Proceedings of the 2024 Winter Simulation Conference
H. Lam, E. Azar, D. Batur, S. Gao, W. Xie, S. R. Hunter, and M. D. Rossetti, eds.

INCORPORATING THE COM-B MODEL FOR BEHAVIOR CHANGE INTO AN AGENT-BASED MODEL OF SMOKING BEHAVIORS: AN OBJECT-ORIENTED DESIGN

David Tian[1], Hazel Y. Squires[2], Charlotte Buckley[1], Duncan Gillespie[2], Harry Tattan-Birch[3], Lion Shahab[3], Robert West[3], Alan Brennan[2], Jamie Brown[3], and Robin C. Purshouse[1]

[1]Dept. of Automatic Control and Systems Engineering, University of Sheffield, S1 3JD, UK
[2]Sheffield Centre for Health and Related Research, School of Medicine and Population Health, University of Sheffield, S1 4DA, UK
[3]Institute of Epidemiology and Health, University College London, WC1E 7HB, UK

ABSTRACT

Modeling trajectories in cigarette smoking prevalence, initiation and quitting for populations and subgroups of populations is important for policy planning and evaluation. This paper proposes an agent-based model (ABM) design for simulating the smoking behaviors of a population using the Capability, Opportunity, Motivation - Behavior (COM-B) model. Capability, Opportunity and Motivation are modeled as latent composite attributes which are composed of observable factors associated with smoking behaviors. Three forms of the COM-B model are proposed to explain the transitions between smoking behaviors: initiating regular smoking uptake, making a quit attempt and quitting successfully. The ABM design follows object-oriented principles and extends an existing generic software architecture for mechanism-based modeling. The potential of the model to assess the impact of smoking policies is illustrated and discussed.

1 INTRODUCTION

Tobacco smoking remains a leading cause of morbidity and mortality worldwide (Reitsma et al. 2021). In England, smoking prevalence has been declining since around 2000, driven by a range of tobacco control measures including: price and tax measures; bans on tobacco advertising, promotion and sponsorship; smokefree policies; health warnings on tobacco packages; monitoring and surveillance; and measures to combat illicit trade in tobacco products (Beard et al. 2019). Novel tobacco control policies will be needed to reach the UK government's target of a smokefree England by 2030 (defined as smoking prevalence below 5%), especially among priority subgroups of the population. This is true for many other countries in a similar stage of the tobacco epidemic. These policies include mass media campaigns, increasing the age of sale for buying tobacco products, improving access to stop smoking services, and smartphone app interventions (West 2017; Naughton et al. 2023).

1.1 Agent-based and Individual-level Modeling of Smoking Behavior and Policy

Agent-based modeling (ABM) has previously been used to model smoking behaviors—see Zhong et al. (2023) for a review. Existing models have considered how social influence mechanisms can explain trends in initiation (Huang et al. 2021) and cessation (Sukthankar and Beheshti 2019). Meanwhile, the 'Tobacco Town' family of policy models have considered the impact of retail availability restrictions on tobacco purchasing decisions (Luke et al. 2017). However, none of these existing models represent a sufficiently rich set of mechanisms to appraise a broader range of tobacco control policies, including those targeting individual behavior change and legislative changes, such as those being considered in England.

The Sheffield Tobacco Policy Model (STPM) is an individual-level simulation model that has been used to inform smoking policy making in England (Gillespie and Brennan 2023). STPM is capable of generating population trends in smoking prevalence by simulating the life-course trajectories of smoking for successive birth-cohorts of individuals. In STPM, a population is divided into numerous subgroups, defined by combinations of age, sex, and quintiles of the Index of Multiple Deprivation—a small-scale geographic indicator of socio-economic conditions. For each subgroup, STPM simulates trends in smoking rates as a function of trends in the annual state transition probabilities of smoking initiation, quitting or relapse. However, STPM, like other microsimulations, is unable to explain how new interventions and policies would affect individual-level state transition probabilities since it lacks theory-based behavioral mechanisms that link to policy. Thus predicting the impact of the interventions over the longer term has previously been based on limited evidence or theory.

1.2 Theory-informed Approach to Agent-based Modeling

A number of ABM studies have used a theory-informed approach to model design—see Antosz et al. (2023) for a recent review. In this work, we leverage the Mechanism-Based Social Systems Modeling (MBSSM) architecture (Vu et al. 2020). MBSSM exploits the macro-micro-macro conceptual framework for social mechanisms proposed by Hedstrom and Swedberg (1996) to incorporate multiple mechanisms of behavioral theories within ABMs of social phenomena such as addictive behaviors. The MBSSM architecture consists of macro and micro entities, which allow changes to be modeled at the population/structural and individual/agency levels respectively. The following three mechanism types describe the relationships within and between the macro and micro entities:

- *Situational mechanisms* in which the micro-level traits of an agent are modified depending on their macro-level situational context, e.g., individuals responding to the marketing of a new product.
- *Action mechanisms* in which the behaviors of an agent are determined by their micro-level traits, e.g., individuals acting on new information.
- *Transformational mechanisms* in which the collective behaviors of individuals at the micro-level are used to update the macro-level context, e.g., a new fashion among some individuals changing the generally perceived culture in the population.

MBSSM also allows for macro-macro mechanisms, in which direct relationships between changes to contexts at the macro-level are modeled without any abstraction of micro-level generative mechanisms, e.g., the marketing of a new product directly changing the perceived culture in the population.

