

IEA Discussion Paper No.72

UNDERSTANDING THE BASIC ECONOMICS OF TOBACCO HARM REDUCTION

Carl V. Phillips
August 2016

With some exceptions, such as with the publication of lectures, IEA Discussion Papers are blind peer-reviewed by at least one academic or researcher who is an expert in the field. As with all IEA publications, the views expressed in IEA Discussion Papers are those of the author and not those of the Institute (which has no corporate view), its managing trustees, Academic Advisory Council or senior staff.

Contents

About the author	04
Summary	06
Introduction	08
What we can learn from economics	11
Preference, purchase price, and health: some basic implications	14
A rational choice?	22
Optimal rates of taxation	29
Discussion	40
Acknowledgements and Competing Interests	44
References	45

Carl V. Phillips, MPP PhD, is a consultant and independent researcher in epidemiology and economics. He is a former professor of public health and medicine (University of Alberta, University of Texas), with training in economics (PhD and MPP in public policy, Harvard), health policy research (fellowship, University of Michigan), philosophy of science (fellowship, University of Minnesota), mathematics and history (BA, Ohio State University).

Phillips was a pioneer in the field of tobacco harm reduction, and has written extensively about it from the perspectives of economics, epidemiology, political philosophy, ethics, and scientific epistemology. He has been studying and documenting the poor quality science that dominate tobacco control for fifteen years. His work in that field includes health outcomes epidemiology and modelling, behavioural and economic surveys and modelling, and the only evidence-based calculations of the comparative risk of different tobacco products. His award-winning work in epidemiology includes identifying core failures in the field and suggesting improved methodologies.

Summary

- The branch of economics that studies preferences and choices as they relate to costs and benefits is the natural starting point – indeed the only apparent candidate for a rational starting point – for understanding these phenomena as they relate to tobacco/nicotine use.
- Many implications of substituting low-risk alternatives (e-cigarettes, smokeless tobacco) for smoking, usually referred to as ‘tobacco harm reduction’, can be derived from even simple economic analysis. But such analysis has been largely avoided, limiting understanding and creating potentially harmful myths. Illiberal policies are hidden behind indefensible (and never actually defended) assumptions that contradict basic economics.
- Substitution of a low risk product would be welfare-enhancing for most smokers. Some smokers will still prefer smoking to any available alternative, despite the much higher risk. But there are vanishingly few smokers for whom abstinence is a better choice than switching to a low-risk alternative. Thus there is no apparent ethical justification for anti-smoking measures that push for abstinence rather than switching.
- The availability of low-risk tobacco/nicotine products will inevitably increase total consumption as compared to a world where cigarettes are the only option. This is the inevitable and rational effect of lowering the costs of a consumption choice. It is properly counted as an additional benefit, though it is widely derided as a cost. Public supporters of low-risk products who condition their support on those products not attracting any new users are either being naïve or cynically imposing conditions they know cannot be met.

- For any remotely defensible goal, including minimising population health risk, the optimal level of excise tax on low-risk products is zero (assuming that is the lower bound; a subsidy would be better still). This is sometimes presented as if it were immediately evident from the comparative risk, but that is not actually a valid claim. However, simple economic analysis shows that it is the case.
- Many of the proposed bases for denying the economic model are vacuous upon closer examination. Those that are valid, such as a history of tobacco use creating a desire for continuing use, can be incorporated into the model in a straightforward way. The concept of 'addiction', in particular, is best understood in the context of the model and its implications do not alter the economic analysis.
- While the basic economic model is clearly not perfect, it is close enough to be informative and is the only apparent candidate as an analytic foundation for further analysis. There is no apparent defensible theory of choice and benefits that provides a viable alternative to standard economic thinking.

Introduction

Discussion and debate about tobacco policy, particularly tobacco harm reduction (THR), has increased dramatically with the increasing popularity of e-cigarettes. Harm reduction is a concept that encompasses consumer empowerment and freedom and minimising caused-harm (e.g., eliminating criminal penalties), but particularly focuses on offering or encouraging lower-risk alternatives. The major success story for THR is the substitution of snus (smokeless tobacco) for most would-be smoking among Swedish men, and to a lesser extent Swedish women and Norwegians. More recently, the explosion of popularity of e-cigarettes revitalised THR, after repeated failures to export the Swedish success with snus. While smokeless tobacco is currently more popular, e-cigarettes are the rising star that is proving attractive to many smokers who never previously considered THR. The pharmaceutical industry's nicotine replacement therapy (NRT) is also a potentially viable THR product, but has such low quality in the eyes of most consumers that few are interested in long-term use. Thus most of the analysis here is probably best thought of as being about e-cigarettes, though it applies equally to smokeless tobacco or NRT in any population where those might be considered attractive alternatives.

Much of the THR discussion – including by THR proponents – is mired in economic myths about tobacco/nicotine¹ consumption. In particular, there is the implicit notion that consumption of these products is somehow so unlike other consumer choices that everything we know about welfare

1 The construction 'tobacco/nicotine products' is used, with apologies, to collect consumer products that contain tobacco leaf (cigarettes, smokeless tobacco) along with products that put the nicotine extracted from tobacco in another substrate (e-cigarettes, NRT). Because there is oddly heated debate about whether the latter category should be included in the simpler phrase, 'tobacco products', this awkward terminology has developed to bypass that argument.

economics² should be ignored, in favour of ad hoc stories about what people do, why they do it, and how they will react to incentives. Those ad hoc departures from economic principles are insidious. They are a case of the Sapir-Whorf hypothesis about language shaping cognition, which when engineered intentionally is better known as Orwell's concept of thought control via control of language: the anti-economic assertions are embedded in the language surrounding the discussion, forestalling useful analysis.

There is remarkably little understanding of what welfare economics tells us about the consumption of tobacco/nicotine products. When consumer choices are attributed to 'marketing' or volition-free behaviour, it is difficult to rationally analyse preferences and their implications, let alone benefits and net welfare effects. The dominant (implicit) model of tobacco/nicotine consumption is effectively *demonic possession*: People act in ways that have no rational explanation, because some arbitrary force is controlling them. Once a demonic possession theory is built into the bedrock of the discussion, it is possible to justify any policy to exorcise the demons (how could anyone object to that?). Worse, it is possible to make up stories about the effects of policies, and to deny people the opportunity to make free choices on their own behalf.

The basic welfare economics model, as a framework for understanding the world and an approach to research, is extremely flexible. If someone believes that 'addiction' or 'marketing' are complications that must be considered, they can incorporate them into the model. Indeed, having a model forces proponents of such claims to explain what impacts they have, rather than just invoking them as vague evil demons. All models are imperfect, of course, but they can usually be made as good as is needed, and they are clearly superior to presenting no model at all (which usually means having a model, but avoiding articulating it because it is indefensible). There is simply no excuse for not using the science of welfare economics, which has served us so well, as the starting point for analysing these particular consumer choices.

Even a basic understanding of economics informs several common questions which seem to baffle most commentators. A deeper exploration can help answer more complicated questions that have been largely

2 Welfare economics, which can also be called consumer economics, is the study of costs and benefits of consumption decisions, and thus also of the resulting impact on welfare (or to use the economics jargon, on 'utility') at the individual level, as well as the associated incentives and decisions.

ignored because it was not possible to conceptualise them without a useful model of human behaviour. The framework in this paper will help readers with an interest in THR gain a general understanding of the concepts and lay the groundwork for future analyses, but it also produces a few important results immediately. It offers general lessons that apply to any attempt to hide illiberal policies behind hand-waving departures from basic economics. There is nothing novel or scientifically controversial about the analytics. There are a few departures from the simplest economic models, where the standard simplifications are not useful, but the basic points can be found in introductory textbooks. For those who may not be familiar with some of the underlying economics concepts, footnotes explain some terms and concepts and expand upon a few complicated points that might distract from the basic presentation.

What we can learn from economics

Three common questions can be addressed using the framework presented here, with many other possible extensions³:

1. What is the net welfare benefit to smokers who adopt THR (or would-be smokers who choose a low-risk product in the first place), and how does that welfare compare to choosing either smoking or abstinence?
2. What will happen to total consumption prevalence when low-risk alternatives to smoking become established? Will people who never smoked be inclined to initiate use of these products?
3. What is the optimal excise tax ('sin tax') structure on the various categories of products?

