

DATA AND DESIGN: ADVANCING THEORY FOR COMPLEX
ADAPTIVE SYSTEMS

A Thesis Submitted to the
College of Graduate and Postdoctoral Studies
in Partial Fulfillment of the Requirements
for the degree of Doctor of Philosophy
in the Department of Computer Science
University of Saskatchewan
Saskatoon

By

L. Kurt Kreuger

©L. Kurt Kreuger, January 2018. All rights reserved.

PERMISSION TO USE

In presenting this thesis in partial fulfilment of the requirements for a Postgraduate degree from the University of Saskatchewan, I agree that the Libraries of this University may make it freely available for inspection. I further agree that permission for copying of this thesis in any manner, in whole or in part, for scholarly purposes may be granted by the professor or professors who supervised my thesis work or, in their absence, by the Head of the Department or the Dean of the College in which my thesis work was done. It is understood that any copying or publication or use of this thesis or parts thereof for financial gain shall not be allowed without my written permission. It is also understood that due recognition shall be given to me and to the University of Saskatchewan in any scholarly use which may be made of any material in my thesis.

Requests for permission to copy or to make other use of material in this thesis in whole or part should be addressed to:

Head of the Department of Computer Science
176 Thorvaldson Building
110 Science Place
University of Saskatchewan
Saskatoon, Saskatchewan
Canada
S7N 5C9

Dean
College of Graduate and Postdoctoral Studies
116 – 110 Science Place
University of Saskatchewan
Saskatoon, Saskatchewan
Canada
S7N 5C9

ABSTRACT

Complex adaptive systems exhibit certain types of behaviour that are difficult to predict or understand using reductionist approaches, such as linearization or assuming conditions of optimality. This research focuses on the complex adaptive systems associated with public health. These are noted for being driven by many latent forces, shaped centrally by human behaviour.

Dynamic simulation techniques, including agent-based models (ABMs) and system dynamics (SD) models, have been used to study the behaviour of complex adaptive systems, including in public health. While much has been learned, such work is still hampered by important limitations. Models of complex systems themselves can be quite complex, increasing the difficulty in explaining unexpected model behaviour, whether that behaviour comes from model code errors or is due to new learning. Model complexity also leads to designs that are hard to adapt to growing knowledge of the subject area, further reducing model-generated insights.

In the current literature of dynamic simulations of human public health behaviour, few focus on capturing explicit psychological theories of human behaviour. Given that human behaviour, especially health and risk behaviour, is so central to understanding of processes in public health, this work explores several methods to improve the utility and flexibility of dynamic models in public health. It is undertaken in three projects.

The first uses a machine learning algorithm, the particle filter, to augment a simple ABM in the presence of continuous disease prevalence data from the modelled system. It is shown that, while using the particle filter improves the accuracy of the ABM, when compared with previous work using SD with a particle filter, the ABM has some limitations, which are discussed.

The second presents a model design pattern that focuses on scalability and modularity to improve the development time, testability, and flexibility of a dynamic simulation for tobacco smoking. This method also supports a general pattern of constructing hybrid models — those that contain elements of multiple methods, such as agent-based or system dynamics. This method is demonstrated with a stylized example of tobacco smoking in a human population.

The final line of work implements this modular design pattern, with differing mechanisms of addiction dynamics, within a rich behavioural model of tobacco purchasing and consumption. It integrates the results from a discrete choice experiment, which is a widely used economic method for study human preferences. It compares and contrasts four independent addiction modules under different population assumptions. A number of important insights are discussed: no single module was universally more accurate across all human subpopulations, demonstrating the benefit of exploring a diversity of approaches; increasing the number of parameters does not necessarily improve a module's predictions, since the overall least accurate module had the second highest number of parameters; and slight changes in module structure can lead to drastic improvements, implying the need to be able to iteratively learn from model behaviour.¹

¹Chapter 2 was lightly altered from the approved version due to image copyright restrictions, most notably Figure 2.11 and an explanatory paragraph.

ACKNOWLEDGEMENTS

There are many I would like to thank for making this work possible. Most notably, I would thank my advisor, Dr. Nathaniel Osgood. He has been not only a tremendous source of academic advice and learning, but has also offered a profound support in life beyond the lab. He showed great patience and encouragement as I grew in my learning. He provided me with many wonderful opportunities to learn both in the lab, and in real-world projects. His wise counsel guided me and gave me opportunities to make my own mistakes. He is truly a mentor. In my own future work, I will hope to mirror these qualities.

There are several other academic mentors I would like to mention. Dr. Kelvin Choi provided a clear and insightful perspective on the tobacco use landscape. His eagerness to work across disciplinary lines encouraged my steps into this field. Dr. Geoff McDonnell provided many hours of consultation on modelling philosophies and approaches, and very useful critical feedback of many of my modelling attempts.

Throughout my time in this research, I have been blessed to work alongside very capable and helpful students. Foremost among them is Weicheng Qian, who was very helpful in developing the modular approach, and spent many an hour exploring a wide range of topics important to my research, spanning the gamut from experimental design and calibration to model debugging.

I would also like to thank the members of my committee: Drs. Michael Horsch, Cheryl Waldner, Regan Mandryk and Ian Stavness. They provided very useful guidance in how to improve my work. I have much to learn and look forward to that process.

Certainly this project, and many others, would be impossible without a strong support staff. There are too many to mention here, but Merlin Hansen was truly a magician in helping me with various troubleshooting questions in running the many simulations on the lab machines. Gwen Lancaster was always helpful in answering questions and proactively reminding about administrative matters.

Finally, I would be able to achieve very little without the dedication and support of my family and friends. My parents instilled in me the quest for knowledge and to be of service to the world. My wife, Amelia, not only encouraged me to pursue my desires, but helped me achieve whatever measure of personal excellence I may reach in that pursuit. My close friend, Parisa, always challenged me to reach farther.

To all these people, and the many more I haven't named, thank you. I hope I can pass a measure of your support on to those I work with in the future.

STATEMENT OF WORK

This work is the result of collaborations between many individuals. As a result, it is useful to share the breakdown of work for the various projects. As my advisor, Dr. Osgood supported the work across all projects with regular consultation.

For the work in Chapter 3, I was responsible for all the model construction. Dr. Osgood assisted in a code review to validate the correct implementation of the particle filtering algorithm. He also provided feedback when interpreting the initial results, and helped in running simulation experiments on his computer.

In Chapter 4, Winchell Qian and I built the module software interface, which allowed me to develop the model used to test it, including the two addiction modules defined. I also developed the experimental design.

In Chapter 5, a few experts were consulted. Dr. Swait took the lead on the survey experimental design, with Dr. Choi and I consulting on needs. Dr. Choi also provided analysis and data, and consultation (from a domain expert perspective) on model construction and validation. Xing Zhang and I performed the DCE analysis, with him focusing on developing the R code used, and I providing pre-processing. I sourced the survey population, constructed the survey and the generation mechanisms used for it, and performed user testing. I developed all addiction modules, implemented the DCE within the main model, performed validation, designed the experiment, conducted the calibrations, and conducted the analysis.

To my new daughter, Lua. Before you could even see light, you were mine.