Vu et al. (2020) developed an object-oriented software architecture for MBSSM which can be implemented in any ABM software development tool that supports object orientation; implementations exist for Repast HPC (Collier and North 2013) and Repast4Py (Collier and Ozik 2022). MBSSM has previously been used to develop theory-informed ABMs for alcohol use behaviors using social norm theory, social role theory, the theory of planned behavior, and social contagion theory—see, e.g., Vu et al. (2023). However, in present work, we consider how the COM-B model of behavior change (Michie et al. 2011) can be used to design an ABM of smoking behaviors. The COM-B model represents the observation that behavior at any given moment will occur only when an individual has the capability (C) and opportunity (O) to engage in the behavior and is motivated (M) to enact that behavior over alternatives. Capability refers to an individual's psychological and physical capacity to engage in a particular behavior. Opportunity refers to the physical or social environment with which people interact. Motivation refers to the mental processes that energize and direct behavior. Motivation can be categorized into reflective motivation (e.g., conscious decision-making and inferences) and automatic motivation (e.g., feelings and habits). COM-B provides a coherent framework to organize and describe the most important influences on smoking and smoking cessation, e.g., (West and Brown 2013), and has been widely used for this purpose, e.g., (Gilbert 2023).

To our knowledge, the only existing ABM that has incorporated the COM-B model is by Atkinson et al. (2018). They used COM-B as a framework for developing a set of rules that influence the likelihood that an agent will consume alcohol during their activities of daily life (with motivation to consume updated hourly in the simulation). However, the design is closely tied to its alcohol use application and the source code is not in the public domain, making reuse for other purposes challenging.

1.3 Aims and Overview of the Paper

This paper proposes a design for representing COM-B in an ABM, grounded in a case study of simulating smoking behaviors for the purposes of tobacco policy appraisal. Three forms of the COM-B model are proposed for the smoking behaviors: (1) the initiation of regular smoking; (2) making a quit attempt; and (3) quitting successfully. The paper is organized as follows. Section 2 describes how the proposed ABM models the dynamics of smoking behavior. Section 3 presents the software architecture of the proposed ABM and an example implementation of the ABM using Repast4Py. Section 4 presents the form of the COM-B model for the three smoking behaviors. Section 5 discusses the utility of the model and concludes.

2 THE DYNAMICS OF SMOKING BEHAVIOR

2.1 Agents and Behaviors

Each agent represents an individual person. The smoking status of the agent is defined in Table 1 using five mutually exclusive states, based on definitions in the Smoking Toolkit Study survey of smoking and smoking cessation in England (Fidler et al. 2011): never smoker, smoker, new quitter, ongoing quitter and ex-smoker. The ongoing quitter state is a tunnel state that encodes the memory of how long a quit attempt has lasted for, encoded as 'ongoing quitter i', where $i = 1, \ldots, 11$ is a count of the number of months of maintained abstinence from smoking. After 12 months of abstinence have been maintained, the agent transitions to being an ex-smoker.

COM-B mechanisms are used to explain all transitions between states, with the exception of the transition from ex-smoker to smoker (i.e., long-term relapse to regular smoking). In this case, there is insufficient evidence to develop a COM-B explanation and instead we resort to using purely data-driven transition probabilities from the STPM model. Note that the majority of relapses occur within a year of beginning a quit attempt (Gillespie and Brennan 2023).

Table 1: The smoking states of an agent.

States	Description
never smoker	a person who has never smoked regularly (i.e., smoked for a year or more)
smoker	a person who is smoking regularly and did not begin a quit attempt in the past month
new quitter	a person who was smoking regularly and began a quit attempt in the past month
ongoing quitter	a person who has successfully maintained smoking abstinence for at least one month and for less than a year
ex-smoker	a person who smoked regularly more than a year ago

2.2 The State Transition Model

The ABM simulates the smoking behaviors of a 16+ years old synthetic population at each time step (tick), t, based on an individual-level state transition model (Figure 1). Each tick represents one month. The state transition determines an agent's state at the next tick based on sampling a probability as shown in Figure 1, conditional on their present state. The probabilities for *regular smoking*, *quit attempt* and *maintain quit* are derived from COM-B models. These probabilities vary by individual and over time, since the influencing COM-B factors are both individual-level and dynamic over the simulation run, e.g., the maintain quit probability can change as an agent progresses through the ongoing quitter tunnel state.

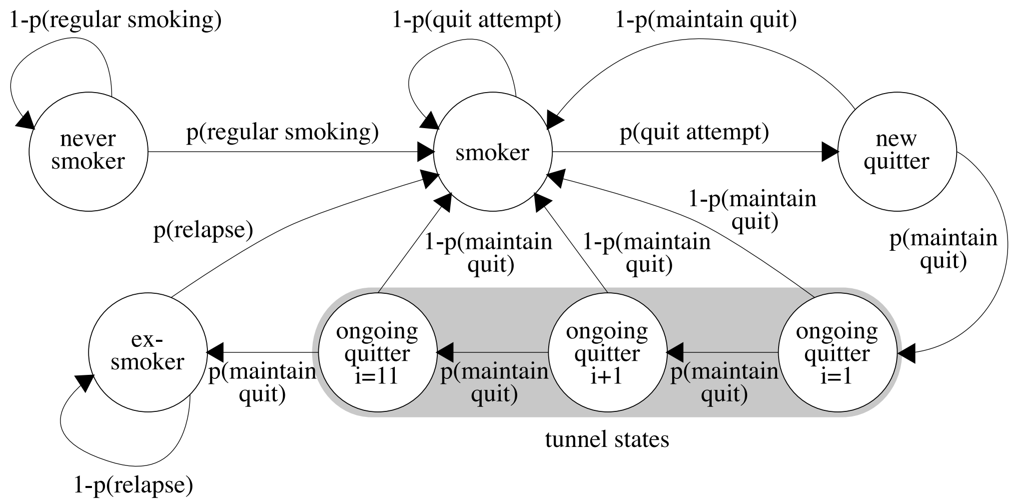

Figure 1: The state transition model modeling the state transition of an agent at a tick t. In the ongoing quitter tunnel state, the number of fully completed months of maintained abstinence is indicated by i.

