How Much Ongoing Smoking Reduction is an Echo of the Initial Mass Education?

Carl V. Phillips, PhD
Marewa Glover, PhD

Objectives: In this paper, we attempt to quantify the "echo" effects of the downward shock in US smoking prevalence from mass education starting about 1965 through 2010. **Methods:** An agent-based population simulation replicates the observed effects of the initial education shock on smoking prevalence, and then estimates ongoing echo effects based on empirical estimates of the effects of parental smoking on initiation and peer-group quitting contagion. Further simulations estimate what additional echo effects would explain the entire historical reduction. **Results:** About one-third of the observed prevalence decline through 2010 can be attributed solely to fewer parents smoking after the initial education shock. Combining peer-group cessation contagion explains well over one-half of the total historical prevalence reduction. Plausible additional echo effects could explain the entire historical reduction in smoking prevalence. **Conclusions:** Ongoing anti-smoking interventions are credited with ongoing reductions in smoking, but most, or perhaps all that credit really belongs to the initial education and its continuing echoes. Ensuring that people understand the health risks of smoking causes large and ongoing reductions. The effect of all other interventions (other than introducing appealing substitutes) is clearly modest, and quite possibly, approximately zero, after accounting for the echo effects.

Key words: simulation; smoking; health education; tobacco

***Am J Health Behav.*™ 2022;46(1):84-95**
DOI: https://doi.org/10.5993/AJHB.46.1.8

A reader of the corpus of program evaluation literature from tobacco control will be struck by the observation that if one accepts every claim of success and substantial effectiveness in the literature, then in Western countries, where many such interventions have been done, one would conclude that the current smoking prevalence must be close to zero. A closer examination of the methods and data reveals that the claims of success are not so well supported. However, the summary claims in most papers in the literature, and even more so, the overarching summaries (eg, WHO Framework Convention on Tobacco Control (FCTC) reports on the successes of its MPOWER interventions) imply that every category of intervention and almost every evaluated specific intervention had a substantial impact. Given that smoking remains common in many places that have implemented many interventions, the common claims must be wrong, in part, or entirely. Thus, it is worth asking if something other than ongoing interventions might explain the reduced smoking prevalence.

One problem with the literature is a common problem in population sciences – many exposures that are hypothesized to cause a particular outcome often occur together. Doing an experiment is rarely an option for population-level interventions, let alone for assessing the causes of trends. The independent effects (to say nothing of countless interactions) need to be statistically separated, which is usually an unrealistic goal.

For example, it is almost impossible to sort out the health benefits of eating more of a particular food, or even a category of foods, because the same people who eat more of one healthy food also eat more of other healthy foods and less unhealthy food. In addition, healthier dietary choices are associated with other healthier behavioral choices due to

Carl V. Phillips, Epiphi Consulting, Nashua, NH, United States. Marewa Glover, Centre of Research Excellence: Indigenous Sovereignty & Smoking, Auckland, New Zealand.
Correspondence Dr Phillips; cvphilo@gmail.com

common causes, including wealth, environment, and individual preferences. It is almost inevitable that a study that tries to estimate the independent effects of one of these behaviors will assign it some, or all the credit for the effects of its many correlates. Usually there is an attempt to "control for" other variables, but these attempts never do so fully. If we were to combine the estimates in the literature that are declared to be the benefits of each individual healthful behavior, we would "learn" that people with a realistic constellation of particularly healthy behaviors will live hundreds of years longer than those making unhealthy choices.

The same problem exists for evaluating tobacco control interventions. Every tax increase, media campaign, prohibition, or other action takes place in an environment where other interventions exist. If we assume some of these have some effect, then it is almost impossible to have enough data and a sufficiently correct model to avoid giving undue credit to the specific intervention under study. But co-occurring interventions are not the only challenge. Ongoing interventions take place in a context of the ongoing secular trends created by the initial downward shocks in smoking prevalence that resulted from education efforts, starting in about 1965 in Western Anglophone populations and at various times in other populations. Estimates of the effect of a current intervention easily can claim credit for secular trends that are an echo of that initial education. Indeed, the typical common claim of credit for the collection of tobacco control interventions, that actions that took place during the 3 decades leading up to 2010 should be credited with the entire decrease in smoking prevalence during that period, steals all the credit owed to ongoing effects of the initial education, claiming that credit for later interventions that may not have had any substantial effect.

We know that choosing to smoke is socially contagious – the more people around someone who smokes, particularly their parents, the more likely they are to start smoking.[1] Parental smoking is the most consistent strong predictor of whether a teenager (of a particular age, in a particular population) will start smoking.[2] Smoking prevalence among siblings, peer groups, and the wider community affects uptake via overt and subconscious social signaling. All of these are taken as fact in the scientific literature and in tobacco control politics, where they are cited as motivation or points of leverage for interventions. But one important implication – that a downward shock or trend in smoking prevalence will, by itself, cause further downward trending for more than a generation – is generally ignored.

