiding admitting to even the most obvious negative consequences of those methods.) Though there are almost certainly some such would-be smokers who start only because there are good cessation methods, it seems impossible that the magnitude of the longevity effect via this pathway would be significant: By hypothesis, affected individuals are radically different from the typical first-time smoker; they are highly motivated by the goal of quitting smoking at a particular time, presumably at a young age before smoking has caused a life-threatening disease, and so are particularly likely to do so, and at the same time anticipate having difficulty quitting. Presumably any such individuals who would

be tipped into smoking by THR is already influenced by other cessation methods being (erroneously) touted as highly effective. Moreover, most young smokers report expecting to quit before it damages their health even without knowing about THR, so it is not clear that many others would be motivated by the prospect of quitting being easier. Plus, this scenario is limited to people making a carful health choice – those who are motivated to avoid being "nicotine addicts" might be tempted by a new pill that facilitates abstinence, but would not consider THR the way out they are hoping for.

That said, the observation about young smokers who are sure they are going to quit (even though many do not) creates some reason for concern that those who are tempted into smoking by an easier opportunity to quit might still not quit, and we cannot be certain that the number who would be tempted is really trivial. Thus, it would be reassuring to have some empirical evidence about the extent to which better cessation methods cause additional persistent smoking. But lacking evidence that the effect is substantial, it seems to put only a modest dent in the social longevity benefits from promoting THR. As with most specific points in this section, those who chose to start and continue smoking would be a result of enhancing free choice and changes of preference, and thus appear to have only positive effects on welfare.

Thus, given the posited health benefits of switching, the portion of smokers that would find THR products acceptable, and very limited negative health consequences, it is clear that policies promoting THR would tend to increase average population longevity. The more aggressive the policy (assuming more aggressive does not diminish effectiveness) the better, and thus each of the above policies is more strongly supported than the ones below it.

What if the empirical assumptions are incorrect? As noted, the "maximize longevity" arguments are somewhat sensitive to empirical quantifications. If the estimate of risk reduction were far too optimistic – say it is really a reduction of only 50% rather than 99% – and the health-cost sensitivity is far greater than expected (which is to say, more people would take up nicotine if the risk was lowered) – such that products with 50% risk reduction attracted as many nonsmokers as smokers – then there would be no net gain. It does not appear that anyone has ever proposed numbers as close to large enough for this to occur.

What about the claim that some THR products (though not promotion of THR) are a gateway to smoking? If this were true, it could arguably support a prohibition of THR products (perhaps creating a system of prescriptions available to smokers and former smokers), but it is difficult to see it justifying misleading people (thus, policy 6). However, this is only justified if the costs from THR products causing smoking are more than the net benefits from the other effects of promoting THR. While this is certainly possible, it does not appear that anyone has ever claimed it is the case – at least not with any attempt to legitimize the claim with quantification. Indeed, it appears that the best solution to this problem would be to aggressively provide accurate information (not in support of THR necessarily, but basically the same information): If people who never would have started smoking, and thus seem to have some disinclination for it, find

themselves drifting from their adopted THR-product habit to smoking, they presumably would be discouraged by learning that switching to smoking would move them from minimal risk to high risk. It is the misperception that THR products are high risk that makes users believe that they might as well smoke. Thus, promoting THR seems like it would reduce, rather than increase, any gateway effect from existing legal products.

In general, the arguments about THR policies having a net negative health effect seem to demand one extreme position or another. If it really were the case that too many new consumers might take up nicotine and the risk reduction was not great enough, or if THR products really did cause more smoking, then the optimal "maximize longevity" policy might be to ban THR products, even if smoking were to remain legal. However, this would have obvious health costs in the form of current THR product users switching to smoking (some might quit or avail themselves of the black market, but smoking would be easier; nonsmoking THR product users often smoke when they find themselves with access to cigarettes but not their preferred products), which are likely far greater than the benefits. Perhaps THR products could be made available only to people who can prove they already use those products or smoke, though the practical implications of this start to border on the absurd.