3 THE SOFTWARE ARCHITECTURE OF THE AGENT-BASED MODEL

3.1 The MBSSM Software Architecture

MBSSM has a *core* architecture formed of a small number of classes with relationships between them (Figure 2). MBSSM is not itself an executable model, but provides a template for how such models can be created efficiently. The MicroAgent class represents an agent and has operations doSituation() and doAction() which trigger a situational mechanism and an action mechanism, respectively. The Theory abstract class models a mechanism-based social theory using situational and action mechanisms. A MicroAgent object interacts with Theory object(s) through a TheoryMediator object, which provides a placeholder for determining which Theory-derived classes should be used to operationalize the situational mechanisms and action mechanisms in the context of a particular model. The MacroEntity class represents a macro entity such as a social network. The Regulator abstract class models a macro-macro mechanism and a transformational mechanism. A MacroEntity object interacts with Regulator object(s) through a MacroMediator object. The Model abstract class has abstract operations for creating MicroAgent objects in the simulation, managing scheduling events and collecting simulation outputs. Here we extend MBSSM to represent COM-B theory. In theoretical terms, COM-B refers to the balance of C, O and M that is necessary for a behavioral action (or inaction) to occur *in the moment*; however, here, COM-B is abstracted to a transition between behavioral states (see Table 1) that arises at some point over the period of a tick. The choice of a monthly representation of behavior is driven by the availability of data to realize the model empirically—the monthly Smoking Toolkit Study. To support this abstraction, we define three levels (categories) of COM-B attributes that pertain to each agent:

- *Level 0* consists of the behavior (e.g., initiating regular smoking).
- *Level 1* consists of the latent composites that represent an agent's C, O and M (e.g., an agent's latent capability to initiate regular smoking) (Bollen and Noble 2011).
- *Level 2* consists of observable factors that contribute to C, O or M (e.g., e-cigarette use as a contributory factor to an agent's latent capability to transition from never smoker to smoker and is defined as an entity (BCIO:050377) in the Behavior Change Intervention Ontology (BCIO) (Schenk et al. 2023)).

Agents may also have other attributes that are not components of the COM-B model but are required for other purposes (e.g., reporting outcomes for a socio-demographic subgroup) or that act as proxies for

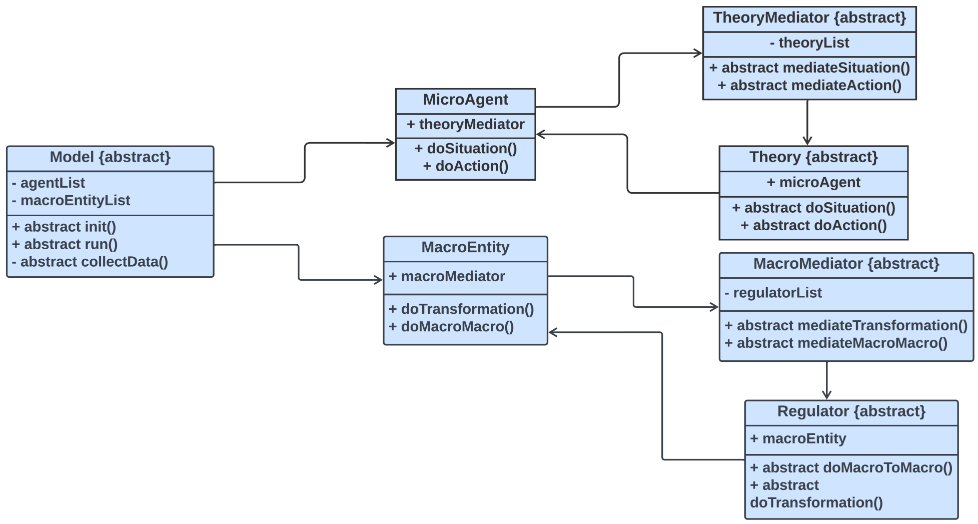

Figure 2: The Unified Modelling Language (UML) class diagram representing the MBSSM software architecture.

one or more otherwise unmeasurable Level 2 attributes. For example, the proportion of smokers in the social network is a proxy of the Level 2 attribute 'Interaction with smokers in social network' which is unmeasurable. We refer to these attributes as *personal characteristics*.