Similarly, smoking cessation is a contagious behavior.[3] This is particularly clear for switching to a lower-risk alternative, wherein the person quitting smoking demonstrates to their social contacts that the choice is appealing and educates them about the alternative. However, even if the choice of cessation method is not affected by social-contact education, the demonstration effect of quitting itself is still powerful. Seeing a friend quit smoking takes it from being an abstract possibility to a concrete example of success. In addition, simply having fewer people who smoke in one's social circles encourages quitting. Each of these, and all of them together, creates a positive feed-forward effect from any smoking reduction.

Thus, a one-time permanent downward shock in the popularity of smoking – like that caused by initial education about the harms from smoking – causes a long tail of transition to a new lower equilibrium, echoes of the initial shock. If many people quit smoking, then many more who would have started smoking had they come of age earlier will not do so and others will be motivated to quit over time. The subsequent cohorts coming of age not only will experience the effect of the downward shock, but also be subject to less social contagion. There will be a new equilibrium, but it will only be reached slowly, with a substantial portion of the effect taking more than a generation. This will happen with or without any further efforts to discourage smoking. Subsequent interventions could still have effects beyond the secular trend toward a new equilibrium, of course, but it makes no sense to try to quantify those effects without trying to estimate the background effects of the echoes alone.

Looking at a graph of the decline of smoking in the United States (US). and other Western countries over the decades, it does not look like a public health eradication success story, as with polio or lead in paint and gasoline. It looks more like a graph of how much airtime music from 1965 gets on the radio or how many World War II memorial gatherings occur. We would not assume

that the ongoing decline of those phenomena is the result of new influences. It is an inevitable time trend. The cohorts who were most invested in the behavior die off and the social contagion effect is far below replacement value, so popularity wanes.

Western smoking prevalence peaked (as best we can tell from the imperfect historical data) about 1965 in the US at a bit under one-half the adult population.[4] Then, there was a shock to the system that included the alarming 1964 Surgeon General report summarizing the growing body of empirical evidence that smoking had substantial ill-health effects. The new understanding inspired a decade of mass advertising of that message, via media, clinical advice, warning labels, and social chatter. From 1965 through the mid-1970s recognition that smoking is extremely unhealthy went from expert knowledge, to joking recognition ("cancer sticks – ha ha!"), to accepted common knowledge. At the end of this process, it was widely ingrained that smoking is extremely harmful, foolish, and practically suicidal.

Efforts to make this knowledge universal across the population substantially reduced smoking prevalence. But that information shock – the collective events that began with the Surgeon General's report and immediate reaction and continued with educational efforts over another decade – may have been the only actions ever taken that substantially reduced smoking, up until the time that low-risk alternatives start to substantially replace cigarettes in a population. If so, this would be rather useful to know, because it would suggest that efforts to bring a population to widespread understanding of the major health impacts is the only tobacco control intervention worth trying to replicate in populations where it has not yet happened, the only one that is not a waste of resources and political goodwill.

This might seem like an absurd suggestion considering an extensive body of literature that supports the claim that the effects of ongoing interventions (restrictions on supply and use, cessation services, and denormalization strategies, etc) explain the continuing decline in smoking prevalence. Why would someone ask "could echoes have caused this," if we already know that other interventions caused it? It turns out that we do not already know this. A comprehensive review of the literature the FCTC cites to support the claim that their "MPOWER" interventions have substantially reduced smoking prevalence concluded that most of that literature did not provide the claimed evidence, or at best, offered only tenuous support for it.[5] In summary, the review concluded that there was fairly strong evidence that regressive taxes modestly reduce smoking and that restrictions on where people can smoke reduced smoking intensity, but no compelling support existed for the claims that other interventions had any effect. It does not appear plausible that the measured effects could come close to explaining the entire trend. This motivated us to ask, if even the strongest proponents of the claim that ongoing interventions explain the continuing decline cannot cite convincing evidence that this is the cause of the continuing decline, then what else might explain it?

This is not to suggest there is affirmative evidence that none of these interventions have much effect – absence of evidence is not evidence of absence. The observation is just that there is little to no solid evidence various interventions had much or any effect, and thus, it is plausible that most or even all the credit they are claiming is really owed to echo effects. Proximity and representativeness biases – the tendency to attribute outcomes to something that is closely related in terms of time, intent, or other characteristics -- make it easy to believe that the ongoing interventions simply must explain the decline, and thus, echoes cannot. The most stated, and seemingly, convincing reasoning that echo effects do not explain the ongoing decline is circular: Smoking prevalence continues to decline, and this is not attributed to echo effects because there is little to no acknowledgment of echo effect, so it must be due to ongoing interventions. Because it is due to ongoing interventions, it cannot be the result of echo effects.

The purpose of the current analysis is to estimate how much of the downward drift in smoking in the US population from the mid-1970s to 2010 is an echo of the shock caused by mass education in the decade before that period, and to assess the hypothesis that it might all have been. This is done by modeling the behavior that would have occurred had the initial shock and its echoes been the only forces acting on the population and comparing that to what occurred. The analysis period ends in 2010 when vaping became a popular alternative and caused an increase in quitting and a reduction,

or at least downtick, in smoking initiation.