The first is immediately illustrated by the model, though it is not a product of the theory alone (in particular, it requires introducing what we know about the low risk of alternative products). The second is a simple result of the model. The third is addressed as an example of an extension of the basic model.

For practical purposes, the market for tobacco/nicotine products consists of two types of products: (1) Cigarettes, the currently dominant product, which substantially hasten the deaths of about a third of lifelong users, and also create substantial disability and non-fatal illnesses; (2) Low-risk smoke-free alternatives, including Western-style smokeless tobacco products, smoke-free inhaled products (e-cigarettes and a few emerging

3 E.g., for an extension to population dynamics, see <http://ep-ology.blogspot.com/search/label/THR%20modeling>

variations), and pharmaceutical-style products (NRT). There is a lot of noise about different risk levels among different low-risk product categories and within those categories (as well as comparisons to abstinence). But for practical purposes these are indistinguishable and the products can simply be thought of as having the same low risk, and so are just combined as 'low-risk products' in the present analysis. Unlike most areas of harm reduction, the risk from the low-risk alternatives is so low that substitution has basically the same health risk as abstinence.

The only low-risk product for which we have useful epidemiology, smokeless tobacco, causes only about 1/100th the disease risk from smoking, based on the only existing attempt to calculate an evidence-based estimate of overall comparative risks (Phillips et al. 2006), and calculations for specific diseases support that estimate (e.g., Lee and Hamling 2009). Estimates for other products must be based on what we know about smokeless tobacco, but since there are no reasons to believe they differ much, this is adequate. There is no affirmative evidence that the risks are different.⁴ This estimate of harms ignores the apparent health benefits of nicotine consumption (e.g., protection against neurodegenerative diseases), so it is plausible that the net effects are actually positive. Thus it cannot be claimed with confidence that use of THR products is less healthy than abstinence. To simplify the present analysis, it is assumed that the low-risk products pose a net health risk. Most of the conclusions remain true, or are even more so, if the net health effect is beneficial. It is worth noting that even if low-risk products pose some risk, a lifetime of use poses lower risk than from continuing to smoke for just a few more months and then becoming abstinent (Phillips 2009).

While there are tobacco products whose risks fall somewhere in between low-risk products and cigarettes, none currently play a major role in THR. Variations in risks among different varieties of cigarettes are undoubtedly substantially greater than the differences among low risk products and between those risks and zero.⁵ But since it is not entirely clear which varieties are lower risk, and the risks are sufficiently similar for present purposes, this is ignored. Thus there are three options: smoking, using a low-risk product, and abstinence.

4 A popular claim at the time of this writing is that e-cigarettes are merely 95% less harmful as smoking – i.e., pose five times the gross disease risk of smokeless tobacco – but this is not supported by any evidence.

5 For more on this, see: <http://antithrlies.com/2015/12/23/utter-innumeracy-six-impossible-claims-about-tobacco-most-public-health-people-believe-before-breakfast/>

(Combinations of products are also possible and quantities matter; these are set aside for the simple model.)

Cigarettes and low-risk products are substitutes, in both the common-language and economic senses of that word. Increases in the consumption of low-risk products come mostly at the expense of cigarettes, though the analysis shows this will not be the case entirely. Lower price or lower risk for one product causes more substitution by would-be users of the other product. Almost all users of one of these products are more inclined than nonusers to use another, though some would prefer abstinence to switching. Few consumers will find any two products to be perfect substitutes in terms of functionality, taste, cost, etc., but for simplicity, the base analysis assumes that every smoker considers some low-risk product a perfect substitute apart from differing health costs and purchase price. Recognition that this is not the case is then introduced as the analysis becomes more complicated, though it is kept informal to avoid a substantial increase in the mathematical complexity.

Preference, purchase price, and health: some basic implications

Some who are not familiar with basic economic models may believe that the mere presence of health costs renders consumer economics inapplicable. But even the simplest model allows for consumer costs to exceed the purchase price, as in Figure 1.

Figure 1

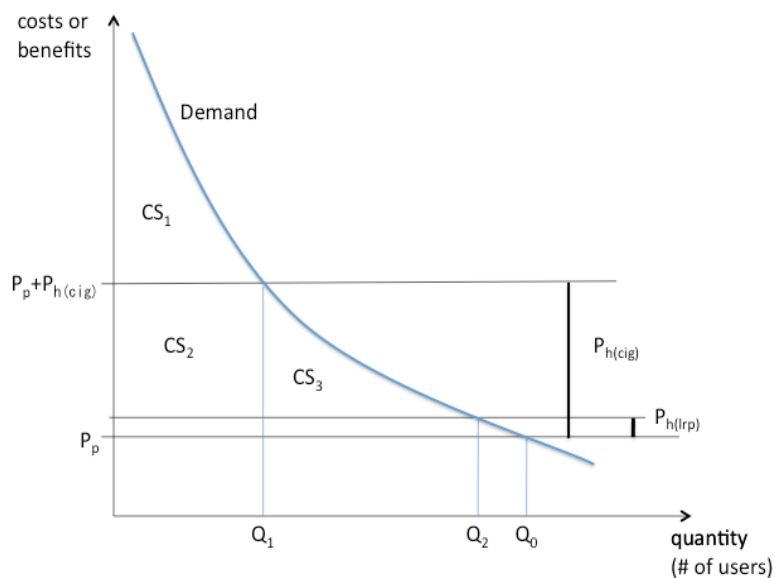


Figure 1 is a simple demand curve model. The horizontal axis is the number of people in the population who use the product⁶ and the vertical axis quantifies costs and benefits of consumption. The demand (D) curve relates the total cost to the number of people choosing to consume. Its vertical height is the maximum cost each would be willing to pay, with the downward slope constructed by ordering consumers by decreasing willingness-to-pay. Any point on the demand curve can be interpreted as either the gross benefit to one particular consumer from consuming the product (with the net benefit then being calculated by subtracting the cost) or, equivalently, as the cost level at which the consumer would choose to forgo the product. The shape illustrated for the demand curve is stylised, with the actual shape determined by real-world preferences. It is chosen based on the knowledge that a few people to the left get enormously more benefits from the product than others do, and there is a tail of people who get a little benefit but not much. There are also people further to the right (not illustrated) who get negative benefit – that is, they would not want to consume the products even if they were completely costless because they do not like the experience. Those to the left are unlikely to choose abstinence, even if the cost is driven very high, while those further right are likely to have already quit or never to have started using the products.

The horizontal lines represent the costs of consumption faced by each consumer. The quantity Q_1 , where the top horizontal line (the total cost of smoking, explained below) crosses the demand curve divides the population into those to the left who get net benefit from consumption, who would rationally choose to consume the product, and those to the right for whom cost exceeds benefit, who would rationally choose not to.

As with any analytic model there are built-in assumptions, several of which are designed to clarify the basic implications of the model (and could be relaxed in a more detailed analysis): Assume that the purchase price, P_p , is the same for all consumers in a population and all products. (To the extent prices vary across categories, that is equivalent to the effects of taxes, analysed below.) Assume that the health risk from a particular product is the same for every consumer. While some will eventually get the diseases caused by smoking and others will not, the *risk* can be

6 Most often this q-axis represents total quantity consumed, ignoring whether a few people consume a lot or everyone consumes a bit. But in this case it is more useful to set aside variations in individual consumption quantity and focus on the number of users.

reasonably estimated to be about the same.⁷ To allow this simplified two-dimensional analysis, it is also necessary to assume all consumers value health risks the same. This is clearly not true – different consumers value a particular risk of harm to health differently – though given that the low risk is approximately zero, that has little effect on any qualitative result. The setting aside of concepts like addiction, which some readers might already be objecting to, is *not* a simplifying assumption. Rather, as is addressed below, the concept actually does not require any departure from standard economic analysis.

Each individual makes a consumption decision based on how their total benefits compared to the total costs (free-market price, taxes, health costs), and total demand is determined by aggregating the individual choices. The purchase price, P_p , is represented by the bottom horizontal line in the graph. Since we are focused entirely on consumers, there is no reason to separate out the free-market price (the revenue collected by the sellers) from the taxes (which are a large majority of the purchase price in some jurisdictions), and so only the sum of the two costs is considered.