3.2 Object-Oriented Design of The Conceptual Agent-based Model

To produce a fully realized ABM design, we must extend from the core MBSSM architecture and associate the core with generic ABM software library functionality. For the latter, we assume that MicroAgent is derived from some base Agent class available in an ABM software library, that management of Agent objects is handled by an AgentPopulation class and that discrete-event scheduling functionality is available from a Scheduler class. The UML class diagram of the software architecture of the smoking ABM is illustrated in Figure 3. **COMBTheory** is an abstract subclass of Theory that represents the COM-B model and is a more specific class than Theory. A COMBTheory object is composed of the **Level1Attribute** objects: compC, compO and compM objects representing the latent composite C, O and M and the **Level2Attribute** objects which compose the Level1Attribute objects. COMBTheory is designed using the *template method design pattern* (Gamma et al. 1995) because any implementation of doAction in a COM-B model should consist of the same operations: (1) creating the Level 1 attributes from the Level 2 attributes; and (2) computing the probability of the behavior based on the Level 1 attributes (Figure 4). How these operations are carried out depends on the particular type of COM-B model being developed, so the operations in COMBTheory.doAction are all abstract. doAction first calls the abstract operations makeCompC, makeCompO, makeCompM to aggregate Level2Attribute objects into Level1Attribute objects, then calls the doBehaviour abstract operation to perform the agent's behavior and set the agent's state of the next tick based on the state transition model. Note that doSituation also remains an abstract operation because how Level2Attribute objects are influenced via situational mechanisms depends on the particular COM-B model being considered.

Now we can define the concrete types of COM-B theory that are relevant to our application example. **RegSmokeTheory** is a concrete subclass of COMBTheory implementing the regular smoking model in its doSituation, makeCompC, makeCompO, makeCompM and doBehaviour operations (Figure 5). Similary, **QuitAttemptTheory** implements the quit attempt model and **QuitSuccessTheory** implements the quit

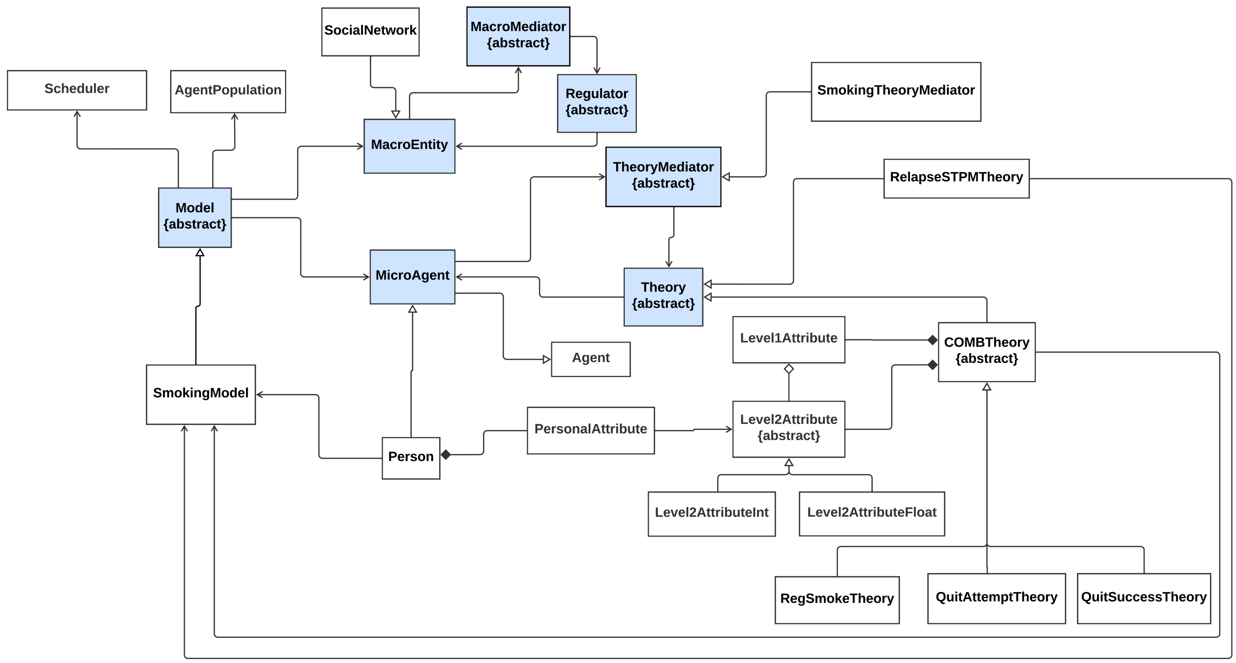

Figure 3: The class diagram of the software architecture of the agent-based model: blue boxes represent the core MBSSM software architecture; white boxes are the COM-B design and ABM software libraries.

success model. Meanwhile **RelapseSTPMTheory** is a subclass of Theory for representing the STPM relapse model and immediately implements the abstract operations doSituation and doAction of Theory.

Level 2 attributes may be categorical or continuous variables and are sometimes based on personal characteristics that have their own natural units, e.g., proportion of smokers in an agent's social network may be a personal characteristic varying continuously on the range [0,1] that may have a categorical representation as an indexed group in one COM-B model, e.g., category 1=[0,0.5], while simultaneously having a continuous representation [0.1] in another COM-B model. To allow this flexibility, we include Level2AttributeInt and Level2AttributeFloat subclasses. The **PersonalAttribute** class represents a personal characteristic. A PersonalAttribute object may then be associated with Level2Attribute object(s) of the same concept. When the value of a PersonAttribute object is set by the setValue operation, its associated Level2attribute object(s) are updated accordingly. The **Person** subclass inherits the attributes and operations of MicroAgent class of MBSSM. A Person object is composed of PersonalAttribute objects and has an updatePersonalAttribute method to update the personal characteristics. Memory of previous states is implemented through use of a buffer. **SocialNetwork** is a subclass of the MacroEntity class of MBSSM with an attribute adjMatrix to store the adjacency matrix representing the social network which is a directed graph connecting the agents (vertices) of the population. Each edge represents an interaction between two agents. The updatePersonalAttribute method of **Person** is called by the **Scheduler** to update the PersonalAttribute 'proportion of smokers in social network'. The updatePersonalAttribute method calls the outDegreeOfNode method of **SocialNetwork** to compute the number of social connections for each agent.