Any attempt to simulate the behavior of a counterfactual population cannot demonstrate with certainty that the simulated counterfactual is accurate. This is true for every epidemiologic regression model and also is true for this systems dynamics model. However, as with a regression model, if the suppositions that went into creating the structure and counterfactual contrast are plausible, then the implications of the results are informative.

GENERAL METHODS

This section provides an overview of the methods, while the specific methodologic details employed for each version of the model are presented as each version of the model is presented. This helps clarify the exposition and avoids forcing the reader to try to match passages from this section to each individual simulation.

The model is designed to replicate the historically observed prevalence decline for 1965-1978, when the initial education shock was playing out, and there was little in the way of additional interventions, and then extrapolate its effects. The analysis consists of a dynamic agent-based (also known as complex systems, or microsimulation) population dynamics simulation using the open-source programming system NetLogo Version 6.2.0[6] (https://ccl.northwestern.edu/netlogo/; the final versions of the programs are available elsewhere).[7, 8, 9] The agents are individuals who age and reproduce new cohorts. Time periods are one year.

Each agent has an initial propensity to smoke. There is a global threshold such that anyone with a higher propensity begins as a smoker and those with a lower propensity begin as a nonsmoker. Note that for parsimony, characterizing labels are substituted for proper person-first language. This is not dehumanizing in this case because we are talking about simulated agents and not real people.

The shock is modeled as a threshold drop in 1965. Social contagion effects influence individual propensities, as described for each version of the model. Anyone who shifts to the nonsmoker side of the threshold has a 10% probability of quitting each year. The latter structural choice produces a reasonable replication of the 1965-1978 curve. It could be replicated by parameterizing the shock as a series of threshold drops over the historical period of sharp drop (and slightly reducing the contagion effects), but modeling a single threshold change with non-immediate reactions is more parsimonious and seems more realistic. It is in keeping with the conceptualization of previous modelers[10, 11] who identify the 1964 Surgeon General's report as a singular shock that precipitated a series of changes.

The model maximizes simplicity to minimize the possibility of complexity obscuring the clarity of its implications. Specifically, the population consists of 10,000 simulated agents, 18-80-year-olds (ages selected at random, with a flat distribution for simplicity). It is simplified to a one-parent one-offspring model, with each 41-year-old adding an 18-year-old offspring in the population and each 80-year-old exiting the population. Smoking behavior is simplified, such that there is only starting as a smoker and quitting, not initiation. This can be thought of as assigning someone's persistent smoking behavior in early adulthood back to the agent at age 18, ignoring the noise created by the stopping and starting that is common in late adolescence. Historical prevalence statistics effectively do the same, with the behavior changes contributing only a temporary bit of noise that then disappears.

The parent's smoking status affects that of the offspring, which is the only social dynamic in the first version. Then, a social network for the contagion effect of quitting is added. These 2 simulations estimate the effects of the 2 most important and reasonably well-quantified sources of echo effects. In the third simulation, additional effects of lower population prevalence are added. These are more speculative, but as described below this is not a problem given the different goal of that simulation.

A burn-in period (typical practice in modeling like this, wherein the model runs with inputs chosen to force the target state for starting the analysis period) generates an equilibrium that produces the historical 1965 prevalence; it is not meant to represent pre-1965 data. The burn-in generation exits by 1965, having spawned a generation of offspring who experience the shock (the Shock Generation, who consist of older Boomers and their elders) who then start producing a new generation (the Post-Shock Generation, who consist of younger Boomers, Generation X, and a few older Millennials at the end). Because there are random elements in the simulation dynamics,

what we refer to for parsimony as an equilibrium here is not an equilibrium by the strictest sense of the term. The actual values will randomly depart from the quasi-equilibrium.

The outcome measure is smoking prevalence, with separate statistics for the Shock and Post-Shock generations. It is not possible to determine the exact historical numbers because available statistics are inconsistent and far less certain than widely believed. Even different government measures differ markedly,[12,13] and estimates by generation must be interpolated from age-based reporting. The most important numbers for the present model are based on picking an approximate average among various available statistics – average smoking prevalences of 48% in 1965, 35% in 1978, and 20% in 2010. Figure 1 shows simulated data with those points and smoothed curves in between. This is an illustration to ease interpretation of the later figures. It does not represent one specific real historical data series, which are all considerably noisier in addition to being mutually inconsistent. Some of these show an apparently real uptick in smoking in Generation X in the early 1990s that settled back to the trendline by 2010, which the present models do not attempt to replicate.

