Kicking the habit is hard: a hybrid choice model investigation into the role of addiction in smoking behaviour

John Buckell*, Nuffield Department of Population Health, University of Oxford
David A Hensher, Institute for Transport and Logistics Studies, University of Sydney Business School
Stephane Hess, Choice Modelling Centre and Institute for Transport Studies, University of Leeds

* corresponding author: Richard Doll Building, Old Road Campus, University of Oxford, Oxford; john.buckell@ndph.ox.ac.uk

Highlights:

- We model addiction as a latent variable
- We study the impact of addiction on smokers' tobacco product choices
- We study how preferences and willingness to pay vary with addiction
- We predict the impact of lowering nicotine in cigarettes on smokers' short-term choices

Abstract

Use of choice models is growing rapidly in tobacco research. These models are being used to answer key policy questions. However, certain aspects of smokers' choice behaviour are not well understood. One such feature is addiction. Here, we address this issue by modelling data from a choice experiment on US smokers. We model addiction using a latent variable. We use this latent variable to understand the relationship between choices and addiction, giving attention to nicotine levels. We find that more addicted smokers have stronger preferences for cigarettes and are unwilling to switch to e-cigarettes. Further, addicted smokers value nicotine in tobacco products to a much greater extent than those that are less addicted. Lastly, we forecast short-term responses to lowering nicotine levels in cigarettes. The results suggest that current nicotine-focused policies could be effective at encouraging addicted smokers to less harmful products and lead to substantial public health gains.

Keywords: tobacco; addiction; hybrid choice model; experience-conditioned choice model; willingness to pay and accept; stated choice experiment

JEL codes: C35; I12; I18

Conflict if interest: the authors have no competing interests

Acknowledgements and funding: Research reported in this publication was supported by grant number P50DA036151 from the National Institute on Drug Abuse (NIDA) and FDA Center for Tobacco Products (CTP). The content is solely the responsibility of the author(s) and does not necessarily represent the official views of the National Institutes of Health or the Food and Drug Administration. Stephane Hess acknowledges support by the European Research Council through the consolidator grant 615596-DECISIONS.

Introduction

In the economics of risky behaviour, and in particular tobacco, the use of discrete choice models has proliferated in recent years (Regmi et al., 2017; Pacula et al., 2019). Standard choice models that are commonly used try to explain differences in sensitivities/behaviour towards products' attributes as a function of individuals' characteristics, i.e. through differences across sociodemographic groups. In so doing, there is the possibility that elements of the cognitive decision-making process are suppressed, and these elements can be important for understanding choice behaviour (Vij and Walker, 2016). Encouragingly, a trend in the risky behaviours literature, which follows the broader choice modelling literature (Hensher, 2015; Balbontin et al., 2019), has been to use more advanced choice models in attempts to understand more complex behaviours of individuals than permitted with standard models. Examples include incorporating measures of risk preferences into models of choice behaviour (Ida and Goto, 2009); the segmentation of individuals into groups or types of individual (Marti et al., 2018); accounting for optimization errors in individuals' choices (Kenkel et al., 2017); accounting for aversive choice behaviours (Buckell and Sindelar, 2019); and embedding real-world behaviours into models based on experimental data (Buckell and Hess, 2019). This information is crucial for policymaking; failure to capture these behaviours can lead to misguided policy recommendations (Buckell and Sindelar, 2019). Given that the complexion of tobacco markets has changed markedly in the past few years with the arrival of e-cigarettes, and that a variety of new policies are being enacted, it is now critical to understand the complex cognitive processes underpinning smokers' choices so as to better inform regulation.

Perhaps surprisingly, one aspect of behaviour that has not been incorporated into choice models is addiction (although there are studies of habitual behaviour such as modal choice in transport, where individuals are captive to a single mode). It is well-known that addiction plays a central role in smokers' behaviours (West, 2009; West and Brown, 2013; Wehbe et al., 2018). Indeed, economists have developed theories for, applied models to, addiction using longitudinal smoking data (see comprehensive literature reviews in Chaloupka and Warner, 2000; Cawley and Ruhm, 2011; DeCicca et al., 2020). While there are many studies of addiction-based, longer term tobacco use behaviour, the role of addiction in smokers' shorter-term choice behaviour is not well understood. Moreover, many studies have found heterogeneity in how specific subgroups (defined by observed individual characteristics) react to individual attributes. But it is unclear as to whether these are true preferences, or that the individual characteristics are linked to addiction, and it is the underlying addiction that impacts on choices. This is a key issue because many policies are directed at both addiction (i.e., nicotine policies) and vulnerable subpopulations, and the effectiveness of policies can be improved with a greater understanding of addiction-related choice behaviour.

Attempting to incorporate addiction into choice models is beset by a number of issues. Firstly, addiction is both multifaceted and unobservable (Collins and Marks, 1991; Shadel et al., 2000). Typically, only indicators of addiction are used, such as the number of cigarettes smoked per day. The number of cigarettes smoked is not a measure of addiction, but a function of addiction (and the underlying direction of causality is unclear). Secondly, there is not a one-to-one correspondence between these indicators and addiction. These are, at best, imperfect measures for the underlying metric of interest. For example, a more addicted smoker with high will power (or lower income) could smoke fewer cigarettes per day than a less addicted smoker with low will power (or higher income). In this case, using cigarettes per day as a direct measure of addiction is incorrect and it should only be used as an indicator of addiction.

There are many other such indicator measures. Thus, it is not always clear which indicators should be used to measure addiction or whether to use the full set of available indicators. Using only a subset of the indicators risks overlooking key information or misattributing effects. On the other hand, using all indicators poses significant problems for modelling, because using many indicators that are also likely to be highly correlated leads to a proliferation of parameters and technical issues such as collinearity. Several indices have been developed that sum these indicators, such as the Fagerstrom Test of Nicotine Dependence (FTND; Heatherton et al., 1991) for cigarettes, or equivalent measures for e-cigarettes (more than 10 now exist, see Bold et al., 2018). This is not an ideal solution because unweighted summation, or equal weighting, may be inappropriate (Fayers and Hand, 2002). In addition, these indices and their constituents are still product-specific and thus only capture addiction towards a specific type of cigarette (though some studies have sought to attend to this issue, e.g. Shiffman et al., 2004; Shimman and

Sembower, 2020). Moreover, any indicators that are collected but are not included in the index will be discarded; it is less than ideal to disregard potentially useful behavioural information.

Whether using a single indicator or multiple ones, the issue of causality remains. The indicators are a function of addiction rather than a direct measure thereof. Moreover, there is likely to be correlation between these indicators and other unobserved effects at the individual level that influence choice behaviour. Using these indicators as error-free variables in the model thus potentially leads to endogeneity bias, i.e. breaching the independence assumption of the explanatory variables and the error term. It is for these reasons that using latent variables to try to identify the underlying berhavioral drivers of these indicators is becoming more common, because identifying these underlying drivers can help to avoid these issues (Shiffman et al., 2004; Strong et al., 2015; Strong et al., 2017).