CONTENTS

Permission to Use	i
Abstract	ii
Acknowledgements	iii
Statement of Work	iv
Contents	vi
List of Tables	viii
List of Figures	ix
List of Abbreviations	xi
1 Introduction	1
1.1 Complex Adaptive Systems	1
1.2 Public Health	2
1.3 Dynamic Modelling	4
1.3.1 Contributions	6
1.3.2 Dissertation Outline	8
2 Background and Related Work	9
2.1 Background	9
2.1.1 Motivation	9
2.1.2 Classification Dimensions	9
2.1.3 Literature Search Method	10
2.2 Simulation-based Behavioural Frameworks	11
2.2.1 Ad hoc	11
2.2.2 Intentional	14
2.2.3 Normative	16
2.2.4 Cognitive	21
2.3 Social-cognitive	25
2.3.1 Health Belief Model	26
2.3.2 Theory of Planned Behaviour	28
2.4 Observables	30
2.4.1 Static Decision Tree	30
2.4.2 Dynamic Behaviour	31
2.4.3 Activity-Based Modelling	33
2.4.4 SIENA	34
2.5 Aggregate Models	35
2.5.1 SimSmoke	35
2.5.2 System Dynamics Models	36
2.6 Discussion	39
3 Particle Filtering in Agent-based Models	43
3.1 Introduction	43
3.2 Particle Filtering Algorithm	45
3.2.1 Intuition and Derivation	45

3.2.2	Final Algorithm	47
3.3	Ground Truth Model	47
3.3.1	Introduction	47
3.3.2	Ground Truth Model	48
3.4	Methods and Results	50
3.5	Discussion	51
3.6	Summary	55
4	Modular Design Pattern	57
4.1	Introduction	57
4.2	Technique Description	58
4.2.1	Modularity Through The Java Interface	59
4.3	Results	63
4.3.1	Markov addiction model	64
4.3.2	PCT addiction model	64
4.4	Discussion	66
4.5	Summary	67
5	Tobacco Price-minimizing Strategy Model	68
5.1	Introduction and Motivation	68
5.2	The Model	69
5.2.1	Environment	70
5.2.2	Discrete Choice Purchasing Behaviour	71
5.2.3	Addiction Modules	75
5.3	Experiment	81
5.3.1	Validation	82
5.3.2	Calibration	83
5.3.3	Bootstrapped DCE Sensitivity Analysis	84
5.4	Results	85
5.5	Discussion	87
5.6	Summary	91
6	Conclusions	94
6.1	Future Work	96
	References	98

LIST OF TABLES

5.1	Counts for survey responses per sub-population and year.	82
-----	--	----

LIST OF FIGURES

2.1	BDI design, reproduced with permission from [7]	15
2.2	EMIL design, reproduced with permission from [7]	17
2.3	LNA interactions. Agents start as smoker or non-smoker. With each interaction, given their and their partner’s smoking state, they choose the option with the highest payback, calculated according to the equations and their values. α and β are parameters meant to bias maintenance of the current state.	19
2.4	Internal representation of a PECS agent. There are 3 layers. The 1st layer, the sensor and the perception, observes the world. The 2nd layer is the internal representation of the agent, captured by the PECS components. Z and F indicate the internal state and transition functions for each PECS component, respectively. The 3rd layer, the behaviour and actor blocks, captures action with the world.	22
2.5	Consumat heuristics, reproduced from [7], in relation to when they are invoked. Note that this depends not only on the problem, but on the agent. Details of the heuristics are not discussed here. See [7] for a more comprehensive discussion.	24
2.6	Formulation of the Health Belief Model. The 4 elements along the bottom represent the various perceptions. Cues to action represent external triggers to action. Health motivation captures confidence in one’s ability to act.	27
2.7	The Theory of Planned Behaviour, including the functional form of elements. Intention is the primary driver of behaviour.	29
2.8	Static Decision Tree from [70]. Transition probabilities are shown on the arrows. There are 5 unique observable behaviour states: never user, cigarette user, smokeless user, dual user, and quit user.	31
2.9	Markov chain used by Killeen [50], where possible agent states are: u for use, w for withdrawal, or a for abstention. Transitions between states are fit to the data.	32
2.10	Smoking states and transitions of the basic SimSmoke model [58]. Initiation and cessation rates are categorized by age, sex, and race/ethnicity. Relapse rates are categorized by age and years as ex-smoker. Years as ex-smoker is categorized in several year groupings: <1, 1-2, 3-5, 6-10, 11-15, and >15.	35
2.11	This figure has been redacted for copyright reasons. It depicted the New Zealand Tobacco Policy Model with system dynamic elements, showing 5 smoking categories—never smokers, youth smokers, adult smokers, recent ex-smokers, and non-recent ex-smokers—and various interactions between them, such as how the adult smoker population influences the creation of youth smokers through role modelling. It can be found in [97].	37
2.12	Final ISIS diagrammatic model. [14] Originally published as Figure 1 in A. Best et al. 2006. Systemic transformational change in tobacco control: An overview of the Initiative for the Study and Implementation of Systems (ISIS). In Innovations in health care: A reality check, ed. A. L. Casebeer, A. Harrison, and A. L. Mark, 189205. New York: Palgrave Macmillan. Reproduced with permission from Palgrave Macmillan.	38
2.13	Stock-flow diagram from Feedback Rich Model [104].	39
3.1	Agent state chart in the ground-truth model	50
3.2	Box plot of error metric over 9 simulations using 3 network types (random, small world, distance-based) with 3 particle filtering setups (no weight updating, updating every 1 time step, and updating every 5 time steps).	51
3.3	a) Box plot of error metric over 3 simulations using a distance-based network with ground truth connection distance of 41.5 and particle connection distance drawn from a uniform distribution from 20 to 60. Results shown for no weight update period, 1 time step updates, and 5 time step updates. b) Box plot of error metric over 3 simulations using a ring-lattice network with connection count of 2 per agent. Results shown for no weight update period, 1 time step updates, and 5 time step updates.	52

3.4	Box plot of error metric over 2 simulations using a ring-lattice network.	52
3.5	Similar particle filter using a compartmental particle model. This was obtained with permission from the original author's source material.	53
3.6	Example output for small world network ($n = 10$, $\beta = 0.9$). Weight update occurs every 10 time steps starting at $t = 1$	54
4.1	Our hybrid technique.	58
4.2	Agent smoking statechart.	59
4.3	IAddictionModel code.	60
4.4	System dynamics diagram of our implemented PCT module.	62
4.5	Histogram of final smoking status.	64
4.6	Histogram of daily smoking rates for smokers and population-wide average of daily smoking rates over time.	65
5.1	Example of stylized city. There are 3 radial sectors. From the inside out these represent urban, suburban, and rural areas respectively.	70
5.2	Sample DCE survey question.	73
5.3	Impact of inventory on baseline consumption hazard	76
5.4	Stock and flow depiction of the Bobashev Control Theory approach.	79
5.5	Best objective values for each module and sub-population.	85
5.6	Comparison of the best-fit module by sub-population between the historical mean daily smoking rate and the simulated daily consumption hazard. Calculated from the mean of 50 bootstrap iterations. The error bars represent the max and min of all bootstrap sensitivity experiments.	86
5.7	Comparison between historical and predicted values of the proportion of current smokers of the best-fit module, by sub-population. Calculated from the mean of 50 bootstrap iterations. The error bars represent the max and min of all bootstrap sensitivity experiments.	86
5.8	Comparison between historical and predicted values of the proportion of former smokers of the best-fit module, by sub-population. Calculated from the mean of 50 bootstrap iterations. The error bars represent the max and min of all bootstrap sensitivity experiments.	87
5.9	Fractional difference between historical and predicted values summed across all sub-populations for each module, including the total difference.	88