Figure 1
Trend in US Smoking Prevalence, 1965-2010 (Smoothed Simulated Data)

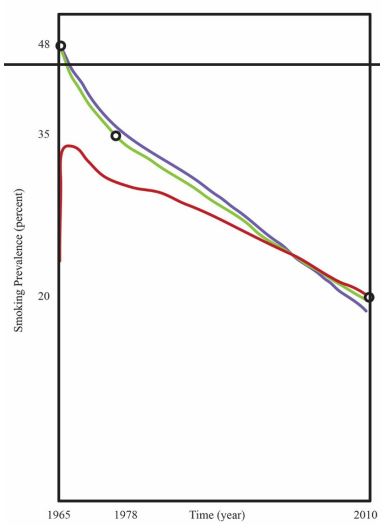

Note.
(purple) pre-shock generation ("Boomers")
(red) post-shock generation ("GenX")
(green) population average

There is legitimate disagreement about what the right historical trendline is, and so one should note that the qualitative and the approximate quantitative results presented below do not depend on accepting these historical estimates. The interpretations simply would need to be adjusted slightly for alternative estimates. In particular, the 2010 prevalence estimate used here errs on the side of optimism, with some authoritative estimates putting the prevalence several points higher than 20%, and thus, the stated quantitative conclusions are arguably too conservative.

Simulation 1 – Echo from Reduced Parental Smoking

The most robust predictor of individual smoking initiation, apart from place and time, is whether a parent smokes. Historical estimates reviewed in Leonardi-Bee, Jere, and Britton[1] suggest having a parent who smokes at the time of coming-of-age approximately doubles the probability of smoking uptake. The first simulation estimates the echoes caused by this alone. This presents the simplest possible introduction to the concept.

Specific methods – simulation 1. The offspring has a higher propensity to smoke if the parent smokes at the time they come of age. There is no other social contagion. The parental effect is set such that offspring of smokers are always twice as likely to smoke as those of nonsmokers. That is achieved by each offspring of a nonsmoker parent coming of age having a propensity that is a random integer from 0 to 99, and those of a smoker parent having (x + random number from 0 to 99-x), where x is determined by solving the system of equations [Ps = 2Pn, Pn = (99-T)/100, Ps = ((99 - (T-x))/(100-x)], where Ps = the probability an offspring of a smoker smokes, Pn = the probability an offspring of a nonsmoker smokes, and T is the threshold, which is determined empirically to match the pre- and post-shock prevalences as previously described.

Thresholds of 57 before the shock, increasing to 71 in 1965 replicate target historical values. In the modeling jargon, these inputs are used to "tune" the model to get the target historical values. Specifically, they result in a pre-shock equilibrium of 48% adult smoking prevalence and a post-shock prevalence of 35% in 1978.

RESULTS OF SIMULATION 1

Figure 2 shows the smoking prevalence in this simulated population from the initial equilibrium, through the shock, and through 2010. By construction, the initial equilibrium and immediate-post-shock prevalences are at the target values described in the methods.

Figure 2
Output from Simulation 1

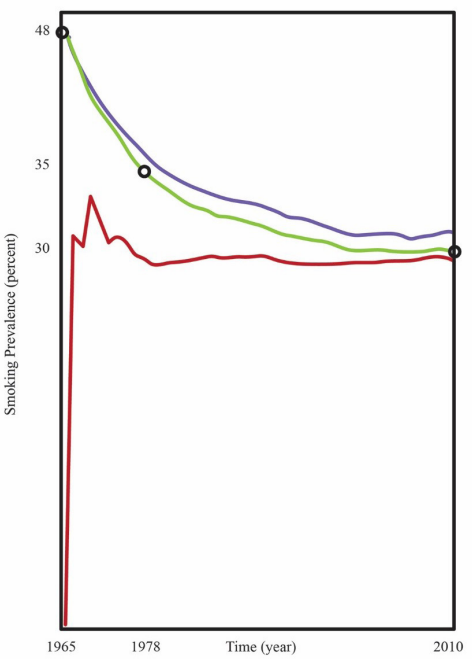

Note.
(purple) pre-shock generation ("Boomers")
(red) post-shock generation ("GenX")
(green) population average
Every run is slightly different due to random processes, but the year-percentage function presented by the three noted points is the result in almost every run. The initial instability of the new generation is the inevitable result of the small early population size.

Over the post-shock period, the smoking prevalence drops another 5 percentage points, to 30%, which is about one-third of the historically observed decline from 1978 to 2010. This is just the effect of parental contagion. It is primarily the effect of a generation whose prevalence rapidly dropped from about 50% to about 35% producing offspring with lower smoking initiation than previous generations, and who have just started to produce a new generation that continues this effect. The new generations partially cohort-replacing some of their forebears. The generations alive at the time of the shock continue to react slowly to the information shock (the geometric decay from quitting being probabilistic after the threshold is crossed), reaching about 32% prevalence.

If this dynamic continued, the new even-lower-prevalence generations would produce progeny who continued to replace the oldest generations. This would eventually reduce prevalence to a new equilibrium a few percentage points lower still. This, of course, has no practical implications, but it illustrates the ongoing feedback in the real-world system.

In this scenario, no one quits smoking except due to the shock to the threshold. Smoking diminished over the lifecycle during this period (as observed in the 2014 Surgeon General's report[14] and elsewhere). This could be added as a background phenomenon, which would increase the effect in this simulation, because it increases the effect of younger generations on overall prevalence. However, we also know that smoking cessation is contagious, and it can be introduced as a second set of echoes.