In this paper, we develop a choice model capable of handling the present issues in measuring and incorporating addiction into the analysis of smokers' choices. This model draws from two areas in the choice modelling literature: hybrid choice models and experience-conditioned choice models. Hybrid choice models allow for a latent, or unobserved, variable to be specified and estimated with the choice model in a system of equations (Abou-Zeid and Ben-Akiva, 2014). As such, it is well-suited to capturing addiction, the nature of which is inherently latent (see e.g. Shiffman et al., 2004, Strong et al. 2015, and Strong et al., 2017 for latent variables to measure addiction in tobacco); whilst we are not the first to take a latent variable approach to measuring addiction, we are the first to use it in a hybrid choice framework to analyze smoking choices. Figure 1 is a schematic of the modelling framework. Within the system, the latent variable is used to explain observed variables, such as indicators of addiction. Thus, these variables do not enter the choice model directly, avoiding possible endogeneity issues. Moreover, this framework can accommodate any number and form (i.e. the nature of the data) of indicator, using all the available information. The latent variable is then used in the utility function in place of the indicators (which would be a more traditional approach). Because there is only a single addiction variable, having to specify a large number of parameters directly in the utility function is thus also avoided. Of course, additional parameters are required for the measurement model components that are used to explain the values observed for the indicators.

The experience-conditioned choice models (Hensher and Ho, 2016; Hensher et al., 2019; Balbontin et al., 2019) are based on the notion that preferences and choice behaviours are in part a function of a given individual's prior experience with a given product or service. In the context of smoking, this implies a behavioural assumption that tobacco product choices are to some extent determined by the past use of tobacco products. This is highly appealing because the idea is applicable to both the addiction to nicotine and formation of habits that are associated with longer term, habitual tobacco product use (Wehbe et al., 2019). For this reason, we test addiction-conditioning in our model[1].

Developing this model allows us to overcome a set of difficult empirical issues and specify a behaviourally appealing model of short-term smoking choices. This model embodies a more sophisticated depiction of smokers' cognitive decision-making processes than in previous work. We use this model to study the relationship between addiction and choice behaviour, examine smokers' willingness to pay for nicotine in tobacco products, and predict the impact of lowering nicotine in cigarettes (which has recently been proposed by the US government). With this, policymakers are better informed in key issues around smoking and addiction. The remainder of the paper is set out as follows. In section 2, we set out the model and its features. In section 3, the results from the model are presented. Section 4 summarizes and discusses.

Methods

Experiment and Data

[1] In preliminary modelling, we tested a series of specifications based on experience-conditioned choice model versus a conventional hybrid choice model (i.e. with interactions between the latent variable and elements of the utility function). We found that the hybrid choice model led to better fit of the model and we thus used this approach. However, we note both the similarity to, and relevance of, the experience conditioning model for this application. The model in this form allows for addiction to condition utility, but in varied directions (as opposed to conditioning impacting all attributes in a single direction which is the case in the standard experience-conditioned model).

Data are taken from a labelled smoking choice experiment in which 1,531 adult smokers chose between cigarettes, e-cigarettes and an opt-out option[2], labelled "none of these" (Buckell et al., 2018; see Fig 2 for a sample choice scenario). Products were described by four attributes: nicotine, flavour, health harms and price. These attributes and levels were defined according to literature reviews, pilot studies, consultation with subject matter experts, and according to policies that the FDA could implement. Restrictions on attribute levels were made to make the experiment more realistic, for example, fruit and sweet flavored cigarettes are not available on the market, so we did not allow them in the experiment; attribute level balance was maintained in the design as far as possible by the design software, Ngene (Choice Metrics, 2018). Table 1 shows the products, attributes and levels. The design was Bayesian D-optimal, using priors obtained from a pilot study of 87 respondents. The design was based on the main effects only (i.e. without interactions). Individuals each answered 12 choice sets, which balances concerns of learning and respondent fatigue (Hess et al., 2012). A total of 36 choice sets were divided into three blocks, and respondents were randomized to each block in the ratio 1:1:1. Choice sets were presented in the same order within blocks. Sampling was based on quotas, defined using the Behavioural Risk Factor Surveillance System (BRFSS) data in 2013/14, based on gender, age, education and region, to make the sample representative of US smokers. Table 2 shows the descriptive statistics for individuals in the sample. The sample size was sufficient to ensure statistical power for the main parameters and is larger than most DCEs in health (de Bekker-Grob et al., 2015). Data quality measures, including minimum time thresholds, forced responses, attention checks, cheap talk, duplicate ID checks, and practice choice scenarios for respondents, were taken to promote data quality. Failure of any of these checks resulted in respondents being ejected from the survey.

A survey was collected alongside the experiment. In this survey, sociodemographic information on respondents was collected. We also collected revealed preference data on respondents' tobacco behaviours. This data includes product use and products' attributes such as prices and flavours. For addiction, a number of indicators of addiction were collected. These were daily smoking, number of cigarettes smoked per day, time before first cigarette of the day is smoked, time since last having smoked a cigarette that day, time since last having smoked a cigarette in the last few days/weeks, the number of quit attempts in the past year, e-cigarette use, frequency of e-cigarette use, and current urge to smoke (i.e. craving). These measures include those which may be considered outcomes such as cigarettes per day versus symptoms such as craving (cf. Strong et al., 2015; Bold et al. 2018).

Summary statistics for these measures are shown in table 2.

Choice Models

In a standard random utility model, the utility U_{nit} that individual *n* derives from product *i* in choice *t* comprises a systematic component, V_{nit}, and a remaining error term, ε_{nit} that follows an iid type I extreme value distribution, such that:

$$U_{nit} = V_{nit} + \varepsilon_{nit} \tag{1}$$

The systematic component incorporates two parts, with:

$$V_{nit} = \delta_{ni} + \beta_n x_{nit} \tag{2}$$

where δ_{ni} is a constant capturing product-specific preferences for alternative *i*, where this includes cigarettes, e-cigarette and the opt-out. β_n is a vector of estimated sensitivities capturing the impact of changes in explanatory variables x_{nit}, where this includes nicotine, flavour, health harm and price. Both δ_{ni} and β_n are person-specific, where the former is also different across alternatives.

[2] Since there are two alternatives for each label (i.e., each product and the opt-out), there may be difficulty in interpreting preferences for each label separately. To aid interpretation, we use generic constant terms for each label in our models.

We next allow for deterministic and random heterogeneity across individuals in the values of both δ_{ni} and β_n, both directly, and through the latent addiction variable. In particular, we have:

$$\delta_{ni} = \mu_{\delta_i} + \lambda_{\delta_i} z_n + \sigma_{\delta_i} \xi_{n\delta_i} + \tau_{\delta_i} \alpha_n \tag{3}$$

and for the coefficient associated with attribute $x_{ni_k t}$:

$$\beta_{nk} = \mu_{\beta_k} + \lambda_{\beta_k} z_n + \sigma_{\beta_k} \xi_{n\beta_k} + \tau_{\beta_k} \alpha_n \tag{4}$$

In this specification, μ_{δ_i} and μ_{β_k} capture mean values in the sample population for δ_{ni} and β_{nk}; λ_{δ_i} and λ_{β_k} capture shifts in their values as a function of socio-demographic characteristics, z_n; σ_{δ_i} and σ_{β_k} capture random heterogeneity, where $\xi_{n\delta_i}$ and $\xi_{n\beta_k}$ follow standard normal distributions across individual respondents; and τ_{δ_i} and τ_{β_k} capture the impact of the latent addiction variable, α_n, a point we return to below.