SmokingModel defines the overall structure of the smoking ABM. Following the MBSSM framework, it is a subclass of the Model abstract class and provides implementations of the abstract operations. The **SmokingTheoryMediator** is a subclass of the TheoryMediator abstract class and determines which COMBTheory-derived classes should be used to operationalize the situational mechanisms and action mechanisms. The **Agent**, **AgentPopulation** and **Scheduler classes** are incorporated into MBSSM using a generalization relationship between Agent and MicroAgent class and association relationships between the Model abstract class, AgentPopulation and Scheduler (Figure 3).

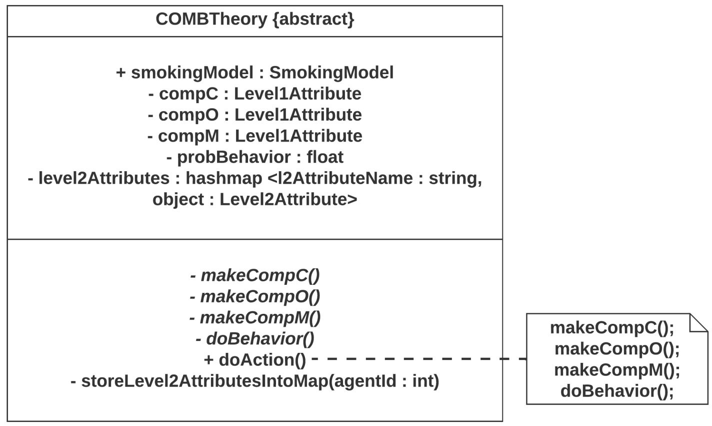

Figure 4: The class diagram of the COMBTheory abstract class.

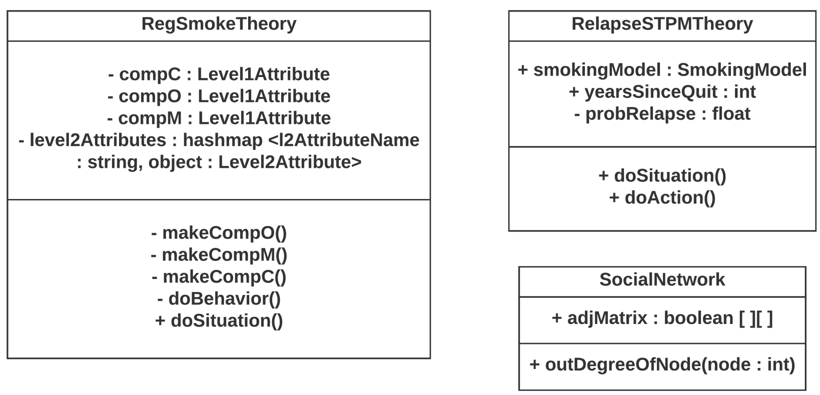

Figure 5: The class diagrams of RegSmokeTheory, RelapseSTPMTheory and SocialNetwork subclasses; the class diagrams of QuitAttemptTheory and QuitSuccessTheory are identical to that of RegSmokeTheory.

3.3 Example Implementation in Repast4Py

MBSSM can be implemented in any programming language that supports object-oriented design. It was originally implemented in C++ using RepastHPC libraries (Collier and North 2013). To improve accessibility, it was recently ported to Python using Repast4Py libraries (Collier and Ozik 2022). This is the version we have used for the COM-B and smoking model implementation. The Agent, AgentPopulation and Scheduler classes of the ABM software architecture in Figure 3 correspond to the Agent, SharedContext and SharedScheduleRunner classes of Repast4Py. We chose implementation in Repast4Py for its accessibility to public health economic modeling practitioners and its scalability to empirically representative populations of agents (8,150 in the case of the STPM population in 2012).

4 THE COM-B MODEL OF SMOKING BEHAVIORS

The C, O and M for each of the smoking behaviors (smoking initiation, quit attempt, and quit success) are modeled as latent composite variables. Each of the latent composites is constructed using observable factors associated with the smoking behavior. These factors were identified from comprehensive relevant reviews, e.g., (West 2017), and the research team's expertise; some examples are given in Table 2. The

Table 2: Directions of influence for some example Level 2 attributes with smoking behaviors: A→B represents an association between the Level 2 attribute A and the behavior B.