Simulation 2 – Echo from the Contagion Effects of Smoking Cessation

The contagion effects of smoking cessation are well-known and generally accepted but are much less well-quantified and are more difficult to model. The most informative study of this by Christakis and Fowler[3] suggests that having at least one peer (friend/coworker/sibling) who quit smoking decreases the chances that someone remains a smoker by one-fourth to one-half. This was based on only those peers who were also part of the study, so does not account for all such influences. It is much higher if that peer is a spouse, but the model does not have spouses. However, the data are necessarily somewhat fuzzier in terms of what constitutes the relevant social network and the time period over which most of this effect occurs (that data covered a similar time period to what is being modeled herein). Those authors focused on a view of stable social networks and the dichotomy of having at least one peer who quit, and the observation that this created social pockets of residual smokers after many others had quit. There are other possible ways to conceptualize and model both phenomena. Unlike the previous scenario of parental effects, which is based on the effects of

a well-quantified and easily defined relationship, and thus, offers a solid simple model of a well-established quantification, this scenario should be viewed as a rougher and seemingly conservative estimate of how to model having a quitter within one's social network.

Specific methods – simulation 2. Starting from the previous version of the model, each smoker is now assigned influential peers consisting of 2 random members of their own age cohort who smoke at the time of cohort creation, 2 random agents from any cohort who smoke or have quit, and the agent's parent (who might not smoke, in which case this relationship does not matter). If at least one of a smoker's influencers has quit smoking, then there is a 1% chance each year that the smoker will quit. This results in a quit rate a bit over one-fourth during the simulation time for the average smoker with an influencer who quits, in keeping with estimates of the rough empirical evidence available.

Obviously, this network is small and unrealistic in many ways. It is intended merely to quantify a particular reasonable – seemingly conservative – estimate of how much quitting contagion effect echoes from the round of quitting created by the shock. Real-world influencer relationships are usually reciprocal, and without that feature this model will not produce the clusters that Christakis and Fowler[3] observed, but that is not important for present purposes. Another unrealistic feature is intentional – no one is influenced by anyone quitting for reasons other than the shock. The point of the exercise is to isolate the echoes.

Because the quitting contagion starts to contribute to reduced smoking prevalence in 1965, alongside the direct effect of the changed threshold, the threshold shock (ie, chosen to construct historical numbers) must be re-tuned to a lower difference. Specifically, the shock increases the threshold from the original 57 to only 69 instead of 71.

RESULTS OF SIMULATION 2

After again quasi-equilibrating at 48% in 1965 and dropping to 35% in 1978, by construction, smoking prevalence in this model drops faster than the first model with just the effect of parental smoking. By 2010, prevalence drops to about 26%. This is considerably more than one-half of the total reduction in prevalence from 1978 to 2010 (Figure 3).

Figure 3
Output from Simulation 2

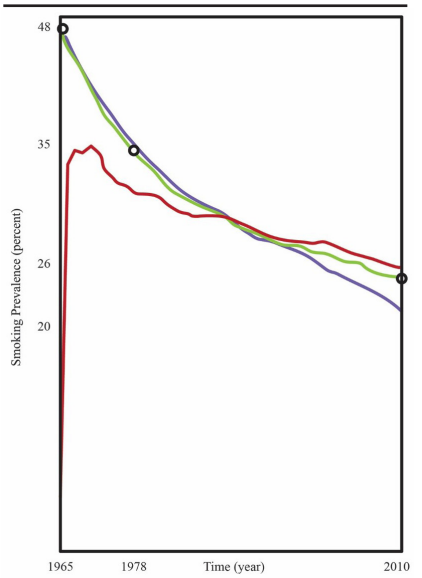

Note.
(purple) pre-shock generation ("Boomers")
(red) post-shock generation ("GenX")
(green) population average
Every run is slightly different due to random processes, but the year-percentage function presented by the three noted points is the result in almost every run. The initial instability of the new generation is the inevitable result of the small early population size.

Compared to Simulation 1 the smoking prevalence of the shock generation is lower because many of them quit over time due to the contagion, and not just due to the shock at the start. The prevalence of their progeny is higher because of the reduced threshold change that is needed to generate the 1978 target prevalence. In this scenario, additional contagion quitting by the shock generation further lowers the prevalence by 1978, necessitating this. The result is patterns for the 2 different generations that are remarkably close to their respective trendlines (smoothed to ignore Generation X's temporary spike above the trend in the 1990s). This unintended result is reassuring about the model given that the targets and resulting parameters were based entirely on whole-population prevalence, with no attempt to tune to replicate historical numbers for generations separately.

This result also replicates the phenomenon that, for the relevant period, there was a steady increase in the probability someone has quit smoking as

they age through late adulthood. Many members of the generation who experienced the shock but kept smoking went on to quit (as happened in real-life but not in Simulation 1) as a result of the quitting contagion echo. The accumulating motivation to quit created by influencers in their lives who had quit drives the smoking prevalence of the shock generation below that of the younger generations by 2010, as actually occurred.