Quantity Q_0 is the number of people who would smoke if there were no health risk. For concreteness, if the purchase price were the free-market price (i.e., the high ‘sin taxes’ did not exist), Q_0 would probably be about half the population. This is based on the observed rate of smoking in populations where it was socially acceptable, consumers had enough wealth to afford it, products were not aggressively taxed, and there was not widespread concern about the health risks (e.g., men in Western countries c.1960). When we add the health costs from cigarettes, $P_{h(cig)}$, on top of the purchase price, the portion choosing to consume drops to Q_1 . That is, when this level of risk becomes known and is incorporated into consumers’ assessment of the total costs, consumers are responding to a cost of $P_p + P_{h(cig)}$ rather than just P_p , so a shift from Q_0 to Q_1 occurs. Again, for concreteness (though nothing in the present analysis is anchored to these empirical observations), this can be thought of as about 25% of the population, a typical smoking prevalence in populations where there is common knowledge of the health effects, but taxes and social opposition are limited and smoke-free alternatives are not popular.

7 A few consumers will know they are at dramatically higher risk of a particular disease (e.g., they have existing breathing problems). In that case, their individual calculus will be quite different, but this is beyond the current scope.

The net benefit for a consumer who chooses to smoke – anyone to the left of Q_1 – is the vertical distance between curve D (her total gross benefit) and the $P_p + P_{h(cig)}$ line (her total cost). This is referred to as *consumer surplus*, and is the net goodness in the world created by her having that option available. The total social welfare benefit created by the product is the sum of each consumer's surplus, the area CS_1 . Notice that this simple economic observation demonstrates the obvious inaccuracy of a claim implicit in much discourse about tobacco product use, that there is no benefit. Of course there is benefit. People make consumption choices only because they provide benefit.

The huge welfare benefits from low-risk alternatives

Now consider low-risk alternatives to cigarettes. Recall the assumption that the purchase price is the same and also assume that every consumer considers some low risk product to be as appealing as cigarettes, differing only in its health effects. These assumptions result in the same demand curve for the low-risk product as the original one for cigarettes.⁸ The health cost of these products is merely $P_{h(lrp)}$. (This is illustrated with a vertical distance that is much greater than the true risk. It is not possible to illustrate the comparative risk of low-risk products in a reasonably sized graphic and be able to distinguish it from zero, which may contribute to the failure to understand just how close to zero the risks are.)

The increase in consumer surplus offered by switching products is quite large. The potential consumer surplus gain from all smokers switching to the (hypothetically perfect) substitute is CS_2 , which is the vertical difference $P_{h(cig)} - P_{h(lrp)}$ – the health benefits of switching – summed across all smokers. In reality, of course, many smokers will find the alternatives produce lower benefits (and some will find the opposite), decreasing (or increasing) the net consumer surplus from switching. But under all the simplifying assumptions, the welfare gain is equal to the health benefits. This point is often overlooked in public health discussions, which focus on improved health or lower risk as an end in itself. The reason lowering risk has value is not some existential goodness, but because people value their health and so their overall welfare is increased.

It is seldom disputed that most smokers would be better off using a low-risk product than smoking. But looking at the actual economics tells us something more: Smokers are generally better off using the low-risk product rather than being abstinent. They preferred smoking to abstinence, generating consumer surplus CS_1 , so the low-risk option is clearly much better than abstinence (total surplus of $CS_1 + CS_2$).

Contrary to the simplifying assumptions, there are smokers who like smoking (in particular) so much that its net benefits are higher than those for any low-risk product, despite the risk difference. But there will be

8 Readers thinking one step ahead will realise that the demand curve for cigarettes will not be the same in a world where low-risk alternatives are available. Thus the word 'original'. The role of emerging makes the standard simple economic analysis, in which everything is currently close to equilibrium and the shocks are relatively minor, inadequate. This comes into play in the discussion of price elasticity, and requires a departure from what can easily be illustrated in a figure.

vanishingly few smokers for whom abstinence is the welfare-maximising option. For this to be the case, they would have to (a) change their preferences such that smoking has negative net benefits, which certainly does happen (people change their preferences and quit smoking). But they would also have to (b) derive so little benefit from a low-risk substitute (either because they really dislike them or completely lost their taste for tobacco/nicotine) that it is not even worth its much lower cost. Thus in a world with low-risk alternatives, it is difficult to see any ethical justification for trying to push smokers to become abstinent rather than switching. It is impossible to reconcile anti-smoking efforts that do not encourage switching with the pretence that anti-smoking policies are intended to make smokers better off.

Another conclusion that is immediately obvious from this simple model is that there will be an upward shift in the equilibrium quantity as cigarettes are replaced with low-risk alternatives. With P_h reduced from $P_{h(cig)}$ to $P_{h(lrp)}$, the new demand increases from Q_1 to Q_2 . Anyone to the left of Q_2 is better off using the low-risk product, rather than being abstinent, even though many of them were better off being abstinent rather than smoking. Indeed, since most of the cost of smoking is the health costs (with the possible exception of jurisdictions where punitive taxes are extremely high and no black market option is available) Q_2 will be quite close to Q_0 , probably even closer than is implied in Figure 1 (recall that the height of $P_{h(lrp)}$ is exaggerated). It could be lower than illustrated where strong anti-tobacco/nicotine social pressure has changed the demand curve so that it drops to zero much sooner than pictured in Figure 1.

Predicting the magnitude of the increase is much more difficult than predicting the direction, of course, and even more difficult is predicting the dynamics over time. Few never-smoking adults will adopt e-cigarettes even if they might benefit from them, since the physical behaviour itself is a barrier. More might adopt snus if it became popular in their culture. More ex-smokers who remember smoking fondly are likely to consider adopting e-cigarettes. The eventual equilibrium will be determined not by switching or new adult adopters, but by the behaviour of new cohorts coming of age. There is much consternation about the popularity of e-cigarettes among teenagers and young adults who have never smoked, but this is exactly what we would expect from rational decision-making. Given the low risk, a young person's decision about whether to seek the benefits of nicotine is really quite similar to the decision about drinking coffee (and almost certainly less consequential than the decision to drink

alcohol). Indeed, given the magnitudes that are (very roughly) predicted by this model, there is surprisingly *little* interest so far.

It may be difficult for many readers to imagine the direction of this effect was ever in doubt. When the quality of a product category increases (reducing a huge health risk is obviously a major quality improvement), more people will choose to consume it. The question ‘will more people use tobacco/nicotine products when they are low risk?’ is exactly equivalent to ‘do some people choose not to smoke because it poses high risk?’, and no one doubts that the answer to the latter is ‘yes’.

Confusion on this point has serious consequences. Some ostensible support for THR and even some regulatory rules are qualified with a demand that introducing or promoting low-risk alternatives must not increase total usage prevalence. This is an impossible condition to meet. (Of course, this might not be an error but rather a calculated tactic by those trying to discourage THR by tricking others into believing this condition is not tantamount to a ban.⁹)

To put this point in more technical terms, CS_3 is the additional consumer surplus that accrues to people who rationally chose to not smoke because the total costs exceeded the benefits, but would rationally choose to use low-risk products. Note the immediate implication of this: This is a *benefit*, not a cost. Nonsmokers who choose to use a low-risk alternative should be counted on the positive side of the ledger – representing additional benefit of these products being available – not as collateral damage as the rhetoric (on both sides) usually implies.

Many commentators are (inappropriately) focused on only the population health effects, as if minimising all risks is the only thing people care about. Thus they consider rational choice to use low-risk products to be a negative. However, there is also good news for them: The size of the new health cost rectangle, the area of the rectangle $P_{h(lr)}$ by Q_2 is much less than the harm from cigarettes, $P_{h(cig)}$ by Q_1 ; it is wider by perhaps as much as double, but with such a small height that it really cannot even be drawn on the graph. Even if many smokers stick with smoking, the net health cost reduction is huge.

9 For why anyone would ever want to discourage THR, see <http://antithrlies.com/2015/07/21/why-is-there-anti-thr-1/>

It is probably obvious that the equilibrium-based analysis is incomplete. In a world of complete knowledge, frictionless decision making, and perfect substitutability, all smokers would have already switched. A dynamic model with incomplete information about options and their characteristics, which are learned over time, and other sources of friction, and that also allows product preferences and health tradeoffs to differ across individuals, is needed to show why this switching is not immediate or complete. It would need to be populated with much more empirical information about individual characteristics.

A rational choice?