For the price sensitivity, we relax the oft-assumed constant marginal utility of income imposed in many health choice models (Reed Johnson et al., 2011). Preferences for the cost attribute p, are treated as:

$$\beta_{np} = \left(\mu_{\beta_p} + \lambda_{\beta_p} z_n\right) \cdot \left(\frac{income_n}{\overline{income}}\right)^{\eta} \tag{5}$$

This drops the random heterogeneity and impact of the latent addiction variable, but adds in an income effect where η is an estimated income elasticity, $income_n$ is a given individual's income and $\overline{income}$ is the sample median.

As per figure 1, the latent variable for addiction is regressed on individual characteristics in the structural equation:

$$\alpha_n = \gamma z_n + \xi_n \tag{6}$$

where z_n are individual characteristics, γ captures the relationship between addiction and observed individual characteristics, and ξ_n is unobserved individual addiction heterogeneity, which follows a standard normal distribution.

In our specification, we allow for an impact of the latent addiction variable on the constants for the different alternatives as well as the parameters associated with individual attributes. This in essence means that our model is a standard hybrid choice model, albeit one where the latent variable relates to experience/addiction. We did test a model in which the entire utility function was addiction-conditioned, as in the original Hensher and Ho (2016) work, but this led to inferior results, largely as a function of the conditioning exerting itself on all sensitivities in the same direction, whether they relate to desirable or undesirable components.

With the elements of the utility function defined, we move to the specification of the choice model, which, given the assumption on the error term, takes the classic multinomial logit (MNL) form:

$$P_{C_n} = \prod_{t=1}^{T} \frac{e^{V_{ni^*t}}}{\sum_{j=1}^{J} e^{V_{nit}}} \tag{8}$$

where P_{C_n} is the probability of the observed sequence of choices for individual *n*, where i^* refers to the chosen alternative.

Next, we examine the relationship between the latent variable for addiction, α_n , and the indicator measures of addiction. The data for the indicator measures takes three broad forms. For each, its probability is modelled in a series of measurement equations. For the two binary variables (daily smoking, e-cigarette use), such as whether the individual smokes every day or not, a logit model is used:

$$P_{binary_n} = \prod_{k=1}^{2} \frac{\left(e^{(\delta_k + \zeta_k \alpha_n)}\right)^{I_{k_n}==1}}{1 + e^{(\delta_k + \zeta_k \alpha_n)}} \tag{9}$$

where δ_k are constant terms to be estimated, and ζ_k are estimated parameters capturing the relationship between the latent variable for addiction and the indicator at hand, I_{k_n}, where the exponent $I_{k_n} == 1$ ensures that the appropriate numerator is used depending on the observed value for I_{k_n}.

For five ordered variables (frequency of e-cigarette use, time before first cigarette of the day is smoked, time since last having smoked a cigarette that day, time since last having smoked a cigarette in the last few days/weeks, the number of quit attempts in the past year), an ordered logit model is used (Greene and Hensher, 2010):

$$P_{ordered_n} = \prod_{k=1}^{5} \left(\sum_{s=1}^{S} \delta(I_{k_n} == s) \left[\frac{e^{\tau_{k,s} - \zeta_k \alpha_n}}{1 + e^{\tau_{k,s} - \zeta_k \alpha_n}} - \frac{e^{\tau_{k,s-1} - \zeta_k \alpha_n}}{1 + e^{\tau_{k,s-1} - \zeta_k \alpha_n}} \right] \right) \tag{10}$$

where $\tau_{k,s}$ are estimated threshold parameters for threshold *s* of categorical indicator I_{k_n}, and ζ_k are estimated parameters capturing the relationship between the latent variable for addiction and the indicator at hand.

Finally, for two continuous variables (number of cigarettes smoked per day, current urge to smoke), a linear model is used:

$$P_{linear_n} = \prod_{k=1}^{2} \frac{1}{\sqrt{2\pi\sigma_k^2}} e^{-\frac{\left(I_{k_n} - \bar{I}_k - \zeta_k \alpha_n\right)^2}{2\sigma_k^2}} \tag{11}$$

where demeaning the variables, as shown by subtracting $\bar{I}_k$, avoids the need to estimate a constant (Daly et a., 2012b). The variance of the error, σ_k^2, is estimated along with other parameters. ζ_k are estimated parameters capturing the relationship between the latent variable for addiction and the indicator at hand.

Finally, each of the model components are combined into a single likelihood function and jointly estimated:

$$LL = \sum_{n=1}^{N} ln \int_\beta \int_\alpha P_c P_{binary} P_{ordered} P_{linear} \phi(\zeta) m(\beta|\Omega) d\beta d\alpha \tag{12}$$

We integrate over the parameter mixing distributions and latent variable, where maximum simulated likelihood is used for estimation (Train, 2009). Because we have more than 5 dimensions, standard Halton draws are rejected in favour of 1500 modified Latin hypercube sampling draws per random component and per individual, due to correlation concerns (cf. Bhat, 2003; Hess et al., 2006). Sobol draws would have also been a useful alternative to Halton draws (Czaikowski et al., 2019); the main point is to avoid using Halton draws. Given the known complexities of likelihood functions with mixing distributions, we use an algorithm for searching for starting values, based on Berliere et al. (2009)[3], to aid in finding a global optimum. All models are estimated using Apollo (Hess and Palma, 2019).

We optimized the specification with several rounds of removing interaction terms that were not statistically significant and updating the starting values. We followed the choice modelling literature in the choice of normalisation for alternative specific constants and categorical variables by deliberately over-specifying the model (attempting to estimate all parameters) and then omitting those with the lowest variance (Walker, 2001). On the basis of this, we normalised to zero the constant for the opt-out, the impact of the addiction latent variable on cigarettes, the low level of nicotine, and the impact of the latent variable on high nicotine.

[3] Procedure detailed in the Apollo manual (version 0.1.0), pp. 123-125

Finally, because we are using the model for forecasting, it is crucial to calibrate the model constants and the scale of utility. We follow the approach of Buckell and Hess (2019). More specifically, the ASCs are calibrated post-estimation using national data on tobacco product market shares. This aligns the base choice shares in our model to real-world market shares of the products. We also adopt the partial calibration developed in that paper; that is, the choice share of the opt-out in the uncalibrated model is retained and then the cigarette and e-cigarette choice shares are calibrated according to RP market shares. This avoids ascribing a specific behavioral interpretation to the opt-out (due to its framing, it could confer several behaviours). For calibrating the scale of utility, we use revealed preference (RP) data on respondents collected in the survey and build a RP choice model equivalent of equation (8). We estimate a common price coefficient in model (8) and its RP equivalent, and estimate an additional SP scale parameter, μ_{SP}. This aligns the scale of the SP model to that of the RP model. Terming the choice probability in model (8) $P_{c,SP}$ and its RP equivalent $P_{c,RP}$, the log-likelihood becomes:

$$LL = \sum_{n=1}^{N} ln \int_{\beta} \int_{\alpha} P_{c,SP} P_{c,RP} P_{binary} P_{ordered} P_{linear} \phi(\zeta) m(\beta|\Omega) d\beta d\alpha \qquad (13)$$