LIST OF ABBREVIATIONS

ABM	Agent-based Model
BA	Behavioural Attitude: A part of the TPB described in Section 2
BCT	Bobashev Control Theory: See Section 5
BMI	Body Mass Index
BDI	Beliefs, desires, intentions: A model described in Section 2
CAS	Complex Adaptive System
CSL	Cognitive Social Learners: A model described in Section 2
DCE	Discrete Choice Experiment: See Section 5
DES	Discrete Event Simulation
DNA	Deliberate Normative Agents: See Section 2
ELM	Elaboration Likelihood Model: A model described in Section 2
FlexP	Perceptual Control Theory with flexible set point: See Section 5
GIS	Geographic Information System
HBM	Health Belief Model: A model described in Section 2
IBM	Individual-based Model
IDM	Inventory-driven Markov: See Section 5
KF	Kalman Filter: See Section 3
LNA	Lightweight Normative Architecture: A model described in Section 2
MATS	Minnesota Adult Tobacco Surveys: See Section 5
MLE	Maximum Likelihood Estimate: See Section 3
MSM	Men who have sex with men
NoA	Normative Agent Architecture: See Section 2
PBC	Perceived Behavioural Control: A part of the TPB described in Section 2
PCT	Perceptual Control Theory: See Sections 4 and 5
PECS	Pysis, Emotion, Cognition, Social status: A model described in Section 2
RA	Relative Agreement: See Section 2
RurBA	Sub-population in Ch. 5: Rural with bachelor's or more.
RurHS	Sub-population in Ch. 5: Rural with high school or below.
RurSC	Sub-population in Ch. 5: Rural with technical school or some college.
SD	System Dynamics
SEIR	Susceptible, Exposed, Infective, Recovered: See Section 3
SIR	Susceptible, Infective, Recovered: See Section 3
SN	Subjective Norms: A part of the TPB described in Section 2
SubBA	Sub-population in Ch. 5: Suburban bachelor's or more.
SubHS	Sub-population in Ch. 5: Suburban with high school or below.
SubSC	Sub-population in Ch. 5: Suburban with technical school or some college.
TPB	Theory of Planned Behaviour: A model described in Section 2
TPM	Tobacco Policy Model: A model described in Section 2
TPMS	Tobacco Price Minimizing Strategy: See Section 5
UrbBA	Sub-population in Ch. 5: Urban with bachelor's or more.
UrbHS	Sub-population in Ch. 5: Urban with high school or below.
UrbSC	Sub-population in Ch. 5: Urban with technical school or some college.

CHAPTER 1

INTRODUCTION

1.1 Complex Adaptive Systems

Forests and grasslands, coral reefs and fish stocks, the stock market, the human immune system, the Internet, the World Wide Web - these can all be treated as Complex Adaptive Systems (CASs). There is growing consensus in a diversity of fields [73, 37, 43, 57] that investigating CASs represents an important new direction in scientific research. There are many elements that distinguish a CAS from those that are merely complicated, but perhaps the most common is that the understanding of the constituent elements of any CAS does not necessarily lead to knowledge of the behaviour of the whole system [40], often phrased as “the whole is greater than the sum of its parts”, a distinction that reflects the phenomenon of “emergence”. This necessitates going beyond, but using the knowledge from, reductionist modes of research to understand the science of the whole.

CASs are systems composed of many interacting, heterogeneous actors with their own diverse rules of behaviour. These systems usually have little or no centralized control, and lack a single governing equation [43]. A primary distinction between complex and complicate systems is one of non-linearity, which not only describes many of the local interactions, but through feedbacks, is also present at the system level [105].

One such behaviour is regime changes – unexpected, sudden, and often drastic changes towards a new behavioural steady-state. We witness such transitions when stock markets crash, when local environments are overtaken by invasive species, or in the extinction of endemic infections in a population due to the development of herd immunity. This reflects the fact that CASs live within a state space that possesses multiple basins of attraction. Due to this non-linearity, analytic solutions are rarely available, and linear or reductionist models are often misleading [57]. If a system is likely to stay within a single steady state, then linearization around that steady state might be useful. But if we seek substantial change or a broader knowledge of behaviour, then we need new tools.

The evolution over time of a CAS is guided by an important balance – diversity and redundancy. It is the diversity of the internal system actors that enables evolution to take place. But more diversity does not always lead to better systems, since too much adaptation can lead to very fragile systems [57]. For this reason, CASs also need to maintain a measure of internal redundancy, which can serve to protect a system from external shocks that would otherwise destroy it. For example, evolution under constant conditions leads to a reduction in genetic diversity, thereby reducing the adaptability of a system [57]. This highlights the

importance of the feedback that exists between the system and its environment, and therefore the need to account for both.

A CAS is distinguished not only its diversity, but also by its connectivity [57]. And while connectivity is often a necessary feature for system resilience, it also often means dependence. Without a corresponding measure of modularity, the health of a system is at risk. For example, foreign trade can increase wealth and improve welfare when it increases the connectivity between different human communities, but this can lead to global financial crises when the system is not sufficiently bulwarked against the collapse of certain key nodes in the network [57].

In many CASs, the operation of cycles on different time scales is observed [57]. While some processes occur relatively quickly, others are much slower. For example, algae and fish populations work on shorter time scales than the coral reefs they inhabit. In the presence of faster cycles, we can miss the slower phenomena, which can lead to mistakes in management [57]. This idea of time scales also suggests that some systems have a maximum rate of change that can be induced. Knowing, for example, that it will take 20 years to affect the desired change in the health outcomes of a human population will not only manage expectations, but it might allow management agencies to plan for the required sustained support of such a program.

This interaction between system and environment is responsible for other types of behaviour as well, such as hysteresis and irreversibility. Hysteresis is the dependence of a system's future on its history, and is present in many CASs. This path dependence can make universal predictions of the future difficult without detailed knowledge of the past. This also leads to the trait of irreversibility, since going back to an earlier state might require recreating that earlier history, which is impossible in some situations, because of loss of information or other key factors.

The purpose of this section is not to argue that we need to avoid simplifications when studying complex systems. Indeed, it can be argued that science is a quest for simplicity [73]. We must, therefore, develop more powerful simplifications of the inherently complex world around us. As part of this process, then, we can reasonable ask, "What methods can be used to bring useful simplicity to complex adaptive systems?"

1.2 Public Health

While CASs can be seen across many fields, this dissertation focuses on the application domains of public/population health (henceforth, public health), public policy and epidemiology. These domains can readily be considered to be complex and adaptive [14]. It is no longer seen as possible to be able to solve all public health crises in isolation from the others, as is supported by traditional reductionist approaches. Many [46, 103, 85] have called for more approaches in public health research in general that can inform and provide insight into these complex systems. In 2015, Carey et al. [19] performed a review of system science literature in public health. One of the main conclusions is that public health has not taken full advantage of analytical system science approaches. What makes these fields so worthy of attention is the degree to which latent

factors play a role. While this is true of many CASs, it is especially true here.

Consider, for example, an outbreak of H1N1. What are the public health agencies able to do to avert or mitigate the outbreak? Much is known about this pathogen and associated illness, including the method of transmission, which species are affected, and its generalized infectiousness, incubation and infectious periods. However, in any given outbreak we might not even know the number of currently infected individuals, perhaps having only an approximation of the number of reported cases. And while we have some awareness of the types of social network conditions that lead to someone being a high risk for spreading the disease more widely, we do have only limited understanding as to the social network structure and dynamics associated with actual human interactions. Nor does our knowledge of immune system functioning give us the ability to predict the probability that a given interaction will result in an infection for the susceptible individual.

Beyond this, interventions can have negative unintended side-effects that counteract, or indeed worsen, the public health situation. If, for example, the chosen response is for anyone with certain symptoms to present themselves at medical clinics, emergency departments, or family doctors, an unintended side-effect could be that the emergency departments get over-filled, which increases wait times, leading some infective individuals to turn away and head home. Or, with the increase in possibly infective people waiting in a room with others who, being in the hospital might be more immuno-compromised, the infection could spread to those least likely to be able to fight it off.

This is not only common in epidemiological situations. When trying to reduce the level of obesity within a population, an education program might be designed in order to empower people to make better life choices. However, this could make it more likely for people to place blame on obese individuals for their condition. If weight can be improved by making better choices, then some might reason that people are obese because they make poor choices. This stigmatization is common in North America, and the rate is increasing [82]. In fact, it is sometimes believed that stigmatization is beneficial to obese individuals as a motivation tool. But it also might increase rates of mental health disorders, leading to behaviours that only exacerbate the problem.

Finally, sometimes the solution to problems involves a combination of interventions rather than any single intervention. For example, in the literature on homeless adults with severe mental illness, some have contrasted two main approaches - housing first (which offers immediate and permanent housing not conditional on treatment compliance) and treatment first (which provides temporary housing conditional on compliance before the possibility of permanent housing) [80]. Studies have shown that individuals who experience a treatment first option are more likely to attend treatment services, but also more likely to leave early and have higher rates of substance abuse. It appears that providing conditional housing reduces the effect of treatment programs, but providing unconditional housing actually increases their effectiveness.