It is useful to circle back to a question that is undoubtedly vexing some readers: Is it legitimate to model tobacco/nicotine use as rational choice? The typical discourse does not even acknowledge that there are any benefits, let alone that they outweigh the costs. But claiming that hundreds of millions of people are making a choice that does not provide net benefits is an extraordinary claim. It defies our most basic knowledge of consumer choices and the supposed support for the claim is extremely dubious. Similarly, the fact that many smokers are happy to have quit does not mean those still smoking are being irrational, or even that those individuals were being irrational before – there is nothing irrational about either heterogeneous preferences or changing one's preferences over time. Proponents of that position seem to just be counting on no one ever mentioning either economics or the explicit testimony from consumers that they are deriving benefit.

Tobacco/nicotine use behaviour, like all human behaviour to some extent, involves important departures from the simple rational economic model. People do not have perfect information, and they sometimes react irrationally to what they know. People use rules of thumb rather than optimising every decision. But it should be remembered that the choice to smoke, in a world where the health risks are well known and aggressively communicated, is one of the weightiest decisions in someone's life, not a casual choice. Thus the assumption of rational choice is the reasonable starting point.

Bases given for denying rationality

Some claims by those who seek to deny that this is a benefit-motivated rational choice are so outlandish that they are not even worth addressing – e.g., that people are making a consumption choice that have enormous impacts on their psyches, lifestyles, disposable income, and health because they are permanently mesmerised by pretty packages or how suave movie characters look when they smoke. These are almost literal claims that consumers act because of demonic possession. Those who seek a legitimate argument, rather than just counting on people to believe any absurd claim, seem to realise they need to offer some valid departure to the standard economic model. On closer examination most such proposed departures are just be more complicated claims about demonic possession.

One claim is that most smokers really want to quit, a claim that is ostensibly supported by surveys in which many respondents answer simplistic questions ('Do you want to quit?') affirmatively. Putting this in economic terms, they are claiming 'most smokers would have higher welfare if they quit.' But quitting is always an option, so the affirmative survey response must not really be true in any normal sense of the word 'want'. Economics favours revealed preference (inferring preferences from choices) over merely asserted preference, and with good reason: talk is cheap. People frequently make claims about their preferences that contradict their cost-and-benefit motivated decisions. Thus stated preferences are generally considered unreliable for economic analysis when revealed preference is available. Claiming that hundreds of millions of people really do not want to make the choice they are making, and that they really prefer an alternative that is available and that they are aware of, is absurd. The behaviour-contradicting answers to those survey questions are probably largely explained by virtue signalling (giving the politically correct answer) and *second-order preferences* (they are saying that they want to *want to* quit smoking, even though they want to smoke). This is explored in detail in Phillips, Nissen, and Rodu (2016), which also shows how this further strengthens the case for promoting THR.

Addiction

The most common approach to trying to justify an anti-economics view is invoking the vague concept of addiction. But such claims do not constitute an argument against the economic model for three reasons. First, most use of the word is vacuous, with no statement of its definition, let alone implications. There is no accepted scientific or even medical/psychological definition of 'addiction'. To the extent that definitions exist, they usually describe a consumption pattern that is massively disruptive to short-term functioning, and thus do not characterise tobacco/nicotine use behaviour. If pushed to offer a definition, those who use the term in the context of tobacco/nicotine use generally offer a vacuous characterisation such as 'persisting in a behaviour despite its costs', which is captured in any economic model since *all* consumption has costs.

Second, as the term is usually used, addiction *is* an economic concept, describing preferences and choices. While it gets discussed as if it were a biomedical concept – because it has implications for health and sometimes involves identified biological pathways – ultimately it describes behaviour, not physiology. Thus it can *only* be analysed via economics.

Third, the viable candidate definitions for 'addiction' demonstrate how it fits within an economic model. There seem to be two categories of defensible candidate definitions. The first starts with the observation that someone's current net benefits from smoking are higher because he has a history of smoking. This in itself is neither harmful nor unusual; benefits of most goods increase as someone becomes more familiar or habituated. However, the history of consumption may also lower someone's baseline welfare, so that some of the benefits of each day's consumption are needed to merely bring the consumer back to the level of welfare he would have had he never used the products.¹⁰ But if this is addiction, it turns out to be moot for purposes of the economics of harm reduction. The preferences of current smokers are what they are, regardless of how they came about. This might mean that the demand curve drops sharply for never-smokers, which would move Q_2 toward Q_1 , but this does not change any of the qualitative implications of the economics. The benefits of *current* consumption are still positive, even under an (implausible) extreme version

10 In Becker and Murphy's (1988) presentation of this economic conceptualisation of addiction, they argue becoming addicted can be rational and welfare maximising, with the lifetime benefits outweighing the accumulated costs even though some ongoing consumption is needed just to get back to baseline.

of this story in which all benefits of smoking are merely compensating for the reduced welfare state caused by past smoking.

A second defensible definition of ‘addiction’, better labeled ‘dependence’, plausibly exists at a level that has practical implications. To put the claim in economic terms, individuals have multiple consumption equilibria and their consumer surplus is affected by their recent past consumption (it is unpleasant to be abstinent today if one was consuming the product yesterday). Thus, short-term optimisation may leave them stuck at local maximum that provides lower welfare than some global maximum.¹¹ That is, ‘addiction’ refers to a *short-term* change in tastes from *recent* past consumption. This might mean that while a smoker’s net benefits from smoking are currently positive, so she wants to smoke today, if she were to avoid smoking for a while this would no longer be the case, and then she would genuinely prefer to not smoke. It is possible that the nonsmoking state has higher net welfare for a current smoker, but she is not changing behaviour either because she does not realise that she would be happier at the other equilibrium or the short-term costs to make the transition are too daunting.

Being a genuine hypothesis, this is testable. The evidence suggest that it describes only a tiny minority of smokers. Failure to recognise that abstinence is not the global optimum for many smokers, and is often not even an equilibrium, seems to be the main reason why most smoking cessation aids are almost useless (see Phillips, Nissen, and Rodu (2016) for more details). For example, nicotine replacement therapy and other clinical approaches to smoking cessation are implicitly premised on the hypothesis that if smokers can just be assisted into a nonsmoking equilibrium they will stay there. Yet these methods cause temporary abstinence in only a small portion of smokers and, more important, the vast majority of those who achieve temporary abstinence resume smoking. With or without pharmaceutical assistance, many smokers go through reasonably long periods of abstinence, only to make the choice to start again. After overcoming the short-term effects of dependence and settling into an

11 *Multiple equilibria* describes a situation where there is more than one consumption choice that is *local equilibrium* – that is, it is preferred over any small deviation away from it. One of these equilibria is the global maximum (better than any other possible choice), but it is possible for someone to be ‘stuck’ at another local equilibrium and not want to move through the less pleasant space in between that and the global optimum where long-run welfare would be the highest. Think of a ball rolling on hilly ground; it will stop at a point where the contour is immediately uphill every direction, but this might not be lowest point in the area.

abstinence, many discover that it is *not* an equilibrium, much less their global optimum. That is, they still feel like their net welfare would be improved by smoking, and it is a constant fight to not do so. Others find themselves in a new local equilibrium but realise they preferred the smoking equilibrium and switch back.

However, it is consistent with the evidence that this story applies to a few smokers. Given the large population, presumably some would have higher lifetime welfare by being abstinent compared to switching to a low-risk product. However, this would require the extreme conditions described above, where his welfare from abstinence is higher than that from using some low-risk product by more than the $P_{h(cig)} - P_{h(lrp)}$ that would be gained by switching. This is an enormous margin that seems to defy 'did not realise' or 'not worth the short term cost' explanations for being stuck in the wrong equilibrium. That is, anyone who would only be a bit better off being abstinent rather than smoking would be even better off using a low-risk product; most everyone for whom abstinence is so much better would have already figured out that quitting was better than smoking, and would have done so. Thus, even for smokers who are trapped by dependence, for whom abstinence would be welfare-improving in the long run, it is still most likely that a low-risk product is welfare maximising compared to either abstinence or smoking.

Inadequate concern about future risk

The one apparent empirically and ethically defensible claim that this rational economic model is inadequate is that many people do not care 'enough' about future risk. This is a generic observation, but it is particularly relevant to questions about smoking because most of the costs (the health impacts) are uncertain and occur much later than the benefits. Some assertions of inadequate concern are really just statements of differing preferences ('smokers are not weighting future health costs as heavily as I would and therefore they must be irrational'). But more legitimate arguments can be offered.