Indeed, the inherent oddity of an ethical argument that we should set THR policy, in isolation, to maximize longevity makes absurd conclusions inevitable if we push too hard on it. Even setting aside the question of why smoking is not banned if society really accepted that the objective of health-related policy is just to maximize longevity (or for that matter, why we have not banned soda, coffee, contact sports, McDonald's, recreational transport, etc.), it is difficult to imagine creating a large minority of card-carrying users who are pushed hard to use THR products rather than smoke, while we try to enforce prohibition on everyone else. Moreover, if there are no other ethical constraints, pushing on the smokers should include telling them whatever it takes to get them to quit smoking (e.g., all smoking cessation methods are perfectly harmless and will also make them more attractive and improve their sexual performance), since honesty and freedom are not part of the goal. Meanwhile nonusers should be told that THR products are pure poison. Strangely, this hypothetical muddle actually bears a remarkable resemblance to current policies in many Western countries, though further analysis of that point is beyond the present scope. Fortunately, this does not need to be unpacked because so long as the empirical assumptions are correct, the health promotion pseudo-ethic clearly supports strong THR policies, and it is possible to reach a sensible conclusion without exploring the weaknesses of that objective function as a normative goal.

Discussion

This analysis presents the ethical basis and logic behind the established affirmative arguments supporting THR policies, as well as delineating possible policies. Such an exercise is necessarily imbalanced, not addressing interpretations of the empirical evidence that differ dramatically from those posited, and excluding ethical goals that are not captured in one of the three ethical bases presented. Though imbalanced, this presentation was intended to be unbiased given the starting points. That is, though this

presents only the affirmative case, it is still possible to address both the strengths of the arguments and bases for disagreement with them. It might appear that there was limited attempt to provide the latter, but it is actually remarkably difficult to find any legitimate arguments that might be made. Thus, in terms of the resulting message, this attempt to analyze the arguments is admittedly difficult to distinguish from simply arguing the case for THR. Without making absurd departures from existing empirical estimates, it is not possible that these can shift the balance of any of the arguments. This is not a claim that there are no compelling arguments against promoting THR, but it does mean that those would have to be orthogonal to arguments presented here – that is, they have to be based on some goals or duties other than informed autonomy, improving welfare, or improving physical health.

> [Note that at this stage in the life of this paper, the author is soliciting any comments. These could include, perhaps most importantly, any better presentations of the arguments against or entries in the ledger that disfavor the affirmative pro-THR arguments presented. Any such arguments that are considered compelling will be added.]

It should be noted that any of the three arguments – if its ethical premises are accepted – is sufficient to make the case for (particular) policies promoting THR. It is not necessary to make all three arguments (though treating the "maximize health" argument as the only argument, as is often done, tends to severely understate the case). The right to informed autonomy alone is a sufficient argument, if that right is accepted. The policy goal of pursuing available welfare improvements is also sufficient. Were it the case that rights and welfare conflicted or that some individuals suffered welfare losses even as others gained, the analysis would be more difficult. But informed autonomy is effectively equivalent to Pareto optimization in a case like this, eliminating this potential complication.

Similarly, if the often invoked health promotion goal led to the conclusion that promoting THR was bad, it would be necessary to argue the merits of welfare or autonomy over forcing people to maximize their longevity. However, since the results point in the same direction (modulo exactly how strong a policy is warranted), no such resolution is necessary. This also means any consequentialist objective function that is based on welfare but that overweights the contribution of longevity to the total will be improved by promoting THR, though the question of which policy it best supports might require more detailed analysis. The quantitative empirical claims necessary to suggest a net negative health impact are such a departure from the best estimates, and it appears that no one with any scientific credibility (even the most vocal opponents of THR) tries to make such quantitative claims.

It is interesting to briefly consider how the above analysis compares to other harm reduction discourse. There are cases, like the use of seatbelts in cars, where there seem to be no serious liberty or competing benefits issues (that is, any benefits of driving beltless are quite trivial and exercising one's ability do so is generally seen as a mere act of unseen rebellion for its own sake). Thus, the only ethical challenges to mandatory

seatbelt laws seem to come from the perspective of extreme liberalism, the right to do something that hurts no one else, even though no one can identify a reason for doing it. By contrast, motorcycling without a helmet is regarded by some as being a dramatically improved experience and an open self-defining statement, and thus mandatory helmet laws are actively (though often not successfully) resisted on the grounds of welfare and individual rights. For transport harm reduction, there does there seem to be any substantial effort by government or its allies to mislead people about the true risks even though there are political fights in this arena. The use of condoms as a harm reduction strategy is supported by arguments very similar to those for THR. Indeed, parallels include the health promotionist faction advising adults to favor sexual abstinence over the low-risk option. However, in contrast with THR, they recognize that people will choose the high risk alternative if not given a low risk alternative, and so it is accepted that even the health promotion pseudo-ethic favors harm reduction over abstinence-only approaches. There is no serious effort to prevent access to condoms in liberal societies, and despite the hype about sexually transmitted diseases, there very few fringe actors try to tell people that sex with a condom is just as risky as without. In contrast with the THR case, the demands for abstinence and the rare cases of product bans or grossly misleading health claims are almost universally recognized as absurd and unethical.