Limitations

Our methods are subject to a set of limitations. First, we are limited by the set of indicator measures that we collected. However, those that we do use cover important aspects of addiction: we have measures for cigarettes and e-cigarettes; we have longer and shorter term measures of addiction; we include measures of symptoms (craving) as well as outcome measures (cigarettes per day). Whilst it is common practice to run factor analyses prior to estimating latent variables, we discarded this preliminary exercise in this case. The factor analysis suggested that there were three latent variables. However, the directions of indicator variables across these latent variables were highly implausible. Moreover, attempting to estimate three separate latent variables would have added considerable complexity to the exercise and exposition. And, in fact, the one that we did estimate resulted in plausible directions of coefficients and implied choice behaviours. Using a single latent factor from a factor analysis was similarly doen in Shiffman et al. (2004). And other studies that conducted factor analysis yielded a single underlying driver of addiction (Strong et al., 2015; Strong et al., 2017). For these reasons, we kept to a single latent variable. As with all hybrid choice models, the latent variable should not be used in forecasts, due to the fact that it is only a cross-sectional measure (Chorus and Kroesen, 2014). Another limitation is the framing of the opt-out as "none of these". This could denote several different behaviours of respondents and ultimately we cannot observe precisely what these meant. However, we note that amongst smokers, other behaviours such as non-cigarette/non-e-cigarette use - which are likely responses when choosing the opt-out here - are much less popular among smokers, which is in keeping with what we find here. But this is at best speculation, and its interpretation is not clear. It is for this reason that we applied the partial calibration (Buckell and Hess, 2019), to avoid this issue impacting on forecasts and to allow for an open interpretation on the opt-out coefficient. It is also not an issue specific to our study. As we have previously suggested, it may be possible to allow different opt-out options with different labels, e.g. "I would rather try to quit" or "a cigar" to aid interpretation (Buckell and Hess, 2019). It would also be possible to use post-experimental questions to help interpret these choices (Reed Johnson et al., 2013). Finally, we were limited by processing power to using 1500 Halton draws in estimation. Whilst this is in excess of many other studies in health, we note recent research which suggests this may be too few draws (Czaikowski et al., 2019). Using 1500 draws was the maximum feasible number of draws permissible with the processing power available. This is likely to be common in applications where researchers do not have access to high-performance computing. In preliminary analyses, we used 500 draws. When we moved to 1500 draws, we did not see any notable difference in parameter estimates. So while we would have preferred more draws, we do not see 1500 as problematic.

Results

Measurement equations and structural equation

Table 4 shows the results from the measurement equations and the structural equation. This is informative in understanding the nature of addiction captured in the model. By studying the ζ parameters, we note that higher values of the latent variable are associated with a lower chance of being a daily smoker; smoking fewer cigarettes per day; a longer time in the morning before smoking their first cigarette; a longer time since last having smoked a cigarette that day; a longer time since last having smoked a cigarette in the last few days/weeks; more quit attempts in the past year; a higher probability of being an e-cigarette user; no increase in the frequency of e-cigarette use (not statistically significant); and reporting a lower level of current craving.

Taken together, these results suggest that a higher value for the latent variable is associated with lower levels of addiction to cigarettes and nicotine (or, equivalently, lower values of the latent variable are associated with higher levels of addiction to cigarettes and nicotine). This is due to, firstly, the lower levels of cigarette use and addiction. But it is of course possible that more addicted smokers could just be switching away from cigarettes to e-cigarettes. In this case, we would expect to see very little change in reported craving and increased frequency of use of e-cigarettes. However, we see the opposite: no increase in the frequency of e-cigarette use and lower reported craving. Moreover, higher values for the latent variable are associated with more past year quit attempts, which again is a sign of lower addiction. For these reasons, we interpret the latent variable as capturing reduced addiction to cigarettes and nicotine, where lower values signify higher levels of addiction. We then move to the structural equation to examine how individual characteristics of vary with addiction.

In the structural equation, individual characteristics are used to explain the latent variable. Higher values of the latent variable (i.e. less addicted) are associated with younger individuals (more likely to be younger and less likely to be older); higher education; higher income (less likely to be associated with lower income); non-white race/ethnicity (positive values for Hispanic, Black and Asian); and lower probability that a family member smokes.

This in keeping with what we would expect to see: these demographic patterns are opposite to those of smokers, which would fit with the idea of higher values for the latent variable capturing lower levels of addiction (Wang et al., 2018). (Though note that the definition of smokers varies between this study and that of Wang et al.; and that we are measuring addiction, whereas Wang et al. 2018 study current use.)

Utility function and impact of addiction latent variable

Estimates of the utility function are presented in table 5. All else being equal, the alternative specific constants indicate that cigarettes are preferred to the opt-out; and e-cigarettes are preferred to the opt-out. Unobserved preferences for cigarettes (relative to the opt-out) vary considerably around the mean, as reflected by the large and significant standard deviation. No preference heterogeneity for cigarettes across sociodemographic characteristics was found. Preferences for e-cigarettes (relative to the opt-out) varied in both unobserved ways (statistically significant estimated standard deviations) and according to age (older adults) as captured by the associated interaction terms of preferences and that characteristic.

For nicotine, medium level is preferred to all other levels (highest relative utility amongst the mean coefficients). Preferences for nicotine levels vary deterministically across some characteristics, but not all (i.e. associated interaction terms of nicotine levels and socio-demographic characteristics) and random ways for no nicotine and high nicotine (statistically significant estimated standard deviations). Our results indicate that female smokers prefer high and medium nicotine products, reflecting an overall preference for higher levels of nicotine in tobacco products. Older smokers chose high-nicotine products less often, reflecting an aversion to higher levels of nicotine in tobacco products; and e-cigarettes less often. Younger smokers chose high nicotine products more often. Unmarried smokers had preferences for medium and high nicotine products. (NB, of course these preferences are only part of the overall picture – characteristics are related to the latent variable, which also impacts on product and nicotine preferences.)

Price sensitivity, as expected, is negative. The income elasticity estimate implies that price sensitivities decline as income increases. In other words, as income increases, individuals are less sensitive to price changes. Tobacco is preferred to all other flavours. And healthier products are preferred to more harmful products, shown by stronger preferences for fewer life years lost.

Next, we move to the impact of the latent variable. From the interactions of the latent variable and product constant terms, lower addiction (higher values for the latent variable) is associated with increased preferences for e-cigarettes. Oppositely, this implies that those that are more addicted prefer cigarettes. With reduced addiction, the lower levels of nicotine are preferred to high nicotine; or, oppositely, more addicted smokers prefer higher levels of nicotine in cigarettes – and progressively so with higher levels of addiction, as captured by the monotonicity in the interactions of nicotine levels and the latent variable.

Willingness to pay (WTP) and willingness to accept (WTA)

Table 6 shows the estimated willingness to pay (WTP) (or willingness to accept, WTA, in the case of negative values) for the attributes in the utility function using the standard approach (Hensher et al., 2015). These reflect the dollar value per 20 pack (or e-cigarette equivalent) that individuals are willing to pay (must be compensated for, WTA) for that level of the attribute, relative to the omitted category, so as to remain indifferent (equal utility) between two products with these differing levels. For health harm, smokers, on average, are willing to pay $8.64 extra on a packet of 20 cigarettes to reduce the health risk from losing 10 years of life to losing 2 years of life. For flavours, smokers, on average, must be compensated to move to non-tobacco flavoured tobacco products to attain the same level of utility as for tobacco products. This differs across flavours. For menthol, WTA is $4.31; for fruit, WTA is $1.79; and for sweet, WTA is $2.32.