All of these situations exhibit hallmarks of CASs – multiple feedbacks and over different time scales, no centralized imposition of order, and high levels of connectedness. And since they deal with human populations, there is a tremendous diversity in behaviours and responses, all of which can be seen as plausible adaptations

to external forces.

In the past 10 or so years, there has been a growth in the interest of system science methods as applied to questions of public health [19], and this includes the dynamic modelling methods discussed below. While traditional statistical methods are still highly valuable, there is growing consensus that more is needed in order to identify the underlying causes of public health problems, rather than focusing on the symptoms; in other words, to put in place fixes that stay fixed.

1.3 Dynamic Modelling

Many authors in public health, epidemiology, and public policy ([43, 46, 57, 73], to name a few) suggest that the way forward includes the use of dynamic models. These are models that are designed to characterize the behaviours of complex systems. They can be broadly classified as either aggregate or individualized. Aggregate or compartmental modelling uses mathematical relationships between significant system variables and rates or probabilities of change in such variables. System Dynamics (SD) is perhaps the leading exemplar of this approach. It focuses on causal feedbacks and accumulations. Various system variables (such as the number of people who are infected) are placed in relation to the forces that affect them (such as infection or immunization) and described using continuous-time differential equations. In public health modelling, the quantity that is aggregated is usually the person, so system stocks are often counters of individuals by status. However, many aggregate models include stocks for latent (non-observed) quantities, such as the strength of a social norm or the strength of public awareness, quantities that might, in turn influence or be influenced by the observed quantities. This increases their flexibility in supporting learning about system behaviour as a whole.

One of the simplest compartmental models is the SIR model, named after the 3 disease states (stocks) that it defines: Susceptible, Infective, and Recovered. It then posits rules for the rates of moving from one stock to the next. For example, susceptible agents become infected based on the size of the infective stock, and a proposed effective contact rate. While very simple, the SIR model is able to describe the changing rate of infection based on the changing populations in each stock, capturing the distinct phases associated with outbreaks. A slight variant, the SIRS model, allows recovered individuals to return to a susceptible state after an appropriate waning of immunity. The SIRS model is able to demonstrate cycles of outbreak. When the susceptible population is initially reduced due to infection, the infection rate reduces after a time delay. But when the susceptible population increases again after immunity wanes, the infection is able to grow quickly again. There are many other variants that are built to examine different types of illnesses, including those without an immune stage (SIS model) for pathogens such as *Neisseria gonorrhoeae* and *Chlamydia trachomatis*, those with an additional exposed stage where there is a notable incubation period (SEIR) such as Measles and Influenza, as well as models including a carrier stage where some people never recover but continue to carry the infection, such as with Methicillin-resistant *Staphylococcus aureus*. See Section 2.5

for several examples of SD models used in public health.

The other broad class of models are focused on the individual level. These individual-based models (IBMs) center on describing individual-level behaviour. Perhaps the most pertinent to system science, the agent-based model (ABM), focuses attention on individual actors, or agents, with the goal of observing and understanding the resulting emergent behaviour. Agents have characteristics that generally possess some degree of heterogeneity, are given defined rules of behaviour, and exist within some type of environment that allows agent interaction, either spatial, social, or both. This specification of the data-generating mechanism also allows researchers to conduct counter-factual simulations. What makes ABMs useful is that the degree of heterogeneity and interaction within the agent population generally make statistical or aggregate prediction very difficult.

An example would be illustrative. An early, but very significant, ABM is the Schelling segregation model [88]. This is, in its original form, simple enough to investigate even without a computer, although being run as a computational simulation allows probing of many more experimental conditions. One simple form of the model posits two types of agents, say X and O. They are placed randomly on a discrete space, such as a chess board, with half X and half O, and a few empty spaces. Neighbourhoods are defined as the 8 cells surrounding any given cell. The single model parameter controls the proportion of neighbourhood similarity sought by a given agent: if it is 0.5, then each agent wants to have at least 4 of its 8 neighbours being of the same type; if it is 0.8, then each agent seeks a neighbourhood with 7 of 8 similar agents. All unsatisfied agents choose to move to a new cell (chosen with uniform probability) that matches their requirements. This process iterates until steady-state is achieved, which can be when everyone is content and no new moves need to be made. However, some situations can lead to a persistent number of unsatisfied agents, though changes to the sum of individual happiness approaches 0.

Even though the model is very simple, it results in a great deal of segregation even with surprisingly small similarity thresholds. Analysis of the model behaviour shows that this is because when a given agent moves to a “better” neighbourhood, they change the neighbourhood that they are leaving, which changes the choice landscape of their previous neighbours. This interaction between agents and their dynamic environments leads to behaviour that is otherwise very difficult to understand without the benefit of hindsight. If there can be very significant differences between expected and actual behaviour with such simple models, it suggests that people can otherwise be quite poor in reasoning through the implications of assumptions in complex systems. Clearly, this model is a drastic simplification of the real world – many more complex ABMs describe more realistic behaviours, including extensions and modifications to this original model – but the original model supported investigation of a specific posited mechanism from which segregation could arise, and provided a means to challenge current thinking. More in-depth analysis of other ABM examples can be found in Sections 2.2 and 2.4.

Comparing and contrasting these two categories of models, we can see how valuable they are when seen as complements. Aggregate and individual methods are both very useful in describing behaviour of

CASs, though often focusing on different aspects. Aggregate SD models excel at describing accumulations and feedbacks, two aspects fundamental to CASs. Their logic is very clear from the visual representation, meaning that non-modellers can still understand much about the model structure and simplifications. They struggle with high levels of population heterogeneity, when spatial context or social networks are important, and in describing individual history dependent behaviour such as learning. IBMs, on the other hand, succeed in very precisely articulating agent heterogeneity, spatial context, and social networks. However, their logic is generally more hidden, meaning that more specialized skills in software design and computer programming are required, though as will be seen, such barriers are being reduced over time. Finally, because of this added heterogeneity and the typical incorporation of stochastic effects, IBMs can require more computational resources. This complementarity suggests that aggregate-individual hybrids might be very powerful. This is discussed in more detail in Sections 4 and 5.

Mabry et al. [64], describe some key benefits of the use of dynamic models in public health. They allow researchers to integrate a broad array of data sources in order to better understand their relationships. They can also suggest what knowledge is missing, and how important it might be. This makes them well suited to suggesting research paths that might be the most valuable. And, through careful experimental analysis, they can help identify important leverage points in the system. Beyond this, researchers can use these models to test interventions or combinations thereof, and explore possible unintended side-effects. Rather than replacing our current tools, dynamic models can play a meaningful role as an additional method of studying complexity.

In spite of these benefits, dynamic models themselves can become quite large and complex. This can cloud a research team’s understanding that models might otherwise clarify. As with other complex software projects, they can require long periods of development, during which project goals might change, leading to further confusion among developers and stakeholders. Changing project goals is also possible since learning is generated over the course of the project as the simulation is run at various levels of complexity. Understanding, then, the role that dynamic models can play in facilitating data gathering, data use, and learning via theory development, and computational experimentation when studying such sparsely observed, dynamically complex systems as public health, *our goal is to employ sequential Monte Carlo methods and modular development patterns to improve our ability to quickly build and learn from the most appropriate model for a given question.*

1.3.1 Contributions

The contributions of this work can be seen in three broad categories. The first describes efforts to advance integration of diverse data sources within dynamic models. Chapter 3 presents the first known work integrating agent-based models with the particle filtering algorithm. This allows a model to be continuously regrounded against incoming data from the real world. While the particle filter does improve the predictive accuracy of an agent-based model in some situations, in others it resulted in poorer accuracy. We outline

some strong possibilities as to why, supporting those hypotheses with experiments. Chapter 5 outlines the use of results from a Discrete Choice Experiment to parameterize agent tobacco purchasing within a model that studies the impacts of interventions on changes to smoking behaviour. We argue that the limitations of the Discrete Choice Experiment are counteracted by the strengths of the dynamic model, thereby strengthening the overall model predictions.