Simulation 3 – Can Echoes Explain the Entire Reduction in Smoking Prevalence?

The purpose of this scenario is to take a cut at addressing the question: "Is it plausible that the echoes explain all of the reduction in smoking through 2010?" It is important to note the epistemic difference between this and the previous 2 simulations. The previous 2 consisted of estimates of the effects particular echoes did have, starting with external information and assessing its implications. The present simulation is intentionally reverse-engineered to provide an example of how increased echo effects plausibly could explain the entire reduction in smoking prevalence through 2010.

Lower population smoking prevalence discourages smoking uptake and encourages quitting via pathways other than those estimated in the first 2 simulations. There have been attempts to estimate some of these effects independently – everything from the specific effects of sibling smoking, to how often someone sees anyone smoking – but these estimates are generally dubious. Trying to add them up would create double counting of the sort noted in the introduction. It is not possible to make an empirical estimate of a global net effect of higher populations' smoking prevalence because of the lack of a suitable counterfactual for comparison. However, it is still possible to posit a collection of pathways, apart from the 2 already modeled, by which lower population prevalence reduces initiation and increases quitting, and investigate their implications.

This simulation implements what seems to be the simplest version of that – for every percentage point that population smoking prevalence is lower than baseline (48%) there is a reduction in smoking uptake and/or some increase in quitting, which stack with the effects in Simulation 2. The reduction in smoking uptake is modeled by reducing smoking propensity at the time of entering the population as a linear function of the difference between the current prevalence and the pre-shock prevalence of 48%. The increase in quitting is modeled as an additional chance of quitting, beyond the 2 already existing quitting pathways, also proportional to that difference. Both effects only begin in 1965, so the burn-in is not affected. The question being answered here is how large each of these needs to be such that they result in a prevalence of 20% in 2010.

RESULTS OF SIMULATION 3

Experimentation with the model shows that a reduction in the new generation's smoking propensity score of 0.7 for every percentage point reduction in population prevalence below baseline (without changing quitting behavior) is sufficient that the echo effects lower population smoking prevalence to about 20% in 2010. This is after re-tuning the model to reduce the shock to the threshold by 2 points to keep the 1978 target prevalence because the original period of decline is now partially caused by this new effect (so the additional reduction needs only to be what comes additionally after that period). To put this in practical terms, recall that the propensity is a scale from 0 to 99 and that the initial information shock moves it by 10 points (down from 12 in the previous scenario). This means that the post-shock generation experiences an echo in the form of its own (slow-motion) propensity shock, with a magnitude about double that of the original shock experienced by its forebears (who are not affected by this). This seems a bit too large to be plausible.

To get the target final prevalence by increasing quitting only, it is sufficient to assign each smoker a chance of quitting of 0.05% per year for each percentage point below baseline prevalence (in addition to already-existing chances of quitting following the shock and from peer contagion). The re-tuning needed for the threshold shock was the same as in the previous paragraph. This affects all generations about equally. The meaning of this magnitude is more difficult to intuit. It perhaps helps to think that when prevalence is down to 20%, the echo from the reduced prevalence gives each smoker a 1.4% chance of quitting each year (on top of the contagion effect). This does not seem implausible by itself, although the change is concentrated in the older generation, which is a bit

imbalanced compared to the historical data.

Combining lower magnitudes of the effect for both changes also gets to the target final prevalence. Taking half of each of the changes from the previous 2 paragraphs – plausible quantifications for the extra echo effects of population prevalence – lowers 2010 prevalence to about 20%, with the 2 generations' prevalences almost equal (Figure 4). Anything along the spectrum of weighted averages of the 2 has a similar outcome (their effects do not combine exactly linearly, but it is close to that).

Figure 4
Output from Simulation 3, Half-and-Half Version

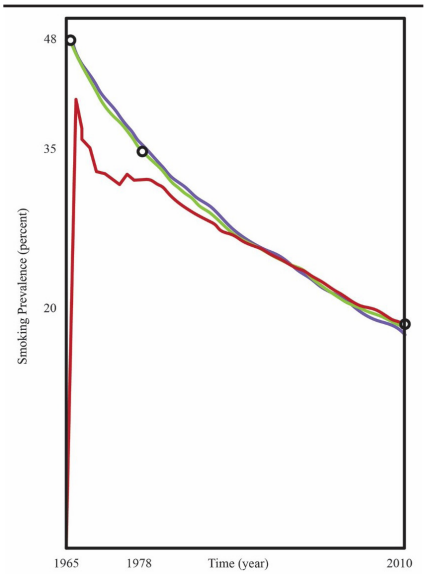

Note.
(purple) pre-shock generation ("Boomers")
(red) post-shock generation ("GenX")
(green) population average
Every run is slightly different due to random processes, but the year-percentage function presented by the three noted points is the result in almost every run. The initial instability of the new generation is the inevitable result of the small early population size.