Claims of inadequate concern about the future are often couched in terms of technically irrational patterns of discounting. For example, people systematically demand a much higher premium to wait one day for a payoff rather than getting it now, as compared to the premium they demand to wait an extra day for a payoff a year in the future. But the assumption that people

ought to treat one-off events as if they were returns on money in an investment account is dubious. For example, in a psychology lab experiment, walking away with £10 today rather than £20 tomorrow saves a trip back to the lab and guards against the possibility of clerical error or that the researchers are practicing some deception that includes not paying. Aphorisms about birds-in-hand are not without justification. Moreover, looming threats, in contrast with profits, actually tend to weigh heavier than just getting them out of the way. The better way to frame this seems to be Schelling's (1984) theory of multiple selves: we impose costs upon our future selves not because of technical financial-style discounting but because, in some sense, we are acting as if that future self is a different person onto whom we are offloading costs for current-self gains. Under that conceptualisation, ethical arguments can be made (though they are far from trivial) for protecting the future-self victim from the current-self actor.

The available research shows that most people, including most smokers, technically overestimate the risks from smoking compared to what the epidemiology shows (and this is older research; subsequent anti-smoking messaging has presumably increased this tendency). However, there are good reasons to believe that even people who cognitively overestimate such a risk act on gut-level heuristics that do not give it appropriate consideration. (For more depth on this, see Slovic (2001)). This can lead to inadequate concern for future risk compared to rational analysis.

Incorporating departures from rationality

Once the claims about departures from rational choice are expressed in concrete terms, it is easy to see that most are moot for present purposes or that only the rarest extreme cases represent departures from the general conclusion. Claims that most smokers really prefer not smoking do not stand up to even casual scrutiny and the data used to support them can be explained without reaching absurd conclusions. If addiction is merely acquired preference then it is a preference, and can be modelled as such. If one's position is that acquisition of this preference often reduces lifecycle welfare and is chosen by adolescents who do not understand its ramifications, that can be argued based on the economics (though such positions are generally merely asserted, not analytically supported). It could be argued, based on the multiple-selves concept, which is particularly compelling in the case of children harming their future adult-selves, that this justifies restrictions aimed at stopping non-adults from initiating

consumption (though, again, this needs to be argued rather than treated as if it goes without saying).

The case for dismissing the implications of the basic economic model becomes more tenuous when we are talking about adults. A reasonable working definition of 'adult' is someone empowered to make decisions with potentially momentous lasting implications – e.g., to cease investing in one's education, take on debt, or join the military. But adults do suffer from multiple-selves tensions, 'not going to happen to me' delusions, and technically irrational discounting behaviour. The challenge is that most every choice that could be explained by one of these could also be explained by some pattern of rational preferences. Positing that the choice was made irrationally, let alone the stronger claim that the individual would have been better off under the control of an external force that prevented or discouraged the choice, might be accurate, but it needs to be analytically justified. It is not sufficient to appeal to dubious answers to simplistic surveys let alone to say, in effect, 'I personally would not make that choice, so those consumers must be irrational', which is effectively equivalent to saying 'I do not like the taste of beetroot, and so anyone who consumes it is irrational.' Moreover, to be used as a justification for policy interventions, the effect needs to be quantified. It is difficult to envision any basis for such analysis – other than assuming demonic possession – that does not begin with the basic economic model.

Optimal rates of taxation

An analysis of tobacco/nicotine tax policy can help to further illustrate the core points about the economics, as well as answering some specific questions about how to optimise taxes to achieve some particular goal. It is easy to see that for any conceivable goal, the most common simplistic proposals are wrong. These include claims that low-risk product taxes should be on par with cigarette taxes as well as claims that taxes should be proportional to the level of risk.

The question of optimal taxation depends on the goal, which could be any of:

1. aligning consumer preferences with the real costs to correct perception errors (i.e., maximising consumers' welfare),
2. maximising net social welfare,
3. minimising health costs,
4. minimising consumption,
5. maximising revenue,
6. some objective function that balances two or more of these.

Goal 3 or 4 seems to be what many commentators have in mind when discussing the goal of taxation, though they seldom make this clear. In many jurisdictions, 5 is the real goal; this is not addressed here, since in addition to seldom being explicitly defended as an ethical goal, the analysis would require estimating consumer proclivity to use the black or grey market and any limits imposed by a sense of fairness, which are beyond

the present scope.¹² Category 6 includes infinite possibilities and, given the lack of anyone clearly stating a particular objective function, it is difficult to know which to analyse.

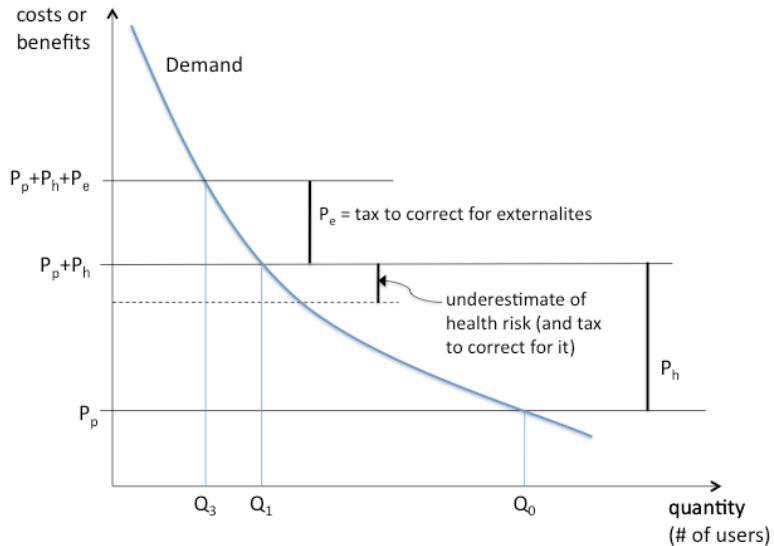
Maximising consumer welfare

Since normal policy analyses of financial incentives are usually about maximising welfare, consider this first. Consumers will fail to maximise their own welfare if there is inadequate concern for future risks. If consumers at the margin underestimate the magnitude of P_h – either literally not knowing the risk or acting as if the risk were lower due to psychological biases – then there will be too many product users. Some consumers who perceive their net benefits to be positive will be mistaken and will be making the wrong choice. A tax equal to the difference between the perceived health cost and the real health cost, as illustrated in Figure 2, would maximise their welfare by dissuading them from making the wrong choice. (Note that for purposes of this analysis we are starting from a base purchase price, P_p , that does not include the taxes and adding them in on top. This contrasts with Figure 1 where we took the taxes as given and included them in P_p .)

The naïve claim that taxes should be proportional to the risk seems to be motivated by a rudimentary understanding of the economics, with the assumption being that consumers completely ignore health costs when making a decision, and thus consumption would be Q_0 in the absence of taxes. That would imply taxes should be proportional to the risk, but since the assumption is clearly wrong this calculation is a moot exercise.

12 Taxes also result in the government being inclined to support the continuing sale of the products rather than interfere with it – i.e., it serves as protection racket money. While this clearly motivates the pro-tax lobbying efforts by some in industry (who gain the benefits of government support while consumers pay the tax), and it is not without merit from a short-run realpolitik standpoint, it clearly falls outside any normative optimisation analysis.

Figure 2



If consumers accurately consider future health risks or are overly fearful of them, then any positive tax rate moves consumers further from their optimal choice. This possibility is worth mentioning because it is actually quite difficult to make the case that low concern for the future is actually a welfare-harming inadequate concern, rather than a legitimate matter of preference. This means that quantifying the optimal level for the tax is basically impossible, so if an arbitrary level cigarette tax did manage to create the incentives that maximise welfare, it did so by blind luck. Still, the inadequate consideration of the future arguments suggest that would entail some positive tax.

The tax would be paid by consumers whose behaviour does not change, and so will decrease their welfare. This introduces the question of *whose* welfare is being optimised and how to trade off among those helped and those harmed: A tax that gets the incentives right and discourages some marginal users who were not really better off smoking improves their welfare, hurts the welfare of other smokers, and enriches the government. The simple response is to observe that the tax payments are a *transfer*, with the government gaining everything consumers lose, and to call that a wash, but it is a bit facile to shrug off a highly regressive tax as a mere transfer. The better conceptualisation is that getting the incentives right

for those on the margin hurts those who continue to smoke, so it is benefiting the former at the expense of the latter.