Level 2 attribute (BCIO ID)	Association with smoking behaviors
Prevalence of smoking in social network (006001)	More smokers→uptake more likely (O, uptake) More smokers→less likely to try to quit (O, attempt) More smokers→less likely to succeed (O, success)
Enjoyment of smoking (006159)	Higher enjoyment→uptake more likely (M, uptake)
Exposure to cessation prompts (BCIO entry not yet defined)	Greater exposure→more likely to quit (O, attempt)
Understanding of smoking harms (006055)	Greater understanding→more likely to quit (C, attempt)
Desire to stop smoking (0001127)	Greater desire→more likely to try to quit smoking (M, attempt)
Number of recent quit attempts (0000729)	More quit attempts→more likely to try to quit (M, attempt) More quit attempts→less likely to succeed (C, success)
E-cigarette use (0000665)	E-cigarette use→uptake more likely (C, O, uptake) Use of e-cigarettes while quitting→ more likely to succeed in quitting (C, success)
Smoker self-identity (0000714)	Stronger smoker identity→less likely to try to quit (M, attempt) Stronger ex-smoker identity→ more likely to succeed (M, success)
Strength of cigarette addiction (0001214)	Stronger cigarette addiction→uptake more likely (C, uptake) Stronger cigarette addiction→ less likely to succeed (C, success)

regular smoking, quit attempt and quit success models are implemented as three logistic regression models, where the probability of the smoking behavior b, where $b \in \{\text{uptake}, \text{attempt}, \text{success}\}$, of agent i at tick t is computed as:

$$p\left(b_i[t]\right) = \frac{1}{1+\exp\left(-\left(\alpha_0 + \alpha_C C_i[t] + \alpha_O O_i[t] + \alpha_M M_i[t]\right)\right)}, \tag{1}$$

with

$$C_i[t] = \sum_{x_C \in \mathscr{X}_C} \beta_{x_C} x_{C,i}[t], \tag{2}$$

$$O_i[t] = \sum_{x_O \in \mathscr{X}_O} \beta_{x_O} x_{O,i}[t], \tag{3}$$

$$M_i[t] = \sum_{x_M \in \mathscr{X}_M} \beta_{x_M} x_{M,i}[t], \tag{4}$$

where C, O and M are the Level 1 attributes representing latent capability, opportunity and motivation respectively; α_0, α_C, α_O and α_M weight the contributions of the latent components to the probability of the behavior; x_C, x_O and x_M are Level 2 attributes representing specific capabilities, opportunities and motivations—from the sets of such attributes, $\mathscr{X}_C$, $\mathscr{X}_O$ and $\mathscr{X}_M$—which are also weighted in terms of their contribution to the latent composites via β_C, β_O and β_M. While we show all Level 2 attributes as time-varying, we expect that a subset will be constant over the lifetime of each agent. The three logit models (1) are implemented in the doBehavior operations of the subclasses RegSmokeTheory, QuitAttemptTheory and QuitSuccessTheory; the formulae (2, 3 and 4) of latent composite C, O and M are implemented in the makeCompC, makeCompO, makeCompM operations of the subclasses.

5 DISCUSSION AND CONCLUSIONS

This work proposed an ABM which uses the COM-B model to simulate the smoking behaviors of a population. The MBSSM architecture was successfully extended to incorporate three proposed applications of the COM-B model—smoking initiation model, quit attempt model and quit success model—within an ABM.

The resulting design is amenable to policy analysis in a way that cannot be achieved using existing models. Considering a potential mass media campaign to encourage quit attempts, the resourcing and reach of the campaign can be defined in a new **MassMediaCampaign** MacroEntity. The effect of the campaign can be modelled by passing a reference to the MassMediaCampaign object to the doSituation method on QuitAttemptTheory. The doSituation method determines if the agent has been reached by the campaign and, if so, implements the 'theory of change' logic of a mass media campaign on the level 2 attributes: in this case, these would be to increase the agent's understanding of smoking harms (a capability), increase the agent's exposure to smoking cessation prompts (an opportunity), and to increase the agent's desire to quit (a motivation). These impacts would consequently change the agent's level 1 attributes and therefore change the probability of the agent making a quit attempt. Within existing approaches, the probability of smoking cessation would be taken from studies of mass media campaigns with limited follow up, whereby it is impossible to predict how that probability may change over time and the same probability is usually applied for each individual. However, using the new design we can explain these probabilities for different individuals. It should then be possible to predict quit attempts for different individuals over time according to their individual characteristics and psychological variables, as well as (via the mechanisms of the SocialNetwork) one person's behavior affecting another person's probability of quitting.

Future work will focus on appraising a series of such interventions. Agents in the ABM will be initialized using a synthetic population created using Health Survey England (Mindell et al. 2012) and Smoking Toolkit Study survey data (Fidler et al. 2011). Model parameters will be calibrated by comparing the emergent smoking transitions in population sub-groups to equivalent observations for the period 2011-2016, with validation using reserved data for the period 2017-2019. We have chosen a pre-Covid-19 time period for the calibration and validation to align the model to the past trends of decline in smoking prevalence. Once the calibration of parameters to this past period is complete, the differences in parameter values in the Covid-19 and post-Covid-19 periods will be investigated. The milestones are: 1) incorporating interventions and policies as macro entities into the ABM software architecture; 2) calibration and validation of the ABM using the synthetic population; and 3) simulating the impact of interventions and policies on smoking prevalence.

A APPENDIX: CONCEPTS IN THE OBJECT-ORIENTED PROGRAMMING PARADIGM

The following gives a brief description of the concepts in object-oriented programming that we subsequently adapt to implement the COM-B theory:

- An object is a thing that has attributes and can perform operations. For example, an implementation of the regular smoking model (1) is an object.
- A class is the set of all the objects of the same type. A class defines the attributes and operations to specify how its objects behave and how they are created. For example, all types of implementations of the regular smoking model form a class named RegSmokeTheory.
- A subclass is a more specific class of a more general class (superclass). The subclass inherits the attributes and operations of its superclass. For example, the Person class is a subclass of MicroAgent class.
- An abstract class serves as a template (skeleton) for a class and cannot create objects. An abstract class has abstract operations which do not have implementations and may have normal operations with implementations. Subclasses can be derived from an abstract class to provide different implementations for its abstract operations. The COM-B model is generic and expressed

as an abstract class COMBTheory; RegSmokeTheory is a subclass of COMBTheory to provide an implementation of the regular smoking model.