For nicotine, WTP/WTA is more modest on average than for other attributes (note nicotine valuations are also related to addiction, since the LV is interacted with these attributes in the utility function). Smokers express WTA for all levels of nicotine; that is, smokers must be compensated to achieve the same utility for moving from medium nicotine strength in tobacco products to any other strength of nicotine in tobacco products. More specifically, WTA at the mean for no nicotine is $0.57; for medium nicotine is $0.68; and for high nicotine is $0 since the estimated coefficient was 0.

From a policy perspective, it is of interest to examine how WTP/WTA for nicotine varies as a function of addiction[4]. The results indicate that key heterogeneity in WTP/WTA for nicotine by addiction is masked by analysis at the sample level. In Figure 3, how WTP/WTA for the levels of nicotine varies as a function of addiction is shown. We have used no nicotine as the reference category for these analyses (which is possible following from the fact that all of the coefficients are based on relative preferences). We have also reversed the scale of the latent variable so that increasing addiction is shown along the x-axis (for ease of reading). We see that increasing addiction is associated with higher WTP for nicotine in tobacco products (positive correlations for each of the attributes). We also see that WTP increases monotonically for higher levels of nicotine (steeper gradients across levels). In some cases the value of WTP/WTA is in excess of average the price of a packet of 20 cigarettes (around $8), underlining the importance of nicotine to smokers.

Forecasting of lowering levels of nicotine in cigarettes

A key policy issue is the extent to which smokers would switch away from cigarettes if nicotine levels were reduced in cigarettes. Table 7 shows two models' forecasts of lowering nicotine in cigarettes[5]. Our preferred specification is the calibrated model. Here, we see that the model predicts that lowering nicotine levels in cigarettes would result in around a 3% decline in the choice share of cigarettes; and 4% and 11% increases, respectively, in choice shares for e-cigarettes and the opt-out option. The table also indicates that smokers are slightly less responsive to the reduction in nicotine levels than dual users, because the reduction in cigarette choice share for smokers is less than

---

[4] For this, a slightly more involved computation for WTP is used. Taking low nicotine as an example, WTP/WTA at the mean is computed as, $WTP_{lownicotine} = \frac{\mu_{\beta_{lownicotine}} + \lambda_{\beta_{lownicotine,unmarried}} unmarried_n + \tau_{\beta_{lownicotine}} \alpha_n}{\hat{\mu}_{\beta_{price}} \cdot \left(\frac{income_n}{\overline{income}}\right)^{\eta}}$. In figure 3, this value is plotted for each draw from the mixing distributions.

[5] NB – we used "low" rather than "no" nicotine to make these predictions. This is because the FDA's stated position is to lower nicotine to "non-addicting levels" and, as such, cigarettes would still contain some (albeit very little) nicotine. Therefore, we think that low nicotine makes for more realistic forecasts.

that for dual users. However, this difference is fairly modest and smokers appear at least somewhat responsive to the lowering of nicotine. Of course, these should be considered as short-term responses in demand; longer term forecasts are not possible with these data.

Diagnostics

Table 3 shows the diagnostic information. The log-likelihood of the joint estimation is shown, along with, for comparison, the log-likelihood of a model where all products are chosen equally (i.e. no information on choices) and the log-likelihood of a model with only the experimental choice shares (Mokhtarian, 2016).. Measures of fit, AIC and BIC are presented along with the total number of estimated parameters, 76.

Discussion and conclusions

In this paper, we developed a model to evaluate the role of addiction in smokers' choice behaviour. The hybrid choice model seeks to better understand smokers' decision-making by allowing addiction to flexibly impact on smokers' choices. This framework allows us to overcome a set of empirical issues that present in trying to measure addiction and to incorporate addiction in choice models in a traditional manner. We used the model to estimate preferences, willingness to pay for nicotine in cigarettes, and to predict the impact of lowering nicotine levels in cigarettes in the US.

We find that the latent variable captures addiction, with higher values explaining lower levels of addiction and lower values capturing higher levels of addiction. Higher levels of addiction are associated with increased use of cigarettes (including more cigarettes smoked per day), lower use for e-cigarettes, fewer quit attempts, and higher levels of reported craving. This addiction was associated with sociodemographic characteristics in a way that corresponds to known, observed patterns of cigarette use nationwide.

In terms of the impact of addiction on stated choices, we find that addiction drives preferences for cigarettes, and away from e-cigarettes. Those that are more addicted prefer higher levels of nicotine in tobacco products.

Analysis of WTP indicates that, on average, smokers prefer medium levels of nicotine, and on average are willing to pay or must be compensated in the range of $0.42-0.57 per 20 pack to be equally satisfied with levels of nicotine other than "low", which was the reference in our model. However, these valuations vary considerably across respondents when the range of addiction is taken into account. More addicted smokers exhibit utility for higher levels of nicotine in tobacco products where WTP is, in some cases, in excess of the price of a packet of 20 cigarettes; less addicted smokers value lower levels of nicotine in tobacco products where WTA is, in some cases, in excess of the price of a packet of 20 cigarettes.

Our results suggest that the short-run response to the FDA's proposed lowering of nicotine in cigarettes would result in a slight shift away from cigarettes; roughly 3% of its choice share. Shifts in choice shares would be towards both e-cigarettes (4% increase in its choice share); and a 11% increase in the choice share of the opt-out - either cessation behaviour or alternative tobacco products (depending on one's interpretation of the opt-out in the experiment - we have previously interpreted as the former based on higher choices of the opt out option among those that attempted to quit in the past year; see Buckell et al., 2019).

These results are likely to have significant meaning for policy. In the US, the FDA has set out its regulatory agenda, the centrepiece of which is reducing the level of nicotine in cigarettes (FDA, 2019). Therefore, these findings are likely to be of direct relevance to current policymaking. The results suggest that this policy is likely to be effective at shifting smokers' choices away from cigarettes; though with limited impact in the short run (NB – we can make no determination on the medium- to long-term impacts of this policy). Whilst there seems to be more of a response to this policy from the dual users, smokers, too, showed some switching away from cigarettes. Since the smokers are most likely at harm, this is encouraging for the public health implications of this policy.

These results are important because, from a behavioural and policy perspective, we have greater insight into smokers' decision-making processes with respect to nicotine preferences and product choices. Our basic utility function parameter estimates are in keeping with previous results elsewhere in the literature (Pesko et al., 2016; Marti et al., 2018; Buckell et al., 2018; Shang et al., 2018; Shang et al., 2018). Here, where measured, the preference estimates for nicotine are typically lower than other attributes. However, these studies do not explicitly study the impact of addiction. And now that we do, the results appear to be markedly different. Thus, the key point is that, even when nicotine is used in choice experiments, the behaviour of smokers towards products and nicotine is likely to be underestimated if addiction is not explicitly modelled.

References

Abou-Zeid, M. , Ben-Akiva, M. , 2014. Hybrid choice models. In: Hess, S., Daly, A. (Eds.), Handbook of Choice Modelling. Edward Elgar, Cheltenham .

Balbontin, C., Hensher, D. A., & Collins, A. T. (2019). How to better represent preferences in choice models: The contributions to preference heterogeneity attributable to the presence of process heterogeneity. Transportation Research Part B: Methodological, 122, 218-248.

Bierlaire M, Thémans M, Zufferey N. A Heuristic for Nonlinear Global Optimization. INFORMS Journal on Computing. 2009;22(1):59-70.

Bhat CR. Simulation estimation of mixed discrete choice models using randomized and scrambled Halton sequences. Transportation Research Part B: Methodological. 2003;37(9):837-55.