For the low-risk product it seems impossible to justify a tax based on inadequate consideration of future risk. The risks are so close to zero that there is almost nothing to be inadequately concerned about. Moreover, even if we assume these products post any net risk, almost every consumer grossly overestimates this by multiples. There is simply no basis for claiming that consumers are making choices that inadequately consider this risk. There is a far better case to be made that consumers (both smokers and non-users) irrationally avoid the low-risk products because they overestimate the risks, which means that the welfare-optimising tax would be negative (a subsidy). This is clearly not going to happen in general, though it is not entirely unrealistic: UK regulators have already approved one e-cigarette variation as a 'medicine', which may result in it being subsidised by the medical system (for smokers only, of course).

Maximising net social welfare

To expand to the full social cost-benefit analysis, we need to consider any net externalities, costs and benefits borne by someone who is not party to a transaction. It is widely believed that smoking costs the rest of society resources (usually described in terms of money) because of non-consumer-paid costs of medical treatment. In reality, the opposite is true. There are medical costs savings due to early mortality, making average lifetime medical cost about the same for smokers and nonsmokers, too close to definitively say which is higher. Thus when other foregone consumption is added in (most of the years of potential life lost due to smoking occur after retirement, so there is little foregone production), the externalities are clearly *positive*.¹³ However, there are also aesthetic externalities (smell of smoke, litter) and some health costs imposed on others (though not nearly as much as is generally claimed), and some argue that the lost 'productivity' of companionship due to early mortality after retirement should be counted.

Assume for illustration that the net externalities from smoking were negative, though it is actually very difficult to defend that position. Then the level of consumption is greater than the optimum and what is called a 'Pigouvian tax' should be imposed to make the price paid by smokers reflect the total social cost and not just his own personal costs. This is illustrated in Figure 2 by the line at $P_p + P_h + P_e$ where the last term reflects the total net cost of the externalities and thus the tax.¹⁴ Figure 2 also illustrates that correcting for both misperception of health costs and externalities can be accomplished by adding the two taxes. Assuming the net tax to optimise social welfare is positive, consumption with both of those corrections included would drop from Q_1 (where there were no taxes) to Q_3 .

But even if the assumption that smoking creates net negative externalities were true, there seems to be no possibility the net externalities from low-risk products are negative. Unlike the expensive lingering diseases caused by smoking (lung diseases, cancer), to the extent that there is any health risk from smoke-free alternatives, it seems to be dominated by cardiovascular stress. This tends to cause a rapid death for a clear net medical cost

13 Needless to say, this does not mean it is good someone dies from smoking. It merely means that the claim that smokers cost the rest of society money is clearly wrong.

14 An additional complication is that there is some *producer surplus* (profits), which increases net social benefits and should be counted on the positive side of the ledger.

savings, to say nothing of the reduction in other consumption. The products create minimal aesthetic impact and no health effects from environmental smoke, and littering is substantially reduced. Thus there is no apparent justification for any Pigouvian tax for low-risk products.

Minimising consumption

Some commentators suggest the policy goal should be to minimise smoking. This is probably usually just a sloppy way of trying to minimise health costs, which is taken up in the next section. But taking that goal as stated, an infinite tax on cigarettes (i.e., a prohibition on sales at any price) would be theoretically optimal. Since political realities and black markets mean that this is infeasible, the goal is best achieved by very high taxes, the maximum that is politically feasible without creating too much of a black market (this is the reality in many jurisdictions). Whatever the level of taxes on cigarettes, if the goal is to reduce smoking, the optimal level of taxes on low-risk alternatives is as low as possible (as heavily subsidised as possible, and at most zero), to encourage substitution.

If the goal were to minimise *total product use*, an infinite tax on all products is theoretically optimal. Assuming only finite taxes are practical, the model suggests a *higher* tax for low-risk products. They offer much higher consumer surplus, so a much larger financial punishment would be required to discourage their use. However, the ease of alternative supplies for e-cigarettes¹⁵ might mean that a high tax would create a black market that actually increased consumption compared to some lower tax level. Though the theory is straightforward to work out, some empirical heavy lifting would be needed to optimise this goal.

One might ask why it is even worth analysing this goal, given that there is no conceivable ethical justification for it. And yet it seems to describe a great deal of actual public policy. Cigarettes are banned in only a few totalitarian jurisdictions in the world, but in many jurisdictions there are longstanding bans on smokeless tobacco products (e.g. snus is banned in the EU outside of Sweden) and in many other jurisdictions they are taxed higher than cigarettes. The emergence of e-cigarettes has prompted bans in many jurisdictions, with a de facto near-ban moving forward in the USA, and many legislators are scrambling to figure out how to impose

15 See <http://antithrlies.com/2015/12/19/casaa-analysis-of-what-will-really-happen-under-ecig-deeming/> for an explanation of this observation.

high taxes on them. In general, as THR becomes more popular, more anti-tobacco efforts are shifted toward discouraging low-risk product use rather than discouraging smoking. These policies are not consistent with a goal of minimising health impacts, let alone maximising welfare, but they are reasonably consistent with the (ethically indefensible) goal of simply minimising total product use.

Minimising health risks

Minimising health risks is the most frequently claimed justification for taxes. It is worth noting that despite it being treated as self-evidently valid, this goal actually has very tenuous ethical justification. The ‘public health’ notion of elevating health to trump all other consumer costs and benefits – which is not limited to discussions of tobacco/nicotine – does not fit any defensible ethical model, and is clearly contrary to revealed preferences. It is often asserted that it is better to cause someone to be abstinent rather than use a low-risk tobacco/nicotine product, no matter how enormous her consumer surplus from such use. But this runs contrary to not just normal policy ethics but to the most fundamental tenet of health ethics, that people should be free to make informed health-affecting choices to maximise their welfare. Nevertheless, since this goal dominates much of the discussion, it is worth analysing in detail.

If there were no political constraints or black markets, the goal of minimising health risks would be identical to the ‘minimise use’ goal: The tax on cigarettes should be infinite, since there is no possible health advantage from letting anyone choose the most hazardous option. Moreover, if bans were effective, the optimal tax on the low-risk products would also be infinite because, with no smoking to compete with, there is no health gain to be had by allowing the use of the alternative (recall that the present analysis assumes that low-risk products still produce some net health risk).

The more useful analysis is to take as given the existing taxes on cigarettes, which are often as high as is politically feasible, and then assess the optimal tax rate for low-risk alternatives. Lower taxes encourage smokers to switch, which benefits health, but there is also a small health benefit from discouraging non-users who might want to start using low-risk products, and so a balance must be struck. (Recall that the goal here is not the ethically defensible goal of maximising welfare; if it were, a non-

user freely choosing to start would represent a benefit, not a cost.) The analysis requires two additions to the model.

First, the products cannot be treated as perfect frictionless substitutes, apart from their P_h . Otherwise, as noted above, the very different values of P_h would mean that everyone with access would have already switched to low-risk products. Since this has not happened, it must be that the non-health net benefits of cigarettes are much greater for many tobacco/nicotine users. This could be primarily because we have not achieved equilibrium, and many current smokers are simply sticking with what they are familiar with, or it could be that many smokers really prefer smoking enough that it overcomes the P_h differential, perhaps partially due to inadequate consideration of the risk. Either way, there is some price advantage that would cause a particular current smoker to switch products, thereby reducing total health costs. Put another way, we will now recognise that smokers who have not switched currently prefer smoking, but would prefer the low-risk product if its price advantage were great enough.

Second, it is necessary to introduce the concept of *cross-price elasticity of demand*. The concept of *own-price* elasticity of demand (often just called 'price elasticity') – the change in consumption that is caused by the change in the price of a good – has been present in this analysis from the start (it is represented by the slope of the demand curve), though the term has not been used. The cross-price elasticity is the change in consumption of a good when the price of another good changes. This is usually described in terms of the percentage change in consumption resulting from a percentage change in price, so a cross-price elasticity of 1.0 would mean that if the price of the low-risk product decreases by 1 per cent then consumption of cigarettes decreases by 1 per cent.