The second category of contribution lies in the use of a modular design pattern within dynamic simulations. To our knowledge, it is the first such research on the use of design patterns focused on improving the design scalability and modularity of dynamic models. The framework is introduced and demonstrated in Chapter 4 using a simple smoking model. In Chapter 5, it is implemented in a richer, more complex model with separate addiction and purchasing behaviour, validation with population data, and pricing interventions. We demonstrate, using examples, how this modular design framework improved and accelerated the development process and the generation of insight from model behaviour.

Lastly, this work includes the construction of a hybrid ABM-SD model. The ABM is event-driven, whereas the SD model is fundamentally a system of first-order differential equations, which is continuous time. Naively integrating these approaches can result in significant increases in computation time. This work explores and compares several computational methods to overcome this limitation. We found that custom numerical methods can improve performance in part because the modeller has access to the time step size. The greatest improvements came from analytic solutions, though an extension was required to accommodate the change of model parameters in time due to the event-driven model logic.

Publications

Chapter 3: Kreuger LK, and Osgood N. “Particle filtering using agent-based transmission models.” In *Winter Simulation Conference (WSC)*, 2015, pp. 737-747. IEEE, 2015.

Chapter 4: Kreuger LK, Qian W, Osgood N, and Choi K. “Agile design meets hybrid models: using modularity to enhance hybrid model design and use.” In *Winter Simulation Conference (WSC)*, 2016, pp. 1428-1438. IEEE, 2016.

Other:

- Daza S. and Kreuger LK. “Examining Siena Model For The Estimation Of Selection And Influence Under Misspecification.” Population Association of America Annual Meeting. Abstract and oral presentation. Chicago, Illinois. 2017.
- Depping AE, Osgood N, and Kreuger K. “They All Look the Same to Me. An Agent Based Simulation of Out-Group Homogeneity. Social, Cultural, and Behavioral Modeling.” SBP-BRiMS 2017. (2017), 6064.

- Kreuger K, Flint R, and Osgood N. “Beyond Drill and Fill: Modeling the Impacts of Risk-based Care on Oral Health Disparities.” *Growing Inequality: Bridging Complex Systems, Health Disparities, and Population Health.* (G.A. Kaplan, A. Diez Roux, S. Galea, and C.P. Simon, eds), Washington, DC: Westphalia Press, 2016.
- Nobari TZ, Kreuger K, Osgood N, Nianogo, R, Whaley, SE, and Wang, MC. “An agent-based model to estimate the impact of increasing affordable housing on obesity risk in early childhood.” American Public Health Association Conference 2016. Abstract and oral presentation. Denver, Colorado. 2016.
- Stanley K, Bell S, Kreuger LK, Bhowmik P, Shojaati N, Elliott A, and Osgood ND. “Opportunistic natural experiments using digital telemetry: a transit disruption case study.” *International Journal of Geographical Information Science* 30, no. 9 (2016): 1853-1872.
- Esfahbod B, Kreuger K, and Osgood N. “Gaming the Social System: A Game Theoretic Examination of Social Influence in Risk Behaviour.” *International Conference on Social Computing, Behavioral-Cultural Modeling, and Prediction.* Springer International Publishing, 2015.

1.3.2 Dissertation Outline

This work is outlined as follows. Chapter 2 presents the background, including a review of relevant literature. It categorizes the various models published in the public health space into 4 broad categories, and uses 6 metrics to compare and contrast the approaches. Chapter 3 describes our published work in the use of the particle filter, a machine learning algorithm, alongside an agent-based model of infectious disease. Chapter 4 presents our publication that motivates and develops a module design pattern, demonstrating a simple example of tobacco addiction. Chapter 5 extends this modular design pattern to a much more mature modelling project, integrating numerous addiction modules alongside the results from a discrete choice experiment we conducted. We end with conclusions and discussion.

All model files can be access through the Git repository addressed in this link:

<https://git.cs.usask.ca/lkk064/kreuger-dissertation-and-model-files>

CHAPTER 2

BACKGROUND AND RELATED WORK

2.1 Background

2.1.1 Motivation

In spite of public health, public policy, and epidemiology being considered as complex and adaptive, the use of a grounded behavioural theory within dynamic models is not universally considered standard. These simulation techniques have their roots outside the public health space, and so suffer from some current limitations specific to this field. Feola and Binder [32] outline some of these limitations to provide reasonable articulation of human behaviour patterns that draw on psychological insight. A recent review of agent-based systems modelling of non-communicable diseases [75] provided many insights into the current dearth of agent-based studies. However, it does not compare any models along behavioural dimensions. Some work has been done recently to address this current lacking. A method for standardizing the description of ABMs, the ODD framework, was recently updated to include a description of the underlying behavioural theory used [74]. But much more is left to be done.

Therefore, we undertake in this work to outline the current state of affairs in simulation modelling of human behaviour as it applies to public health. The remainder of this section outlines the main simulation techniques, how we classified the theories found, and the literature search method. Approaches were placed into 1 of 4 categories, each of which represents a successive section. Finally, in Section 2.6, we provide discussion and synthesis on the main themes of this work.

2.1.2 Classification Dimensions

As this report will outline many simulation models that have been used to describe human behaviour, we require a comparison matrix; a set of pre-defined dimensions with which to compare and contrast each approach. Balke and Gilbert [7] conducted a survey of agent-based behaviour models, and used 5 dimensions relevant to agent behaviour: cognitive, affective, social, normative, and learning.

The cognitive dimension measures “are purely reactive, have some form of deliberation, simple cognitive components or are psychologically or neurologically inspired” [7]. The affective asks if agents are capable of expressing emotions. The social describes any social networks and interactions between the agents, as well

as the ability to represent more complex social concepts like the theory of mind. The normative measures models by how architectures allow agents to reason about formal and informal norms, and whether they can express normative emergence and spread. Finally, the learning dimension focuses on how agents learn about their physical and social environment to be adaptive to these changes.

I include these dimensions but add a few of our own. Firstly, since our focus is on tobacco use in particular, and public health behaviours in general, the role of habit and addiction moves to the fore. It can be argued that habits are a subset of another dimension, such as learning. However, habits and addiction can also be considered to include features of cognition, sociality, and emotion, not to mention genetic and biological factors. Indeed, there is a wide range of definitions for habit. Hodgson [39] argues that habits are “submerged repertoires of potential behaviour” and Becker [9] suggests the definition as simple correlation between past and current behaviour. It appears unclear whether habit and addiction can be completely subsumed by another of the dimensions mentioned so far. Because of this, and the need of approaching questions of public health with this lens, we define this as another dimension. Any explicit mention of habits and addiction for a given framework will be drawn out in this report.

Secondly, we add the dimension of *implementability*. As will be discussed in depth in section 2.6, there is very little discovered overlap between models implemented to address tangible public health questions (those described in section 2.3) and behavioural frameworks built specifically for agent-based models (those described in section 2.2). This is supported by Carey et al. [19] in their review of system science literature in public health. They note that even among the relatively few actual model implementations, validation was rare, saying “If the systems paradigm for policy creation is to be taken seriously, a minimum standard of accountability and repeatability needs to be adhered to by researchers”. This dimension captures what would be required to firmly ground the given framework using proper simulation practices, such as validation, calibration, and sensitivity analyses, as well as the data required to do the foregoing.

Further, we ascribe to the idea that simulation modelling is not primarily meant to provide “crystal ball” predictions, but rather to inform insight into the complex system through iterative model construction. Therefore, we add a 3rd dimension called *scalability*¹. This describes how well the framework might be placed within a greater modelling process, how it might be able to be implemented at various levels of complexity, or how it might be able to provide insight into further data-experiments that could more firmly root its predictions in reality.

2.1.3 Literature Search Method

The purpose of this report is to cover research that intersects with 3 areas: systems simulation, behavioural theories, and public health (specifically tobacco-related work, if possible). Therefore, to find appropriate literature, we conducted a semi-systematic search process, outlined here. First, we searched through Google Scholar and Pub Med for literature that matched various search parameters. I used various combinations of

¹We note that this term is also used in analyzing the time complexity of computational algorithms. The similarity coincidental.