Bold, K. W., Sussman, S., O'Malley, S. S., Grana, R., Foulds, J., Fishbein, H., & Krishnan-Sarin, S. (2018). Measuring E-cigarette dependence: Initial guidance. Addictive Behaviors, *79*, 213-218.

Buckell, J. and S. Hess (2019). "Stubbing out hypothetical bias: improving tobacco market predictions by combining stated and revealed preference data." Journal of Health Economics **65**: 93-102.

Buckell, J., Marti, J., & Sindelar, J. L. (2019). Should flavours be banned in cigarettes and e-cigarettes? Evidence on adult smokers and recent quitters from a discrete choice experiment. Tobacco Control, *28*(2), 168.

Buckell, J. and J. L. Sindelar (2019). "The impact of flavors, health risks, secondhand smoke and prices on young adults' cigarette and e-cigarette choices: a discrete choice experiment." Addiction **0**(0).

Cawley, J. and C. J. Ruhm (2011). Chapter Three - The Economics of Risky Health Behaviors. Handbook of Health Economics. M. V. Pauly, T. G. McGuire and P. P. Barros, Elsevier. **2:** 95-199.

Chaloupka, F. J. and K. E. Warner (2000). Chapter 29 The economics of smoking. Handbook of Health Economics, Elsevier. **1:** 1539-1627.

Choice Metrics. 2018. Ngene 1.2 USER MANUAL & REFERENCE GUIDE. http://www.choice-metrics.com/NgeneManual120.pdf Accessed 05/15/2020

Chorus, C. G. and M. Kroesen (2014). "On the (im-)possibility of deriving transport policy implications from hybrid choice models." Transport Policy **36**: 217-222.
Czajkowski, M. and W. Budziński (2019). "Simulation error in maximum likelihood estimation of discrete choice models." Journal of Choice Modelling **31**: 73-85.

Collins AC, Marks MJ. Progress towards the development of animal models of smoking-related behaviors. Journal of addictive diseases. 1991;10(1-2):109-26.

Daly, A., Dekker, T., & Hess, S. (2016). Dummy coding vs effects coding for categorical variables: Clarifications and extensions. Journal of Choice Modelling, *21*, 36-41.

Daly, A., Hess, S., & de Jong, G. (2012a). Calculating errors for measures derived from choice modelling estimates. Transportation Research Part B: Methodological, *46*(2), 333-341.

Daly, A., Hess, S., Patruni, B., Potoglou, D., & Rohr, C. (2012b). Using ordered attitudinal indicators in a latent variable choice model: a study of the impact of security on rail travel behaviour. Transportation, *39*(2), 267-297.

de Bekker-Grob, E. W., Donkers, B., Jonker, M. F., & Stolk, E. A. (2015). Sample Size Requirements for Discrete-Choice Experiments in Healthcare: a Practical Guide. The Patient - Patient-Centered Outcomes Research, *8*(5), 373-384. Fayers,

DeCicca P, Kenkel DS, Lovenheim MF. The Economics of Tobacco Regulation: A Comprehensive Review. National Bureau of Economic Research Working Paper Series. 2020;No. 26923.

P. M. and D. J. Hand (2002). "Causal variables, indicator variables and measurement scales: an example from quality of life." Journal of the Royal Statistical Society: Series A (Statistics in Society) **165**(2): 233-253.

Greene, W.H. and Hensher, D.A. (2010) Modeling Ordered Choices: A Primer and Recent Developments, Cambridge University Press, Cambridge.

Heatherton TF, Kozlowski LT, Frecker RC, Fagerstrom KO (1991). The Fagerstrom Test for Nicotine Dependence: a revision of the Fagerstrom Tolerance Questionnaire. Br J Addict 86:1119-27.

Hensher, D. A. (2015). "Data Challenges: More Behavioural and (Relatively) Less Statistical – A Think Piece." Transportation Research Procedia **11**: 19-31.

Hensher, D. A. and C. Q. Ho (2016). "Experience conditioning in commuter modal choice modelling – Does it make a difference?" Transportation Research Part E: Logistics and Transportation Review **95**: 164-176.

Hensher, D., Balbontin, C., Greene, W. 2019. Experience as a conditioning effect on choice – Does it matter whether it is exogenous or endogenous? Working paper, University of Sydney

Hess S, Train KE, Polak JW. On the use of a Modified Latin Hypercube Sampling (MLHS) method in the estimation of a Mixed Logit Model for vehicle choice. Transportation Research Part B: Methodological. 2006;40(2):147-63.

Hess, S., Hensher, D. A., & Daly, A. (2012). Not bored yet – Revisiting respondent fatigue in stated choice experiments. Transportation Research Part A: Policy and Practice, *46*(3), 626-644.

Hess, S., Spitz, G., Bradley, M., & Coogan, M. (2018). Analysis of mode choice for intercity travel: Application of a hybrid choice model to two distinct US corridors. Transportation Research Part A: Policy and Practice, *116*, 547-567.

Hess, S. and D. Palma (2019). "Apollo: A flexible, powerful and customisable freeware package for choice model estimation and application." Journal of Choice Modelling **32**: 100170.

Ida, T. and R. Goto (2009). "Simultaneous measurement of time and risk preferences: stated preference discrete choice modeling analysis depending on smoking behavior*." International Economic Review **50**(4): 1169-1182.

Kenkel, D., Peng, S., Pesko, M., Wang, H., 2017. Mostly Harmless Regulation? Electronic Cigarettes, Public Policy And Consumer Welfare. National Bureau of Economic Research Working Paper Series, No 23710

Louviere, J.J., Hensher, D.A., Swait, J.D., 2000. Stated Choice Methods: Analysis and Applications. Cambridge, UK.

Marti, J., Buckell, J., Maclean, J. C., & Sindelar, J. (2019). TO "VAPE" OR SMOKE? EXPERIMENTAL EVIDENCE ON ADULT SMOKERS. Economic Inquiry, *57*(1), 705-725.

Mokhtarian PL. Discrete choice models' ρ2: A reintroduction to an old friend. Journal of Choice Modelling. 2016;21:60-5.

Pesko, M. F., Kenkel, D. S., Wang, H., & Hughes, J. M. (2016). The effect of potential electronic nicotine delivery system regulations on nicotine product selection. Addiction, *111*(4), 734-744.

Reed Johnson, F., Mohamed, A. F., Özdemir, S., Marshall, D. A., & Phillips, K. A. (2011). How does cost matter in health-care discrete-choice experiments? Health Economics, *20*(3), 323-330.

Reed Johnson F, Lancsar E, Marshall D, Kilambi V, Mühlbacher A, Regier DA, et al. Constructing Experimental Designs for Discrete-Choice Experiments: Report of the ISPOR Conjoint Analysis Experimental Design Good Research Practices Task Force. Value in Health. 2013;16(1):3-13

Regmi, K., Kaphle, D., Timilsina, S., & Tuha, N. A. A. (2018). Application of Discrete-Choice Experiment Methods in Tobacco Control: A Systematic Review. PharmacoEconomics - Open, *2*(1), 5-17.

Shadel WG, Shiffman S, Niaura R, Nichter M, Abrams DB. Current models of nicotine dependence: what is known and what is needed to advance understanding of tobacco etiology among youth. Drug and alcohol dependence. 2000;59:9-22.