- UML class diagrams: Unified Modelling Language (UML) (Stevens and Pooley 2006) is a widely used graphical language to design and document object-oriented software. A UML class diagram represents the classes of a software as boxes and the relationships between them as arrows. A box representng a class consists of three parts: (1) its name, (2) its attributes; and (3) its operations. A box representing an abstract class has the word 'abstract' in its name. The abstract operations have italic names.
- An association relationship between classes A and B represents that a class A object (client) calls the operations of a class B object (server). In the UML class diagram of MBSSM (Figure 2), each arrow represents an association between the classes; an arrow points to the server from the client.
- The generalization relationship expresses a relationship between a subclass and a superclass. In Figure 3, RegSmokeTheory (subclass) and COMBTheory (superclass) are linked by a solid line with a closed arrowhead. The arrowhead points to the superclass from the subclass.
- The aggregation relationship expresses a part-whole relationship so that one or more objects O_As of class A are parts of an object O_B of class B and O_As can be parts of other objects. In Figure 3, a Level1Attribute object is an aggregation of Level2Attribute objects (shown by an open diamond in the UML diagram).
- The composition relationship expresses that an object (the whole) strongly owns other objects (its parts) so that if one or more objects O_As of class A are parts of an object O_B and owned by O_B, O_As cannot be parts of other objects. A COMBTheory object is composed of Level1Attribute and Level2Attribute objects (shown by a filled diamond in the UML diagram) because a form of the COM-B model has its own Level 1 attributes and Level 2 attributes.
- The template method design pattern (TMDP) is a well-established software design solution. In TMDP, an abstract class defines an operation (the template method) as a skeleton of an algorithm (Gamma et al. 1995). Some of its steps are deferred to the subclasses of the abstract class. The subclasses redefine certain steps of the algorithm without changing the algorithm's structure.
- The Hashmap data structure serves as a dictionary which stores associations of keys and values as pairs of key and value. The keys and values can be objects of the same class, different classes or data structures. The values are accessed by their keys. COMBTheory has a hashmap attribute (level2attributes) which associates the names of Level 2 attributes with Level2Attribute objects.

ACKNOWLEDGEMENTS

This work was supported by Cancer Research UK (PRCRPG-Nov21\100002). Hazel Squires' input was funded via a NIHR Advanced Fellowship (NIHR301406). The views expressed are those of the authors and not necessarily those of the NIHR or the Department of Health and Social Care. All authors declare no financial links with tobacco companies, e-cigarette manufacturers, or their representatives.

REFERENCES

Antosz, P., D. Birks, B. Edmonds, A. Heppenstall, R. Meyer, J. G. Polhill *et al*. 2023. "What Do You Want Theory For? – A Pragmatic Analysis of The Roles of "Theory" in Agent-based Modelling". *Environmental Modelling and Software* 168:105802.

Atkinson, J. A., D. Knowles, and J. Wiggers. 2018. "Harnessing Advances in Computer Simulation to Inform Policy and Planning to Reduce Alcohol-related Harms". *International Journal of Public Health* 63:537–546.

Beard, E. V., R. West, M. Jarvis, S. Michie and J. Brown. 2019. "'S'-shaped Curve: Modelling Trends in Smoking Prevalence, Uptake and Cessation in Great Britain from 1973 to 2016". *Thorax* 74(9):875–881.

Bollen, A. K. and D. M. Noble. 2011. "Structural Equation Models and the Quantification of Behavior". *Proceedings of the National Academy of Sciences* 108:15639–15646.

Collier, N. and M. North. 2013. "Parallel Agent-Based Simulation with Repast for High Performance Computing". *SIMULATION* 89(10):1215–1235.

Collier, N. and J. Ozik. 2022. "Distributed Agent-Based Simulation with Repast4Py". In *2022 Winter Simulation Conference (WSC)*, 192–206 https://doi.org/10.1109/WSC57314.2022.10015389.

Fidler, J. A., L. Shahab, and O. West. 2011. "The Smoking Toolkit Study: A National Study of Smoking and Smoking Cessation in England". *BMC Public Health* 11:479.

Gamma, E., R. Helm, R. Johnson, and J. M. Vlissides. 1995. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional.

Gilbert, E. 2023. "Using the COM-B Model to Identify Barriers and Facilitators in Smoking Behaviour to Support Both Stop Smoking Practitioners and Service Users at Everyone Health, England". *Tobacco Prevention & Cessation* 9(Supplement):A148.

Gillespie, D. and A. Brennan. 2023. "The Sheffield Tobacco Policy Model (STPM): Full Technical Documentation". Technical Report Documentation version number 2.0.1, The University of Sheffield.

Hedstrom, P. and R. Swedberg. 1996. "Social Mechanisms". *Acta Sociologica* 39(3):281–308.

Huang, W., C. H. Chang, and E. A. Stuart. 2021. "Agent-Based Modeling for Implementation Research: An Application To Tobacco Smoking Cessation for Persons With Serious Mental Illness". *Implementation Research and Practice* 2:1–16.