To complete the equation, we also need a cross-price elasticity for continuing to 'consume' abstinence, as a function of the price of low-risk products. Consumption prevalence of low-risk products is determined by how many people switch from either of the other consumption choices, cigarettes and abstinence. Since low-risk alternatives are not yet very popular in most populations, their use can be thought of in terms of their attraction from the much larger populations of smokers and non-users. This is a departure from the usual conceptualisation, which would be anchored on the current consumption of the good we are focusing on, but it is the better

model in this case (more details in the footnote¹⁶). Cross-price elasticity for abstinence works the same way, with a particular change in the price for low-risk products affecting the portion of the population choosing to switch between abstinence and low-risk product use.

With this we can calculate the health (or any other) impact of a change in taxes for the low-risk products taking everything else as fixed. The impact (measured as an increase in total health risk) for a $p\%$ change in the price of the low-risk product equals,

$$p E_{\text{cig},\text{lrp}} Q_{\text{cig}} (H_{\text{cig}} - R H_{\text{cig}}) + p E_{\text{abs},\text{lrp}} Q_{\text{abs}} (-R H_{\text{cig}})$$

where $E_{\text{cig},\text{lrp}}$ is the cross-price elasticity for cigarette consumption as a function of the price for low-risk products, $E_{\text{abs},\text{lrp}}$ is the cross-price elasticity for abstinence, Q_{cig} is the current number of smokers, Q_{abs} is the number who are abstinent in the population, H_{cig} is health risk from cigarettes, and R the comparative risk that converts H_{cig} into the risk for the low-risk product.

This is much simpler than it might look. All it says is that for a positive value of p (an increase in the price of the low-risk product), the number of smokers is increased by a factor of p times $E_{\text{cig},\text{lrp}}$ (at the expense of low-risk product use) and the number who stay abstinent rather than adopting the low risk product is increased by a factor of p times $E_{\text{abs},\text{lrp}}$. Everyone who makes one of those choices avoids the risk from the low-risk products (thus the minus signs on the R terms), but the smokers who make that choice have the H_{cig} risk they would have avoided if they switched. If a positive tax of $p\%$ is imposed then the first term will be a positive number, an increase in the total health costs due to discouraging

16 The standard own-price elasticity model, in which a price change for a good affects the consumption quantity proportional to the quantity of that good currently consumed, does not work in this case. Attracting 2% of the abstinent population to use a low-risk product would double the users in some populations, but increase it by 10-fold in others. Trying to model this in terms of own-price elasticity based on current low-risk product consumption is hopeless because current baseline level simply does not matter much (though it does affect the social dynamics that would attract new users, but that is a different analysis). Thus it is better to model that choice entirely in terms of the choice to switch from either smoking or abstinence, which have fixed prices in this scenario, as a function of the price of the low-risk products. This simplification is possible because the products are treated as pure substitutes in this model and are the only options. It would not work for many cross-price elasticity analyses – e.g., if we were differentially taxing carrots and potatoes, where they are only partially substitutes and each has an existing base of demand that has nothing to do with the other.

switching, while the second will be negative, a decrease due to discouraging new users.

For example, imagine a population of 1 million people, with 20 per cent smokers and 75 per cent abstinent (the others are already using a low-risk product and so, thanks to the simplification, do not enter the equation). Let $E_{\text{cig},\text{lrp}}=0.05$ and $E_{\text{abs},\text{lrp}}=0.05$. These are small numbers for price elasticities of goods that are good substitutes in the common-language sense of the word (e.g., butter vs. margarine has cross-price elasticity of almost 1), but we would expect that people's long-run decisions about choices this important are not hugely affected by price. If smokers are not already switching to the low-risk products, guessing that a 20 per cent reduction in their price would only attract 1 per cent of them to switch seems to be in the right range. In any case, the absolute magnitude of these numbers does not matter, only their relative values, and setting them equal was a very conservative assumption: In reality, it is almost certainly the case that smokers are more motivated by the comparative price of the products than those who are abstinent (i.e., quite a few smokers might think, 'e-cigarettes are now 20 per cent cheaper than smoking; I will give them a try because of that alone', while extremely few non-users ever think, 'I would start using e-cigarettes if only they were 20 per cent cheaper').

With $R=.01$ and just leaving H_{cig} as a constant, the net population health impact of a 1 per cent increase in the purchase price of alternative products (a 1 per cent ad valorem excise tax) is:

$$.01 \cdot .05 \cdot 200,000 \cdot (.99 H_{\text{cig}}) - .01 \cdot .05 \cdot 750,000 \cdot (.01 H_{\text{cig}}) = 95 H_{\text{cig}}$$

This means that the 1 per cent increase in the price of low-risk products would increase total health costs in this population by 95 smokers worth. Breaking it down, this includes 375 people whom the tax discourages from adopting low-risk products, and so remain abstinent, but whose health benefits are dwarfed by the 100 additional smokers who do not switch. Thus even with the conservative assumption that current non-users are just as price-sensitive as smokers, it is clear that any positive tax is harmful. If the value of p changes, it would change the magnitude of the outcome, but not its sign.

If we made the scenario even more favourable for a tax and dropped the cross-price elasticity of smokers to $E_{\text{cig},\text{lrp}}=.005$ (that is, a higher price for the low-risk products discourages switching almost not at all, and has a

much bigger impact on the decisions of non-users, an utterly implausible scenario), there would still be a net increase in population health risks from any positive tax on low-risk products. If R were increased to .05 (clearly much greater than the observed risk from smokeless tobacco, and it would be astonishing if e-cigarettes were that harmful), the result from the above equation would drop only to $76 H_{\text{cig}}$, still representing a clear negative population health impact from taxing the low-risk products. The only plausible change that would reverse the sign of the impact would be if the total number of smokers in the population was very low.

As has been pointed out for decades, given how low the risk of alternative products is, it is functionally impossible for their *existence* to increase total population health costs, no matter how popular they are with those who would otherwise be abstinent (and thus bans are always harmful to population health). With only a few percent of would-be smokers choosing the low-risk alternatives instead, the net effect would be a reduction in risk even if every non-user adopted the products.

That observation is not sufficient to conclude that for small price changes, a tax could never improve population health. However, the present analysis clearly demonstrates that this too is the case. Until we get to a time where the number of smokers is very low, population health risk is minimised if there is zero tax on low-risk products (or, better, a subsidy, if that were possible), whatever the tax on cigarettes.

It is worth recalling the extreme claim about addiction, that the only benefit from tobacco/nicotine consumption is to relieve the desire created from past consumption, which is often an implicit premise in anti-tobacco policies. This claim is equivalent to saying $E_{\text{abs},\text{lrp}}=0$, since abstainers have no desire to start using a low-risk product at any price, and so low-risk product should be given to adults for free to minimise population health risk, because they would only be used as a substitute for smoking. If we wanted to calculate the optimal subsidy level for realistic values of $E_{\text{abs},\text{lrp}}$, we would need to determine real-world estimates for the parameters in the equation. Of course, there is little reason to do this calculation for most populations since only zero or positive taxes are a realistic possibility. Should the NHS consider subsidising a medicalised e-cigarette product, they would be wise to do some version of this calculation.

Discussion

This relatively simple application of basic economic modelling provides definitive answers to a number of questions that are not intuited by most of those writing or making policy about THR.

Choosing low-risk products, assuming their gross benefits for consumers is anywhere close to that from smoking, appears welfare-enhancing for most smokers and would-be smokers. The welfare benefits from the reduced health risk are so large that it seems unlikely that many smokers prefer the experience of smoking enough, as compared to some low-risk alternative, that they would not be better off switching. The portion for whom that is true is an empirical question, of course, that cannot be answered by theory. But the theory suggests it is appropriate to actively try to combat the misinformation or inertia that keeps many smokers in a lower-welfare state. It is difficult to see any welfare loss that could come from making sure that smokers are aware that there are low-risk alternatives and enticing them to give them a try; those who still prefer to smoke can do so. However, there is a large welfare loss created by efforts to push or manipulate smokers to become abstinent, since almost all would be better off switching rather than quitting entirely.