Shang, C., Huang, J., Chaloupka, F. J., & Emery, S. L. (2018). The impact of flavour, device type and warning messages on youth preferences for electronic nicotine delivery systems: evidence from an online discrete choice experiment. Tobacco Control, *27*(e2), e152.

Shi, Y., Cao, Y., Shang, C., & Pacula, R. L. (2019). The impacts of potency, warning messages, and price on preferences for Cannabis flower products. International Journal of Drug Policy, *74*, 1-10.

Shiffman S, Sembower MA. Dependence on e-cigarettes and cigarettes in a cross-sectional study of US adults. Addiction. 2020;

Shiffman S, Waters AJ, Hickcox M. The Nicotine Dependence Syndrome Scale: A multidimensional measure of nicotine dependence. Nicotine & Tobacco Research. 2004;6(2):327-48.

Strong DR, Pearson J, Ehlke S, Kirchner T, Abrams D, Taylor K, et al. Indicators of dependence for different types of tobacco product users: Descriptive findings from Wave 1 (2013–2014) of the Population Assessment of Tobacco and Health (PATH) study. Drug and alcohol dependence. 2017;178:257-66.

Strong DR, Pearson J, Ehlke S, Kirchner T, Abrams D, Taylor K, et al. Indicators of dependence for different types of tobacco product users: Descriptive findings from Wave 1 (2013–2014) of the Population Assessment of Tobacco and Health (PATH) study. Drug and alcohol dependence. 2017;178:257-66.

Train, K., 2009. Discrete Choice Methods with Simulation. Cambridge University Press, Cambridge.

Vij, A. and J. L. Walker (2016). "How, when and why integrated choice and latent variable models are latently useful." Transportation Research Part B: Methodological **90**: 192-217.

Walker, J. 2001. Extended Discrete Choice Models: Integrated Framework, Flexible Error Structures, and Latent Variables. PhD Thesis, Massachusetts Institute of Technology.

Wang TW, Asman K, Gentzke AS, Cullen, K., Holder-Hayes, E., Reyes-Guzman, C., Jamal, A., Neff, L. & King, B. Tobacco Product Use Among Adults — United States, 2017. MMWR Morb Mortal Wkly Rep 2018;67:1225–1232. DOI: http://dx.doi.org/10.15585/mmwr.mm6744a2

Wehbe, L., Ubhi, H. K., & West, R. (2018). Want, need and habit as drivers of smoking behaviour: A preliminary analysis. Addictive Behaviors, *76*, 135-138.

West, R. (2009). "The Multiple Facets of Cigarette Addiction and What They Mean for Encouraging and Helping Smokers to Stop." COPD: Journal of Chronic Obstructive Pulmonary Disease **6**(4): 277-283.

West, R. and J. Brown (2013). Theory of addiction. Chichester, John Wiley & Sons.

	E-cigarette	Cigarette
Flavour	Plain tobacco Menthol Fruit Sweet	Plain tobacco Menthol
Life years lost by average user	10 5 2 Unknown	10
Level of nicotine	High Medium Low None	High Medium Low
Price	\$4.99 \$7.99 \$10.99 \$13.99	\$4.99 \$7.99 \$10.99 \$13.99

Table 1: Experimental Design

	Mean	Standard deviation	Min	Max
Age	41.09	12.55	18	64
Female	0.55	0.50	0	1
Higher education	0.41	0.49	0	1
Income	52974.67	34199.27	15000	150000
Hispanic	0.08	0.26	0	1
White	0.86	0.35	0	1
Black	0.10	0.30	0	1
Asian	0.03	0.16	0	1
Employed	0.60	0.49	0	1
Married	0.36	0.48	0	1
Family member smokes	0.44	0.50	0	1
daily smoker	0.95	0.22	0	1
number of cigarettes smoked per day	14.61	9.34	1	61
time before first cigarette of the day is smoked: less than 5 minutes	0.32	0.47	0	1
time before first cigarette of the day is smoked: 5 to 30 minutes	0.46	0.50	0	1
time before first cigarette of the day is smoked: 31 to 60 minutes	0.12	0.33	0	1
time before first cigarette of the day is smoked: longer than 60 minutes	0.10	0.31	0	1
time since last having smoked a cigarette that day: less than 15 minutes	0.31	0.46	0	1
time since last having smoked a cigarette that day: 15 to 30 minutes	0.27	0.44	0	1
time since last having smoked a cigarette that day: 30 to 60 minutes	0.19	0.39	0	1
time since last having smoked a cigarette that day: 1 to 3 hours	0.11	0.31	0	1
time since last having smoked a cigarette that day: 3 to 6 hours	0.02	0.15	0	1
time since last having smoked a cigarette that day: more than 6 hours	0.02	0.15	0	1
time since last having smoked a cigarette in the last few days/weeks: today	0.90	0.30	0	1
time since last having smoked a cigarette in the last few days/weeks: yesterday	0.07	0.26	0	1
time since last having smoked a cigarette in the last few days/weeks: in the last week	0.02	0.14	0	1
time since last having smoked a cigarette in the last few days/weeks: in the last month	0.00	0.06	0	1
time since last having smoked a cigarette in the last few days/weeks: longer than a month	0.00	0.05	0	1
quit attempts in the past year: none	0.53	0.50	0	1
quit attempts in the past year: 1	0.32	0.47	0	1
quit attempts in the past year: 2 to 3	0.11	0.31	0	1
quit attempts in the past year: more than 4	0.03	0.18	0	1
e-cigarette use	0.36	0.48	0	1
frequency of e-cigarette use: daily	0.12	0.32	0	1
frequency of e-cigarette use: several times per week	0.14	0.35	0	1
frequency of e-cigarette use: once per week	0.04	0.20	0	1
frequency of e-cigarette use: less than once per week	0.05	0.23	0	1
current urge to smoke (i.e. craving)	5.19	2.49	1	10

Table 2: Descriptive Statistics

LL(0)	-47460.2
LL(choice shares)	-46394.4
LL(final, whole model)	-40302.4
AIC	80746.47
BIC	81346.62
Estimated parameters	76

Table 3: Model diagnostic information. LL(0) – log-likelihood where parameters are zero, i.e. there is no information on preferences and the choices of all products are equally likely. LL(choice share) – log-likelihood with only product constants, i.e. recovering the experimental choices shares.

	Estimate	t-ratio (0)
Structural equation		
young	0.65	4.40
older	-0.24	-3.02
higher education	0.34	5.45
low income	-0.29	-3.97
Hispanic	0.49	4.11
Black	0.58	6.75
Asian	0.57	3.68
family member smokes	-0.27	-5.06
Measurement equations		
Daily smoker	-1.99	-5.03
Cigarettes per day	-6.46	-17.32
Time to first cigarette	1.47	10.44
Last cigarette	1.86	6.48
Long ago last cigarette	0.41	5.43
Quit attempts	0.29	4.28
E-cigarette user	0.31	2.76
Frequency of e-cigarette use	0.20	1.76
Current craving	-0.29	-4.29

Table 4: estimates from the measurement equations and the structural equation of the latent variable. In the structural equation, the parameters are the estimated gammas as in eqn (6). In the measurement equations, the zeta parameters are each taken from separate measurement equations, as in (9)-(11).