DISCUSSION

It seems safe to conclude that a large portion of the claim "look at how much these 3 decades of tobacco control measures reduced smoking prevalence!" is stolen valor. Even if intuiting the quantification of this is not possible, it should be obvious that there must have been some echoes from the initial education shock. As for the quantification, the second simulation of just 2 solidly documentable echo effects, seems to explain more than half of the historical reduction in smoking prevalence (up until the time e-cigarettes became popular). This means that, at best, all the subsequent tobacco control interventions are left competing for credit for a minority of the total reduction. So, not only do program evaluations of individual interventions tend to claim some of the credit due to other interventions (assuming any of them are due any credit at all), but they also claim credit that is due to none of them.

Moreover, seemingly plausible additions of additional echo effects can produce outcomes where the entire net decrease is explained by echoes. Obviously, the ability to simulate such a result does not prove that subsequent tobacco control measures deserve zero credit for the reduction in smoking prevalence. However, it is sufficient to show that this conclusion is plausible. It also means that the ongoing reductions in smoking prevalence do not constitute evidence that ongoing tobacco control measures have any effect on smoking prevalence. Those policies have many effects in terms of hurting the welfare of people who continue to smoke (including regressive taxation of those who can least afford it, stigmatization and loss of social life, struggles with cognitive dissonance, and a loss of trust in medical and health institutions).[15] These are widely said to be justified by the reduction of smoking caused by the policies. However, if it turns out these policies are not really reducing smoking prevalence much, the ethics become much more fraught. We should note that we analyzed only smoking prevalence and not intensity (the quantity consumed by each person who smokes), which also decreased over this period, and that causes health benefits beyond prevalence changes. It is possible that this, too, might substantially be echo effects, although it is also possible that contemporary policies explain this trend, as noted in the introduction. There is simply inadequate empirical data about influences on smoking intensity to assess this. If it were the case that a policy reduced smoking intensity, even if it did not reduce prevalence, that would need to be part of the ethical calculation.

As with every statistical model, the validities of the resulting estimates from the first 2 simulations are, of course, only as good as the validity of the input assumptions. The input assumptions in the present case are transparent and seemingly

defensible. Every population science model, whether a simulation or a regression equation, is imperfect, and thus, we would not suggest that any exact quantitative results should be interpreted as precise. However, having some estimate based on a reasonable model is better than failing even to attempt to quantify a phenomenon, and then assuming that the quantification must be zero. Presumably, no one would explicitly argue that the lack of a perfect estimate means that the true value is zero, but unexamined conclusions that effectively assume this logic are common. For the case of the echo effects, this accidental erroneous logic can be found in the assumption that ongoing reductions in smoking prevalence must be caused by ongoing policies. It even can be found juxtaposed with arguments that social contagion effects are among the reasons for supporting one or more of those interventions, even as the implications of those same effects from the most successful historical intervention are implicitly assumed to be zero.

The information value of the third simulation is somewhat different, showing the plausibility of the hypothesis that no tobacco control measures after the initial education had any effect. The inputs were transparently reverse-engineered rather than being empirical estimates. Still, they are sufficiently plausible to reconcile the observation that there is no convincing evidence that most tobacco control interventions had or have any effect with the observation that smoking prevalence has continued to decline.

To summarize, this analysis and the existing literature suggest that there is strong evidence for the following 2 claims about the US and similar populations: (1) The immediate effects of lowering people's inclination to choose to smoke, via educating them about the substantial health risk, reduced smoking prevalence by about one-third over about a decade; and (2) The only other anti-smoking measure that is remotely close to such effects is encouraging substitution of a low-risk alternative (as observed in Sweden, Norway, and Iceland, and to a lesser extent, in the United Kingdom and the US). A large portion of the ongoing reduction for more than a generation after the education shock (unless widespread substitution starts to occur) was a series of echoes of that shock.

It seems likely that the following is true – most of the ongoing reduction for more than a generation of Americans (after the education shock, before widespread substitution started to occur) was echoes of that shock, and something similar will probably be true in other populations.

The following 2 claims are uncertain, and each seems about equally plausible and supported by about the same quality of evidence: (1) Other anti-smoking interventions may have reduced smoking prevalence a little bit; and (2) All the reduction in smoking prevalence through about 2010, other than the modest measured effects of punitive taxes, might be echoes of the initial education shock.

No descriptive quantitative analysis is sufficient to make policy recommendations in the absence of serious ethical or policy analysis, but there are some implications here that seem wise to consider in designing policies that are intended to reduce smoking. At best, the last few generations of tobacco control efforts have had modest effects all combined, and many of those efforts presumably have had no effect at all. Plausibly, the combined effects have been trivial. Thus, among historically attempted policies to replicate, the only one that seems to have a high probability to have a large effect is to determine if there is any population that lacks understanding of the risks equivalent to Americans in about 1975, and to provide them with accurate information. After that is accomplished, the most effective next step might be to figure out something new to do, such as encouraging product substitution, because it is not clear that the other arrows in the traditional tobacco control quiver have any bite at all.

Human Subjects Approval Statement

This study did not involve human participants; therefore, it did not require ethics committee approval.