The simple model also makes clear that low-risk product use is welfare-enhancing for some people who prefer abstinence to smoking. Thus, if there is availability and honest information about low-risk products, there will be more total tobacco/nicotine use as compared to a world where smoking is the only option. Claims by some THR advocates that there is no evidence that low-risk products will attract some would-be abstainers are simply not defensible. Such claims are likely to become obviously false over the next few years, as e-cigarettes grow in popularity, and it is already evident in Sweden where more men use tobacco/nicotine than

would be expected to smoke in the absence of snus.¹⁷

The better argument is that this prospect is a benefit, not a cost, as illustrated by the model. As with any other good, anyone who freely chooses to adopt a particular consumption option does so because it provides an increase in welfare. No one is made worse off by informed free choice so long as we use a legitimate measure of well-being (i.e. we do not pretend people are making decisions based on overall preferences but really only care about health risks). Moreover, though some consumers would be increasing their risk, there is no conceivable scenario that could lead to a net increase in total population risk. This is not an inevitable implication of free choice, like the utility benefit is, but is a result of just how low the risks of the low-risk products are.

A bit more analysis shows that for almost any of the obvious goals, the optimal tax rate on low-risk products is zero (or negative if a subsidy is an option). The only exception is the ethically indefensible goal of minimising consumption of all tobacco/nicotine products, regardless of welfare or health effects, in which case the tax should be as high as possible, and probably higher for low-risk products than it is for cigarettes. Sadly, many laws and regulations seem to best conform to the prescriptions of that ethically indefensible goal.

All of the results here, with the exception of a few quantitative asides, are based on theory rather than empirical evidence. A theory might be wrong in any particular case. However, what is presented here is based on the most established bedrock theory of the science of welfare, preference, and choice. It would be quite extraordinary if any of the very basic principles were fundamentally misleading in this case. Expanded modeling is needed examine dynamic effects, including the positive contagion effects of low-risk product adoption (which would tend to further favor encouraging adoption) and any supposed 'gateway effect' (though there is no apparent basis for suggesting this occurs; see: Phillips 2015). Some of the implications will change if smoking rates drop very low.

However, limitations of the basic model do not suggest that political activists can deny its results by just waving their hands and saying 'it is not that simple' or relying on the unstated and undefended claim that hundreds of

17 This is apparent by comparing rates of total tobacco use to Swedish females, for whom snus use was not so socially accepted (see chart in Rodu, 2013) or to other northern European countries (OECD 2013).

millions of people are making choices, day after day, that are contrary to their preferences. One of the great advantages of formal models is that they force disagreements to take place in the realm of concrete claims and actual analysis (or at least starkly illustrate when disagreements have no such merit). Those who would challenge these results need to produce a well-defined alternative model or a variation on the present model that produces different results. This appears to never have been done. This economic model is the empirically established and ethically accepted basis for analysing most policy decision, whereas the competing view seems to be little more than a modern-day version of 'they are all possessed by evil spirits and need an exorcism.'

It is always the case that empirical observation trumps theory; if experience shows that the theoretical conclusions are wrong, then the theory needs to be fixed. But this principle refers to solid, replicable, carefully examined empirical results. When a one-off result with an ad hoc analysis contradicts a well-established theory, the better conclusion is that there is probably something wrong with the data or analysis. This is clearly the case for the naïve claim that people who are choosing to smoke really would prefer to not smoke. Moreover, it is important to understand that every interpretation of empirical observations is filtered through (usually unstated and undefended) theory, and if that theory is wrong, then the conclusions are likely to be wrong. In many cases, defensible theory alone provides far better information than any available empirical evidence. Because we had an established germ theory of disease, it was quite clear that AIDS was an infectious disease before HIV was identified; empiricism that was 'informed' by earlier medical theories would have blamed it on witchcraft or the gods hating gay men. Similarly, it has long been clear that many oral cancers are caused by infectious transmission, yet they were still widely blamed on tobacco products, just as cervical cancer was blamed on numerous causes before it was discovered that HPV infection causes almost all cases.

An economic model of tobacco/nicotine consumption is not merely theoretically defensible and empirically supported by analogy, but it passes the test that is often taught as the *sine qua non* for distinguishing science from religion: making empirically refutable predictions, that if refuted would lead to the conclusion that it was wrong. The alternative view that consumption is caused by latter-day witchcraft (the magical thrall of marketing and some vague concept of addiction) does not pass the test; it makes few claims that are testable and those that are testable have

repeatedly proved to be false. Ironically, it is also possible to make testable predictions about those who refuse to acknowledge that tobacco/nicotine use has benefits: When the population shifts toward the new equilibrium of higher usage prevalence of low-risk products, the advocates of the demonic possession theory will ignore the actual evidence that people like the products and attribute the increase in prevalence to industry marketing efforts.¹⁸

People who are not empowered by a bit of scientific analysis can easily be tricked into believing clearly inaccurate claims that just sound good ('marketing causes all product use'; 'taxes should be proportional to risk') or are supposedly backed by data ('the limited adoption of smokeless tobacco in most places means that smokers are not interested in reducing their risks'). Scientific theory, even when it does not offer an answer to a question straightaway, points out how to ask the right questions and where to express doubt. Sadly, the study of tobacco harm reduction is dominated by doctrines, with evidence misinterpreted as supporting absurd claims. The vast majority of these come from those who are opposed to all use of tobacco/nicotine as a matter of personal pique, but some come from other quarters. It is hoped that this presentation can help immunise the readers against some of the naïve conclusions that are produced by this pseudo-science.

18 It is worth noting that I originally wrote this sentence in 2013, for an early draft of text that was later incorporated into this paper. It seems to already be proving correct regarding the younger cohorts where e-cigarette uptake by nonsmokers is greatest.

Acknowledgements and Competing Interests

The author is a longtime advocate of tobacco harm reduction and liberty-preserving public health policy. This paper was author-initiated and received no direct or specific financial support. During the course of development of this analysis which spanned three years (a draft version of part of this work was first circulated in 2013), the author served as the Chief Scientific Officer of The Consumer Advocates for Smoke-free Alternatives Association and received compensation from that organisation, and conducted research that was funded by grants from BAT, JTI, and Imperial Tobacco (none of the grant funding involved work on the present analysis). The author thanks Elaine Keller, Brian Carter, Atakan Befrits, Igor Burstyn, Christopher Snowdon, and an anonymous reviewer for comments on various versions of this work.

References

Becker, G. and Murphy, K. (1988) A theory of rational addiction. *Journal of Political Economy* 96: 675-700

Lee P. L. and Hamling J. (2009) Systematic review of the relation between smokeless tobacco and cancer in Europe and North America. *BMC Medicine* 7: 36

OECD (2013) OECD Factbook 2013: Economic, Environmental and Social Statistics. Chart at: http://www.oecd-ilibrary.org/sites/factbook-2013-en/12/02/01/smoke_g3.html?contentType=&itemId=/content/chapter/factbook-2013-98-en&containerItemId=/content/serial/18147364&accessItemIds=&mimeType=text/html&_csp_=b363ad2af2448f26d48113c09a7ac2b8

Phillips C. V., Rabiou, D., Rodu, B. (2006) Calculating the comparative mortality risk from smokeless tobacco versus smoking. *American Journal of Epidemiology* 163 (11): S189

Phillips C. V. (2009) 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. *Harm Reduction Journal* 6: 29

Phillips C. V. (2015) Gateway Effects: Why the Cited Evidence Does Not Support Their Existence for Low-Risk Tobacco Products (and What Evidence Would). *International Journal of Environmental Research and Public Health* 12(5): 5439-5464

Phillips C. V., Nissen C. M. and Rodu, B. (2016) Understanding the evidence about the comparative success of smoking cessation methods: choice, second-order preferences, tobacco harm reduction, and other neglected considerations. Under review: Available at <http://ep-ology.com/2014/11/24/working-paper-phillips-nissen-rodu-understanding-the-evidence-about-the-comparative-success-of-smoking-cessation-methods-choice-second-order-preferences-tobacco-harm-reduction-and-other-neglecte/>

Rodu, B. (2013) Population-level proof: Snus saving Swedish smokers' lives as tobacco consumption declines. Tobacco Truth, 25 September 2013: <http://rodutobaccotruth.blogspot.com/2013/09/population-level-proof-snus-saving.html>

Schelling, T. (1984) Self-Command in Practice, in Policy, and in a Theory of Rational Choice. *The American Economic Review* 74(2): 1-11.

Slovic, P. (2001) Cigarette smokers: rational actors or rational fools? in P. Slovic (ed.), 'Smoking: Risk, perception, and policy'. Thousand Oaks, CA: Sage

The Institute of Economic Affairs
2 Lord North Street
London SW1P 3LB
Tel 020 7799 8900
email iea@iea.org.uk