	Estimate	t-ratio (0)
Utility function		
Cigarette ASC	2.98	4.92
Cigarette s.d.	2.29	4.88
E-cigarette ASC	1.18	4.40
E-cigarette s.d.	2.38	4.89
E-cigarette * older	-0.95	-3.55
No nicotine	-0.07	-1.40
No nicotine s.d.	0.58	4.19
Medium nicotine	0.05	1.18
Medium nicotine * female	0.13	2.04
Medium nicotine * older	-0.26	-2.99
Medium nicotine * unmarried	0.14	2.11
High nicotine s.d.	0.63	4.78
High nicotine * female	0.18	2.69
High nicotine * young	0.30	2.26
High nicotine * older	-0.37	-3.49
High nicotine * unmarried	0.18	2.37
Price	-0.13	-5.07
Lambda income	-0.28	-5.39
Menthol	-0.54	-4.56
Fruit	-0.22	-3.30
Sweet	-0.29	-3.73
Unknown health harm	0.83	4.54
2 years of life lost	1.08	4.61
5 years of life lost	0.45	3.80
mu_SP	1.02	4.95
Dual user	-1.53	-7.99
Addiction-utility function interactions		
Tau * E-cigarette	0.26	2.01
Tau * No and low nicotine	0.37	3.88
Tau * Medium nicotine	0.14	3.55

Table 5: Utility function and addiction interactions in the hybrid choice model. ASC – Alternative-specific constant (mean of the mixing distribution); s.d. – standard deviation of the mixing distribution; mu_SP – scale parameter for SP relative to RP (t-ratio vs 1 = , meaning that differences in scale between SP and RP were not statistically significant, which follows from the coefficient being close to 1); lambda income is the income elasticity, its coefficient reflects that those on higher incomes are less responsive to price variation as economic theory predicts.

variable	WTP	t-ratio (0)	rob s.e.	LCB, 95%	UCB, 95%
No nicotine	-0.57	-1.47	0.39	-1.32	0.18
Medium nicotine	0.42	1.12	0.35	-0.26	1.10
Menthol	-4.31	-10.31	0.42	-5.13	-3.49
Fruit	-1.79	-4.38	0.41	-2.59	-0.99
Sweet	-2.32	-5.59	0.41	-3.13	-1.51
year2	8.64	12.67	0.68	7.30	9.98
year5	3.56	6.23	0.57	2.44	4.68
Unknown	6.62	11.03	0.60	5.44	7.80

Table 6: Estimates of willingness to pay (WTP), $ per 20-pack of cigarettes, for the hybrid choice model. Standard errors are computed using the delta method (see Daly et al., 2012a).

	Cigarette Choice share	E-cigarette Choice share	Optout Choice share
Base choice shares, uncalibrated	0.56	0.33	0.10
	(0.557; 0.562)	(0.333; 0.336)	(0.103; 0.106)
Low nicotine in cigarettes, uncalibrated	0.54	0.35	0.11
	(0.535; 0.541)	(0.344; 0.348)	(0.113; 0.116)
% change in choice shares, uncalibrated	-3.86%	3.50%	9.49%
Base choice shares, calibrated	0.68	0.22	0.10
	(0.671; 0.682)	(09.214; 0.224)	(0.103; 0.106)
Low nicotine in cigarettes, calibrated	0.65	0.23	0.12
	(0.648; 0.659)	(0.224; 0.233)	(0.114; 0.118)
% change in choice shares, calibrated	-3.28%	4.17%	10.68%
% change in calibrated choice shares, if smokes only	-3.04%	4.01%	11.47%
% change in calibrated choice shares, if dual user	-3.23%	4.53%	12.36%

Table 7: predicted choice shares and changes in choice shares from lowering nicotine in cigarettes. 95% confidence intervals are listed in parentheses beneath the product forecasts. Calibration follows previous work (Buckell and Hess, 2019). In the table, *smokes only* refers to those in the sample that report no e-cigarette use; *dual user* refers to those in the sample that report e-cigarette use.

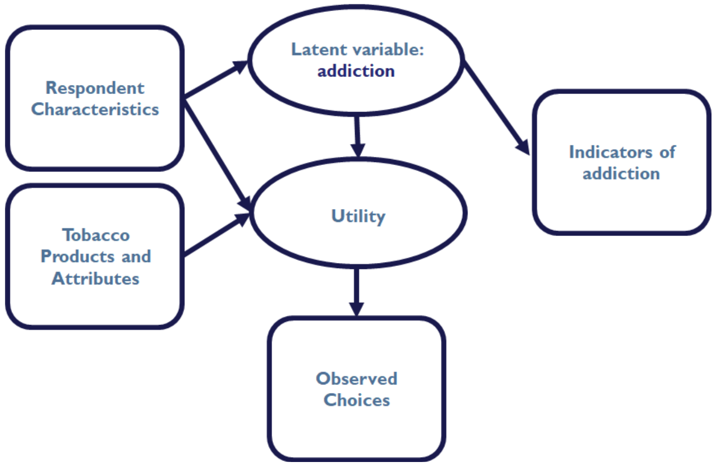

Fig 1. Schematic of hybrid choice model for addiction. Square boxes are observed variables; ellipses are unobserved variables. NB – "Tobacco Products and Attributes", particularly nicotine, could of course impact on addiction. But these will be longer-term effects, rather than short term choices. In this setup, the attributes impact on short term choices via utility; longer term effects impact on addiction and addiction, in turn, influences utility and choices.

Fig. 2: Sample choice scenario. Four options were presented to each respondent in each choice set, along with an opt-out option, “none of these”.

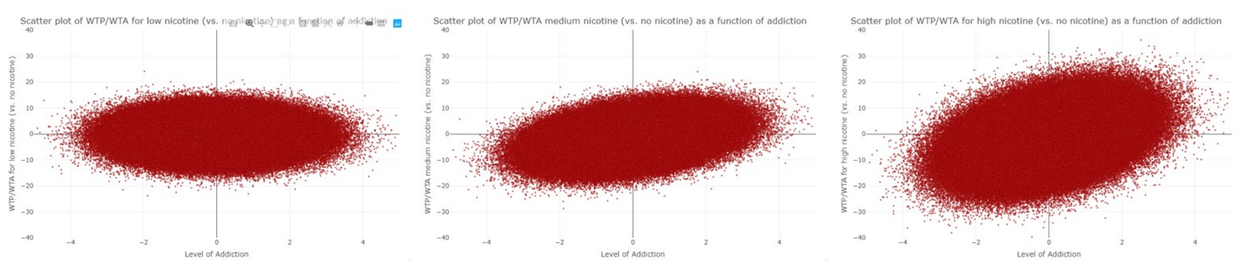

Fig. 3: Willingness to pay/willingness to accept (WTP/WTA) for nicotine as a function of addiction. (i) WTP/WTA for low nicotine (reference: no nicotine) as a function of addiction; (ii) WTP/WTA for medium nicotine (reference: no nicotine) as a function of addiction; (iii) WTP/WTA for high nicotine (reference: no nicotine) as a function of addiction. Addiction is defined as the latent variable (for ease of interpretation we have reversed the scale so that higher score on the latent variable denotes higher addiction). Each point is taken from a draw from the mixing distributions of the parameters and latent variable. NB – WTP for low vs no nicotine will be zero given the parameter restrictions on the model; we show the relationship as there is important random heterogeneity that remains, as can be seen. Overall, WTP increases with addiction for the same level of the attribute (positive correlation) and increase monotonically with increasing levels (progressively steeper gradients).