“agent based”, “systems dynamics”, “tobacco”, “simulat” and “addictions”, as well as various MeSH fields in Pub Med, such as the behaviour Major Topic and the “humans” term. This resulted in 66 papers being included.

For each search result we used the titles, and where necessary abstracts, to classify each paper according to the 3 measures above: what type of systems simulation (e.g. agent-based), what area of public health (e.g. smoking, obesity), and what recognized behaviour theory (e.g. theory of planned behaviour). Each paper was then placed in 1 of 4 tiers. Highest tier papers were those that had at least 2 of the following: simulation technique being agent-based, system dynamic, or discrete event; public health area being smoking; and any explicit behavioural framework. Second tier had at least 1 of the above plus something less related in another area (e.g. an ABM model about drug use, since we are interested in tobacco specifically). Tier 3 had 2 or more less related fields, and tier 4 had any single field.

Papers of tier 1 or 2 (54 of the initial 66) were used in a forward and backward citation analysis, using the paper itself and citation records on Google Scholar. Title and abstracts were used to decide if the paper had any of the 3 areas, and were included if they did, and placed into a tier category. During the review, any relevant literature cited in any paper that was not already included was added to the total list. Further, some work that was initially set at tier 1 or 2 was later found to not be relevant to this discussion, and was therefore left out.

2.2 Simulation-based Behavioural Frameworks

The first broad category of models are defined as *simulation-based*. It is predominantly driven by formal logic and algorithmic reasoning, which can lead to a focus on a model’s behavioural complexity. This category has several dominant branches which we discuss: ad hoc, intentional, normative, and cognitive models.

2.2.1 Ad hoc

Andrighetto et al. [3] describe how game-theory simulations often initially aimed at determining what minimal conditions in a complex system lead to certain macroscopic effects. Much learning was generated around this question, but it also led to a growth in arbitrary and ad hoc approaches, and therefore rare focus on validating against real-world data. There are still many reasons to consider the use of such ad hoc methods. They benefit from low data requirements, possess potential for crisp insight-driven models, and have a relatively quick development process.

Risky Sex Game

Tully et al. [99] model the spread of HIV within a population of men who have sex with men (MSM). It uses the idea of risk perception. Agents are paired randomly and decide whether to play the game-theory “risky sex game”. For pairs who decide to play, one agent is chosen as the initiator, who proposes either protected

or unprotected sex. The other agent chooses to accept or counter. The first agent then chooses to accept the final offer or exit the interaction.

When deciding what to offer, agents need to balance the risk of unprotected sex with the fact that it is more pleasurable. Infected agents cannot be reinfected, so are more prone to offer unprotected sex. Therefore, based in part on the offer history of each agent, all agents define their own probability of infection for every other agent (i.e. the more an agent has offered unprotected sex, the more likely they are infected, and vice versa). Combined with utility functions, this belief enables all agents to choose an optimal outcome for any given interaction. The belief gets updated with each offer and counter-offer of the risky sex game. Unless both agents come to the same optimal outcome, they will exit the interaction.

Any cognition in these agents is essentially hard-coded in the utility function. And while it can generate biases for other agents based on experience, it is a learning approach that takes in only a single type of data and generates very simple output. The mark of these game-theoretic models is that agents are as simple as can be while still leading to emergent population-level behaviour.

Reactance

Hammond in [5] discusses the psychological idea of reactance which, briefly, proposes that when people feel their personal freedoms are threatened, they act in opposition to the perceived threat. A simple model is built to demonstrate it. Agents choose to smoke or not based on a utility function - positive means smoking, negative means quitting. The utility function has 3 components, social, physical, and psychological. The physical factor accounts for physical addiction. Social utility is peer pressure (via a weighted average of pro- and anti-smoking sentiments). The psychological factor comes in with outside-network messaging. Those messages are influenced by individual characteristics, such as degree of skepticism, and notably, reactance level. Highly reactant agents respond to anti-smoking campaigns by being more addicted. When these highly reactant agents are spread throughout the network, the effect of the public health messaging is reduced. When those same agents are clustered, however, they can cause the whole network to act reactively, leading to more addiction within that population. This serves to make the point that network structure can have drastic effects on interventions.

Drug Switching

Bobashev and others [16] build a model that looks at drug switching behaviours. They imagine a population that has access to 3 “reinforcers”; meth, heroin, and others (which is a placeholder for non-drug reinforcers). On a clock cycle of 1 month, agents update their preferences for each of the 3 reinforcers. This is done by assigning scores to each drug. The score function combines 6 parameters (liking, availability, peer influence, price, health problems, and punishment) in a simple ratio of products. Learning is captured by updating the liking function using experience, which, importantly, accounts for growth in desire and user burnout. Agents choose the drug with the best score. Then an intervention is tested where one drug is punished more

often for a period of time, during which use of that drug declines. Expectantly, use in another drug increases during this period. However, long after the increased punishment stage is over, it can be observed that the population has become addicted to both drugs. This model gives insight into possible implementation feedbacks.

Self-help Programs

Hiance et al. [38] describe a model that studies dynamics of self-help programs for addicts. The main aim compares programs with lifetime membership requirements (e.g. AA) with others that graduate out. All agents are former users. Each starts out with a sobriety score; the lower the score, the higher the chance of relapse. Lone agents slowly reduce their sobriety score. Agents that are mentored slowly increase their score. Beyond a specified threshold agents can begin mentoring others, where they can increase sobriety in proportion to the number they mentor. Each week, lone agents seek mentors in a given group. If none are found, they move to another group. After 3 tries, they leave the model (and therefore relapse). Groups have behaviour here as well. Beyond a certain size, they split into smaller groups. Three scenarios are run: a large lifetime program, 2 competing lifetime programs, 1 lifetime and 1 graduation-based program competing.

Temporal Difference Learning

Hammond et al. [36] use a reinforcement learning technique called Temporal Difference Learning to “train” agents on the palatability of 2 different types of food, (a: low quality, high palatability and b: high quality, low palatability). Agents move randomly through a space, whose cells have either food a, b, or both. They have a predicted palatability for each food, and choose food that is likely to give them their preferred value. Each time they eat, the accuracy of their perception improves incrementally. The model shows that agents who are “trained” on only 1 food type (by starting in food-a-only space) take a longer to learn the palatability of the other food when moved to a mixed space, as compared with agents that learn in more balanced environments.

Spread of Opinions and Behaviour

Karanfil et al. [48] present a model studying the spread of opinions and obesogenic behaviour through a network. Agents are placed in a social network with directional, weighted ties. The social ties indicate the direction of opinion influence. Agents use the average of their peers (weighted by the relationship weight) to determine their own opinion, but opinions different enough beyond a tolerance are omitted from the average. The behaviour function builds in hysteresis such that it must pass an upper threshold before obesogenic starts, and it must cross a lower threshold for it to stop again. By altering behavioural thresholds, the opinion tolerance, or by adding forced-opinion agents in the model, this model is able to compare several different intervention conditions.

Below-the-skin

There were 2 papers [41, 5] that describe addiction behaviour at the level below the skin, that is to say using an equation-based approach to describe dynamics of addictive substances within agents. Hoffer et al. [41] use a first-order differential equation to describe the growth and decay of drug concentration within drug-addicted agents. Using threshold-based logic, this translates to above-the-skin drug behaviours. Song [5] proposes mathematical relationships for addiction and cessation probabilities based on smoking rates of adolescents, referencing nicotine metabolism, among other factors. While these do not serve as a complete framework for human behaviour, it is an explicit recognition of the impact on behaviour of physiological factors, especially as it relates to addiction.