The analysis of the arguments for promoting THR emphasizes the fact that the most commonly implied argument – maximizing population average health outcomes – is the ethically weakest and the most technically fiddly affirmative claim. The implications of this are interesting in themselves. Policy ethics questions about a popular, socially embedded, easily understood, potentially high-risk behaviors like nicotine use should be a matter of community engagement, much more so than debates about highly technical questions about minor risks (like environmental pollutants, where stakeholder engagement is accepted as mandatory, as it is in many complicated non-health areas of techno-science policy). Yet in the case of tobacco and nicotine policy, technical analysts and political activists, actors who have no more claim on assessing policies or ethics than other members of society, still claim a near-monopoly over the normative discourse. Technocracies have a habit of mistaking that which they normally measure for what are (or ought to be) society's goals. The only substantial competition for this perspective comes from those who are concerned with purity or abolishing substance use, regardless of health consequences (or welfare, or rights). The people who are by far the most important stakeholders -- nicotine users -- have almost no voice.

Occasionally arguments about users' rights are made, though they are typically dismissed without serious response from THR opponents. Discussions of users' welfare are almost unheard of beyond individual testimonials, and those testifying individuals do not try to push their experience into the political discourse, perhaps because so many are convinced that their experiences do not matter and they should loathe both their consumption choice and themselves, just as the most vocal faction does. Normative discussions about THR thus manage to not only subordinate the core ethics of modern health policy, but also ignore the welfare and opinions of the very people that tobacco policies are ostensibly supposed to help. The present analysis -- by someone who is merely an occasional, highly-non-typical nicotine user -- attempts to incorporate the expressed views of the

major stakeholders, and makes the case for considering their usually ignored benefits, net welfare, and rights, but it cannot fully substitute for the direct voices of the stakeholders.

Competing interests

As should be clear from the overall goal of this paper, explicit statements in the text, and the content of the self-references, places where such information is far more useful than it is buried here, the author is an active proponent of promoting both THR and better ethical reasoning in public health. He is lead author of the cited website that attempts to consolidate arguments in favor of THR, and lead author on some of the specific analyses that are antecedent to the present one. This history of presenting particular arguments might affect his incentives to continue to reinforce those particular arguments. This history might also cause him to unconsciously understate the weaknesses of the arguments despite his best efforts to present the most complete analysis of this topic he could. However, it is not obvious why the author's mere support for THR would affect how he presented the arguments in favor of THR, since the goal was explicitly to present the best possible arguments. The author's personal socio-politics favor liberalism, and he generally tends toward the *de gustibus non est desputatum* principle regarding individual welfare, which might bias him in favor of the first and second arguments compared to the third to a degree that exceeds what the analysis supports, though the presented arguments are intended to capture the basis for these preferences. This project was author-initiated and is not specifically tied to any funding; the author currently receives funding from organizations with clear interests in these matters, including U.S. Smokeless Tobacco Company (an unrestricted grant from USSTC to the University of Alberta has helped support this research) and British American Tobacco, was supported by the Robert Wood Johnson Foundation in the past, and hopes to receive funding from any willing organization interested in THR in the future. No funder played any role in initiating or authoring this analysis. As of the release of the current version, the author will solicit comments on this analysis from many individuals, including some who are employed by organizations that have a political interest on one side or the other of this issue.

References

Balfour, D. (2009). Letter to Samerasekera. Available at: http://www.tobaccoharmreduction.org/wpapers/009.htm

Ballin, S. (2007). "Smokefree" tobacco and nicotine products: Reducing the risks of tobacco related disease. (Available at: http://www.tobaccoatacrossroads.com/2007report/071128_Ballin%20Report_final.pdf).