Conflict of Interest Disclosure Statement

This study was conducted as part of a wider program of research being undertaken by the Centre of Research Excellence: Indigenous Sovereignty & Smoking. The funding for that program of work was obtained following submission of a researcher-initiated application for a funding grant from the Foundation for a Smoke-Free World, Inc. ("FSFW"), a US non-profit 501(c)(3)

private foundation. This study is, under the terms of the grant agreement with FSFW, editorially independent of FSFW. The contents, selection and presentation of facts, as well as any opinions expressed herein, are the sole responsibility of the authors and under no circumstances should they be regarded as reflecting the positions of FSFW. Neither the authors nor the Centre of Research Excellence: Indigenous Sovereignty & Smoking have any commercial interests in vaping, snus, oral nicotine or tobacco products. CVP has a corpus of past work that challenges specific claims that anti-smoking interventions are effective. During the course of his career, CVP has received research grants, salary, or consulting fees from almost everyone with a financial interest in the sale of tobacco products, including national and state/provincial governments of the U.S. and Canada, manufacturers of cigarettes, vaping, snus, and other tobacco products, pharmaceutical companies, and insurance companies, as well as advocacy organizations attempting to eliminate use of these products, advocacy organizations attempting to defend consumers' rights to choose these products, and publishers who focus on these products. His only such funding during his work on this project came from the grant in support of this project. M.G. has, over 10 years ago, received fees from pharmaceutical companies for consultancy re cessation medicines.

Acknowledgements

The authors thank Igor Burstyn and Kimberly Phillips for helpful comments on the conceptualization and details of the simulations.

REFERENCES

1. Leonardi-Bee J, Jere ML, Britton J. Exposure to parental and sibling smoking and the risk of smoking uptake in childhood and adolescence: a systematic review and meta-analysis. Thorax. 2011;66(10):847-855. https://doi.org/10.1136/thx.2010.153379
2. East K, McNeill A, Thrasher JF, et al. Social norms as a predictor of smoking uptake among youth: a systematic review, meta-analysis and meta-regression of prospective cohort studies. Addiction. 2021;116(11):2953-2967. https://doi.org/10.1111/add.15427
3. Christakis NA, Fowler JH. The collective dynamics of smoking in a large social network. N Engl J Med. 2008;358(21):2249-2258. https://doi.org/10.1056/NEJMsa0706154
4. Cummings KM, Ballin S, Sweanor D. The past is not the future in tobacco control. Prev Med. 2020; 140:106183. https://doi.org/10.1016/j.ypmed.2020.106183
5. Phillips CV. Does FCTC policy really advance public health? A review of 'WHO Report on the Global Tobacco Epidemic'. Report prepared for Foundation for a Smoke-Free World. 2019.
6. Wilensky U. 1999. NetLogo. Center for Connected Learning and Computer-Based Modeling, Northwestern University. Evanston, IL. http://ccl.northwestern.edu/netlogo/
7. Philips CV, Glover M. How much ongoing smoking reduction is echoes of the initial mass education? First simulation model. Zenodo. 2022. https://doi.org/10.5281/zenodo.5899482
8. Philips CV, Glover M. How much ongoing smoking reduction is echoes of the initial mass education? Second simulation model. Zenodo. 2022. https://doi.org.10.5281/zenodo.5899525
9. Philips CV, Glover M. How much ongoing smoking reduction is echoes of the initial mass education? Third simulation model. Zenodo. 2022. https://doi.org.10.5281/zenodo.5899544
10. Holford TR, Meza R, Warner KE, et al. Tobacco control and the reduction in smoking-related premature deaths in the United States, 1964-2012. JAMA. 2014;311(2):164-171. https://doi.org/10.1001/jama.2013.285112
11. Warner KE, Sexton DW, Gillespie BW,.et al. Impact of tobacco control on adult per capita cigarette consumption in the United States. Am J Public Health. 2014;104(1):83-89. https://doi.org/10.2105/AJPH.2013.301591
12. Rodu B, Cole P. Smoking prevalence: a comparison of two American surveys. Public Health. 2009;123(9):598-601https://doi.org/10.1016/j.puhe.2009.07.014
13. Rodu B. Tobacco Truth: How Many Americans Smoke? https://rodutobaccotruth.blogspot.com/2011/09/how-many-americans-smoke.html. 2011. Accessed 27 August 2021.
14. U.S. Department of Health and Human Services. The Health Consequences of Smoking: 50 Years of Progress. A Report of the Surgeon General. Atlanta, GA: U.S. Department of Health and

Human Services, Centers for Disease Control and Prevention, National Center for Chronic Disease Prevention and Health Promotion, Office on Smoking and Health, 2014. Printed with corrections, January 2014.
15. Burgess DJ, Fu SS, van Ryn M. Potential unintended consequences of tobacco-control policies on mothers who smoke: a review of the literature. Am J Prev Med. 2009;37(2 Suppl): S151-S158 https://doi.org/10.1016/j.amepre.2009.05.006