Summary

Measuring by our dimensions, agents in these models are relatively simple. They may have some degree of cognition, social norms, learning, or habit, but it is generally specific to the modelling problem. This accounting for a subset of possible human behaviours puts these models in danger of missing non-obvious behavioural mechanisms with noticeable effects on outcomes. It also often requires somewhat arbitrary behavioural and intervention mechanisms. However, the data that is required is very minimal. Also, research in the behavioural area under study is often cited showing evidence for the specific behaviour implemented. This enables each of these models generates some piece of useful insight. This modelling approach can be very scalable if not limited to a given model, but instead seen as part of a modelling process with many models built along the way. The focus of these models is often about how simple assumptions scale to populations, which can inform possible public health hypotheses. It also can be a valuable test of the validity of current assumptions.

2.2.2 Intentional

This next section introduces model frameworks that began to emerge out of the game-theoretic modelling practices. These models often focused on the “mechanisms of agency” [3]. An important early step was in the development of agents with “intentionality”, a mental state that is the “basis for reasoning” [7].

BDI Agents

BDI agents are defined with three characteristics; beliefs, desires, and intentions. Their structure was initially designed to model non-human agents within dynamic environments. BDI agents can have many possible simultaneous objectives with many possible actions at each decision point. The environment surrounding the agents is also dynamic and complex. Further, the agents can only sense the environment locally [84]. Figure 2.1 shows the architecture.

A belief is a logical statement of the current condition of the environment when last sensed by the agent

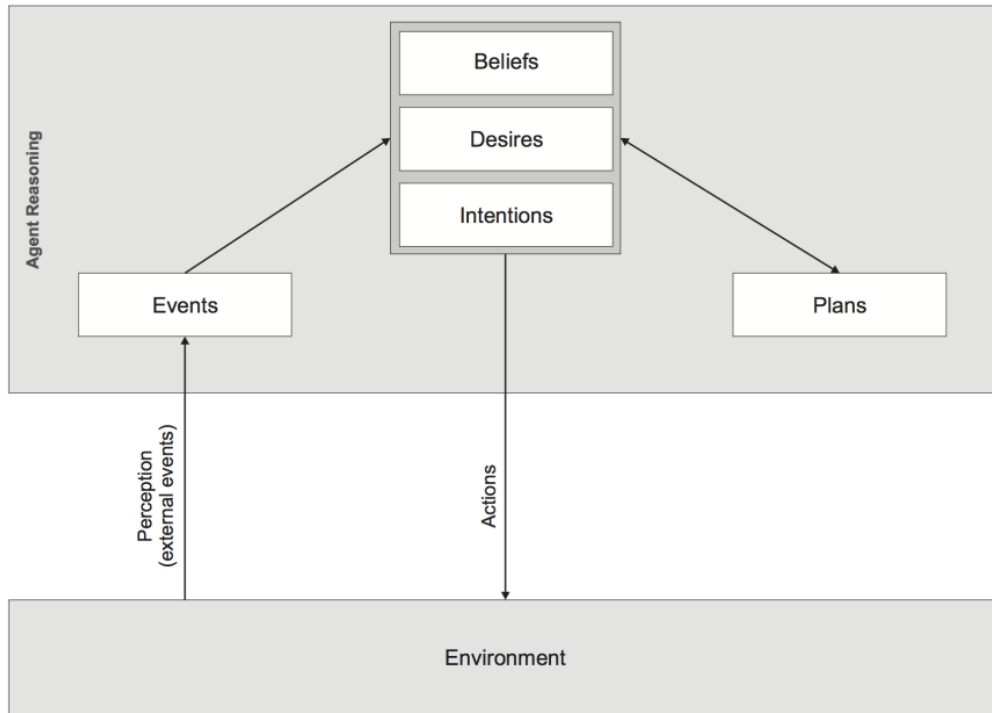


Figure 2.1: BDI design, reproduced with permission from [7]

[7]. It forms the information on the environment to inform the agent’s actions. It is dynamic with each sensing action, and is maintained between each, therefore it also forms a kind of memory [84].

Desires are defined as the knowledge of the priorities and payoffs for various current objectives [84]. These define the possible actions of the agent. There can be many desires, some of which are mutually incompatible, and so the agent will need to decide which desires to pursue.

Intentions are actions that are actively pursued [84]. Essentially, agents store a list of all possible actions. They search through the list to find any whose post-conditions match the current intention and whose pre-conditions match the current beliefs. The agent then chooses the action with the highest payoff.

On the cognitive level, a BDI agent can be reactive and deliberative about its intentions and action plans [7]. There is no focus on normative or affective components, and no specification of agent communication was described [7]. Further, in the original formulation, while agents are allowed only local knowledge, it perfectly represents the environment, therefore the only distinction between the environment and its perception is scope, not validity [31]. In other words, BDI agents are boundedly rational.

Initial implementations of BDI agents were quite technical, and meant for agents to navigate real-time, dynamic, control and management systems. While initially meant for non-human actors, such as air traffic management systems [84], it has nevertheless been used in modelling human evacuation [77], farming cropping plan decisions [93], and in the development of more realistic computer game agents [63]. There has been much work on developing frameworks to account for the above-mentioned limitations, including adding emotions (eBDI), simple norms (BOID), and more complex norms (BRIDGE) (see [7] for an overview of each). Farias

et al. [31] propose a method to address the assumption of perfect knowledge. Others have integrated features of trust [54]. These extensions have been rarely implemented beyond sample simulations.

Furthermore, no work, has been found using BDI (or any of the main extensions) in public health areas, such as epidemiology, health services, or addiction behaviours. The modelling of addiction would be difficult with a pure BDI approach, as addiction invokes emotional and normative elements, uses those elements to bias perception, and requires the agents to possess habitual behaviour.

2.2.3 Normative

Frameworks in this section describe synthetic agents with so-called *norms*. Norms have been discussed by a variety of authors (e.g. [27], [12], [3], [44], and [86]). They are framed as a generally accepted standard of behaviour within a population [27]. They are described in slightly different ways depending on the specialty, such as legal theory, sociology, decision theory, or computer science [44]. But some general themes emerge.

Beheshti and Shukthankar [12] and others outline a norm lifecycle, which includes recognition (when an individual perceives a norm, through passive or active learning [44]), adoption, and compliance or internalization. The process by which norms move along these stages depend on several factors, including the type of enforcement. Dechesne et al. [27] categorize 3 types of norms: legal, social, and private. Legal norms are generally explicitly defined, and enforcement is done through prescribed sanctions. Social norms can be implicit, where agents need to learn by experience, and might be accepted by only a sub-population. A private norm is a norm developed over the course of one’s life. How norms emerge and are enforced (through reward or punishment), what agent features are important, how to handle conflicts between different norms, how agents learn about norms - these are the types of questions investigated in normative research.

Not all normative architectures found were included here. Balke and Gilbert’s survey [7] is a great source for others, including Deliberate Normative Agents (DNA) and the Normative Agent architecture (NoA). Since these were not found mentioned in other literature in this report, we do not include descriptions of their mechanism.

EMIL

The EMIL architecture, overviewed very well in [7], is aimed at describing the normative interaction between an agent and its community. Agents learn about, internalize, and put into practice community norms, thereby also changing them for other agents. It seeks to describe both how norms affect the agent and how they emerge from the agent to the group. Note that within this architecture, norms are defined similar to BDI extensions, where they only serve to restrict goals or intentions, rather than to also prescribe possible actions.

Owing to this focus on norms, EMIL defines 2 types of knowledge - normative and factual (e.g. events), and they are handled separately. Agent have memory for both. Events are remembered on an “event board” and norms are stored in the “normative frame”, which also knows how to infer new norms from experience.