Bates, C., Fagerström K., Jarvis, MJ., Kunze, M., McNeill, A., & Ramström L. (2003). European Union policy on smokeless tobacco: A statement in favor of evidence based regulation for public health. *Tobacco Control, 12*, 360-367.

Beauchamp, T. L., & Childress, J. F. (2009). *Principles of Bioethics*. 6th Edition. New York: Oxford University Press.

Bolton, L. E., Cohen, J. B., & Bloom, P. N. (2006). Does Marketing Products as Remedies Create "Get Out of Jail Free Cards"? *Journal of Consumer Research*, *33* (1), 71-81.

Libin, K. (2007, September 28). Whither the campus radical? *National Post*, p. A1. Available at: http://www.tobaccoharmreduction.org/pol/libin.pdf

Phillips, C. V. (2004). Smokeless tobacco and cigarettes: Gateways, causal pathways and harm reduction. Presentation at SRNT, 20 February 2004. (Available at: http://www.epiphi.com/papers/phillips_gatewaypathways_srnt-talk.pdf).

Phillips, C. V. (2007). Commentary: Warning: Anti-tobacco activism may be hazardous to epidemiologic science. *Epidemiologic Perspectives & Innovations, 4*(13).

Phillips, C. V. et al. (2009). TobaccoHarmReduction.org.

Phillips, C. V . Debunking the claim that abstinence is usually healthier for smokers than switching to a low-risk alternative, and other observations about anti-tobacco-harm-reduction arguments. (Working paper; available at: http://www.tobaccoharmreduction.org/wpapers/007v2.pdf)

Phillips, C. V., Bergen, P., & Nissen, C. M. (2009). A dissection of what passes for scientific review and policy analysis among anti-Tobacco-Harm-Reduction researchers; An annotation of "Is Smokeless Tobacco Use an Appropriate public Health Strategy for Reducing Societal Harm form Cigarette Smoking?" by Scott L. Tomar, Brion J. Fox and Herbert H. Severson from International Journal of Environmental Research and Public Health 2009, 6:10-24. Available at: www.tobaccoharmreduction.org/wpapers/010.htm

Phillips, C. V. (2008). Commentary: Lack of scientific influences on epidemiology. *International Journal of Epidemiology, 37*, 59-64.

Phillips, C. V., Heavner, K., Bergen, P. (in press). Tobacco – the greatest untapped potential for harm reduction. Book chapter.

Phillips, C. V., Rabiu, D., & Rodu, B. (2006). Calculating the comparative mortality risk from smokeless tobacco versus smoking. *American Journal of Epidemiology, 163*(11), S189. Poster Presentation, June 2006, Congress of Epidemiology conference. Available at: http://www.tobaccoharmreduction.org/papers/phillips-comparativerisk-poster-jun06.ppt

Rodu, B. (1995). For Smokers Only: How Smokeless Tobacco Can Save Your Life. Sulzburger & Graham Publishing Ltd.

Rodu, B., & Godshall, W. T. (2006). Tobacco harm reduction: an alternative cessation strategy for inveterate smokers. *Harm Reduction Journal, 3*(37).

Royal College of Physicians. (2007). Harm reduction in nicotine addiction: helping people who can't quit. A report by the Tobacco Advisory Group of the Royal College of Physicians. London: RCP. (Available at: http://www.rcplondon.ac.uk/pubs/contents/bbc2aedc-87f7-4117-9adad7cdb21d9291.pdf)

Thaler, R. H., & Sunstein, C. R. (2008). Nudge – Improving Decisions About Health, Wealth, and Happiness. New Haven & Longon: Yale University Press.

Timberlake, D. S., Huh, J., & Lakon, C. M. (2009). Use of propensity score matching in evaluating smokeless tobacco as a gateway to smoking. *Nicotine & Tobacco Research, 11*(4), 455-62.

Tomar, S. L., Fox, B. J., & Severson, H. H. (2009). Is Smokeless Tobacco Use an Appropriate Public Health Strategy for Reducing Societal Harm from Cigarette Smoking? *International Journal of Environmental Research and Public Health, 6*(1), 10-24.

Zeller, M., Hatsukami, D., & The Strategic Dialogue on Tobacco Harm Reduction Group. (2009). The strategic dialogue on tobacco harm reduction: A vision and blueprint for action in the United States. *Tobacco Control, 18*(4), 324-332.