Luke, D. A., R. A. Hammond, and T. Combs. 2017. "Tobacco Town: Computational Modeling of Policy Options to Reduce Tobacco Retailer Density". *American Journal of Public Health* 107(5):740–746.

Michie, S., M. M. V. Stralen, and R. West. 2011. "The Behaviour Change Wheel: A New Method for Characterising and Designing Behaviour Change Interventions". *Implementation Science* 6:42.

Mindell, J., J. P. Biddulph, and V. Hirani. 2012. "Cohort Profile: The Health Survey for England". *International Journal of Epidemiology* 41(6):1585–1593.

Naughton, F., A. Hope, C. Siegele-Brown, K. Grant, G. Barton, C. Notley *et al*. 2023. "An Automated, Online Feasibility Randomized Controlled Trial of a Just-In-Time Adaptive Intervention for Smoking Cessation (Quit Sense)". *Nicotine and Tobacco Research* 25(7):1319–1329.

Reitsma, M. B., L. S. Flor, and E. C. Mullany. 2021. "Spatial, Temporal, and Demographic Patterns in Prevalence of Smoking Tobacco Use and Attributable Disease Burden in 204 Countries and Territories, 1990 – 2019: A Systematic Analysis from The Global Burden of Disease Study 2019". *The Lancet* 397(10292):2337–2360.

Schenk, P., A. J. Wright, R. West, J. Hastings, F. Lorencatto and C. Moore. 2023. "An Ontology of Mechanisms of Action in Behaviour Change Interventions". *PsyArXiv preprint PsyArXiv: 7jpm3*.

Stevens, P. and R. Pooley. 2006. *Using UML - Software Engineering with Objects and Components, Second Edition*. Addison-Wesley.

Sukthankar, G. and R. Beheshti. 2019. "Using Agent-based Models to Understand Health-related Social Norms". In *Social-behavioral Modeling for Complex Systems*, 633–654. Wiley.

Vu, T. M., C. Buckley, J. Duro, A. Brennan, J. M. Epstein and R. Purshouse. 2023. "Can Social Norms Explain Long-Term Trends in Alcohol Use? Insights from Inverse Generative Social Science". *Journal of Artificial Societies and Social Simulation* 26(2):4.

Vu, T. M., C. Probst, A. Nielsen, H. Bai, P. S. Meier, C. Buckley, , *et al*. 2020. "A Software Architecture for Mechanism-Based Social Systems Modelling in Agent-Based Simulation Models". *Journal of Artificial Societies and Social Simulation* 23(3):1.

West, R. 2017. "Tobacco Smoking: Health Impact, Prevalence, Correlates and Interventions". *Psychology & Health* 32(8):1018–1036.

West, R. and J. Brown. 2013. *Theory of Addiction*. 2nd ed. Chichester: John Wiley & Sons, Ltd.

Zhong, X., X. Li, and S. Mangoni. 2023. "A Review of Agent-Based Modeling Applications in Substance Abuse Policy Research". In *2023 Winter Simulation Conference (WSC)*, 150–161 https://doi.org/10.1109/WSC60868.2023.10407530.

AUTHOR BIOGRAPHIES

DAVID TIAN is a Research Associate in population health modeling at the University of Sheffield. His research interests include development of agent-based modeling software, AI and machine learning methods for population health and healthcare. His email address is d.tian@sheffield.ac.uk.

HAZEL Y. SQUIRES is a Senior Research Fellow in health-related modeling at the University of Sheffield. Her research interests include conceptual modeling, simulation for health-related decision making, and health economic modeling. Her email address is h.squires@sheffield.ac.uk.

CHARLOTTE BUCKLEY is a Research Fellow in population health modeling at the University of Sheffield. Her research interests include the development and utilisation of simulation methods to understand population-level health behaviors and their impacts of policies. Her email address is c.m.buckley@sheffield.ac.uk.

DUNCAN GILLESPIE is a Senior Research Fellow in population health modeling at the University of Sheffield. His research interests include the development of models to support decision-making to improve population health with a focus on tobacco and alcohol policy. His email address is duncan.gillespie@sheffield.ac.uk.

HARRY TATTAN-BIRCH is a Research Fellow in Behavioural Science and Health at University College London. His research interests include smoking cessation research underpinned by behavioral science. His email address is h.tattan-birch@ucl.ac.uk.

LION SHAHAB is a Professor of Health Psychology at University College London (UCL) and Co-Director of the UCL Tobacco and Alcohol Research Group. His research interests include tobacco control policy and novel smoking cessation interventions. His email address is lion.shahab@ucl.ac.uk.

ROBERT WEST is an Emeritus Professor in Behavioural Science and Health at University College London. His research interests include behavior change including addictions, especially tobacco use, and motivation. His email address is robert.west@ucl.ac.uk.

ALAN BRENNAN is a Professor of Health Economics and Decision Modelling at University of Sheffield. His research interests include methodological and applied research and consultancy using mathematical modeling in health and healthcare. His email address is a.brennan@sheffield.ac.uk.

JAMIE BROWN is a Professor of Behavioural Science at University College London (UCL) and Director of the UCL Tobacco and Alcohol Research Group. His research interests include smoking cessation research underpinned by behavioral science. His email address is jamie.brown@ucl.ac.uk.

ROBIN C. PURSHOUSE is a Professor of Decision Sciences at the University of Sheffield. His research interests include the development of complex systems modeling methods for application to policy questions in population health. His email address is r.purshouse@sheffield.ac.uk.