The EMIL agent is structured as in Figure 2.2. It possesses the same elements of BDI, namely normative

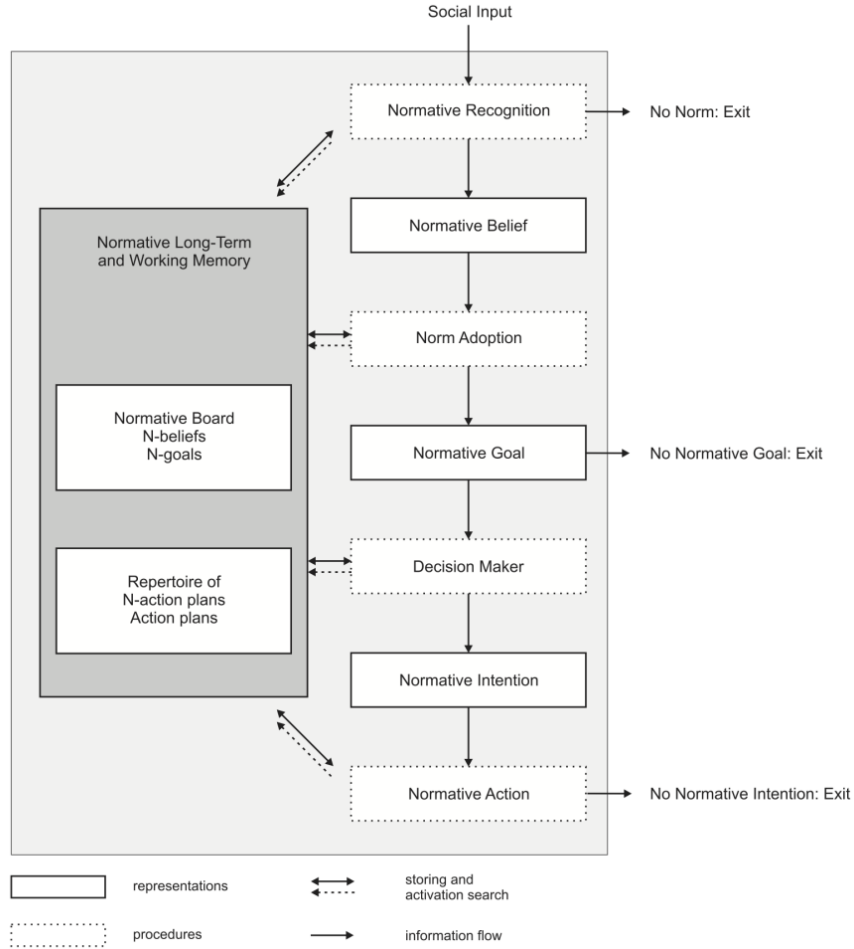


Figure 2.2: EMIL design, reproduced with permission from [7]

beliefs, goals, and intentions. Current normative beliefs and goals are remembered on the normative board. When seeing a new norm through social input, the recognition procedure examines it to determine if it becomes adopted as a normative belief or disregards it. It uses the normative board as well as the new norm’s properties (e.g. norm permissions and obligations, the source of norm stimulus) to decide if this norm will be recognized. Once recognized, the norm becomes a normative belief (the agent believes that it is a norm, not that it agrees with it or will follow it necessarily).

The norm adoption process will search for evidence that the norm conflicts with its current mental state. If there is no such evidence, it has a slight preference to adopt. The decision making process will disregard a norm only if the costs outweigh the benefits i.e. it needs to have a good reason not to follow a norm.

EMIL agents also use both deliberation, through the decision making procedure, and reaction, through shortcuts which allow external stimuli to directly trigger normative behaviour through internalized norms. Agents gain memory and can reason about them. Agents are assumed to undergo a norm learning process, through norm recognition (and deliberation) then adoption. There is no description of emotional responses or behaviours in EMIL. A main feature is the social reasoning and information sharing. Importantly, agents are

able to experience and learn about new norms. Models using EMIL that study human health behaviour have not been found. Many of the uses of EMIL focus on the study of norms from a theoretical perspective. No studies have been found that require the use of public health or particular data sets, as such no quantitative validation is possible.

CafeWilhelmina

Dechesne et al. [27] develop a normative model around smoking built in two phases. They draw from two cultural theories: Hofstede’s cultural dimensions and Schwartz’s cultural values. In both, norms are defined as legal, social, or private. Each agent is given 1 of 3 normative attitudes: lawful agents follow whatever legal norm is in place regardless of the other norms, social agents only follow the action of the majority in a public space, and private agents always follow their private norm. Mid-way through the simulation, a legal norm of “non-smoking” is enacted.

Agents, who are initialized as smoking or non-smoking, meet each other in a cafe. The probability of going to the cafe in any given time step is also set exogenously, and is higher for smokers than non-smokers. They further have a probability to leave, either due to boredom (randomly) or if the prevailing social norm does not match their normative attitude. The prevailing social norm is the defined by majority vote, i.e. the most popular normative position, bearing in mind that each normative agent reaches their conclusion differently.

This model is then extended in the same paper to include smoking behaviour and bar choice (having 3 bars instead of 1 cafe). Agents are also given value dimensions (such as *individualism* or *achievement*), informed from cultural value theory, which guide the behavioural choices of agents, whether agents are smokers or not, how they choose the bar to attend (accounting for the prevailing social norm at each location), whether they will smoke there (accounting for any legal norms), and when they will leave the bar.

Along our dimensions, cognition is reactive, as agents choose norms based completely on static internal values. Emotions are not mentioned. Norms are legal, social, or private, but no norm learning process is described. The only social forces are through norms; there are no social networks or direct communication. Habits are not defined separately from norms. Given that the model is presented in 2 stages, it is implied that some scalability is possible. It also explicitly mentions that data gathering to better parameterize the model is for future work, but that it would likely be possible to use survey data, and some observational studies, to examine whether the theoretical values and cultural dimensions used can be fit to real human behaviour.

LNA

Beheshti and Sukthankar [13] develop the Lightweight Normative Architecture (LNA) as part of a greater effort towards a “general purpose agent-based modelling and simulation system for studying the effects of public policy decisions on a large range of social phenomena”. It was noted that EMIL is quite elaborate

		Node B	
		Smoker	Non-smoker
Node A	Smoker	$ss+\alpha$	sn
	Non-smoker	ns	$nn+\beta$

Figure 2.3: LNA interactions. Agents start as smoker or non-smoker. With each interaction, given their and their partner’s smoking state, they choose the option with the highest payback, calculated according to the equations and their values. α and β are parameters meant to bias maintenance of the current state.

in the cognitive mechanisms, and can be difficult to implement. LNA is hoped to address some of those concerns.

The LNA is described alongside a model on tobacco smoking. It uses the above-mentioned 3 stages of normative learning: recognition, adoption, and compliance. This is captured by giving each agent a *smoking-value* (sv), between 0 and 100, which indicates the agent’s acceptance of the quit-smoking norm. Two thresholds separate this 0 to 100 continuum into the 3 stages of normative learning.

Similar to CafeWilhelmina, agents are given personal values from Schwartz’s cultural value orientations, namely *individualism* and *achievement*, and 3 others (*regret*, *health*, *hedonism*). Agents are also placed in a social network. The diffusion of smoking behaviours through the network is described using game-theoretic interactions, characterized by the interaction table and payoff matrix in Figure 2.3. The payoff values depend on four of the personal values; individualism, achievement, health, and hedonism, as in Equation 2.1 where prime terms are complement (i.e., $100 - x$). Each interaction allows an agent to change their smoking status at a given time in order to maximize the payoff. sv is influenced by the smoking status of an agent and their friends at a given time.

$$\begin{aligned}
 ss &= ind' + ach' + hlt' + hdn \\
 sn &= ind + ach + alt + hdn' \\
 ns &= ind + ach + hlt' + hdn \\
 nn &= ind' + ach' + hlt + hdn'
 \end{aligned}
 \tag{2.1}$$

Agents are then placed within a previously-built transportation ABM, in which agents navigate a university campus based on a daily schedule of classes, socializing, travel home, and eating. The physical environment affects each agent’s sv through observing others smoke, no smoking signs, anti-smoking adver-

tisements, and a miscellaneous category. Each interaction with a physical cue results in a simple Q-learning update step which controls how the agent’s experience influences their learning.

As with CafeWilhelmina above, cognition is purely reactive, and affectivity is not mentioned. Norms are captured, though legal norms are ignored. Social norms emerge from game-theory interactions and a physical environment Q-learning process, following a norm internalization process. Private norms are captured indirectly through the specification of values. A social network is used only to arbitrate the game-theory interactions, and to calculate the *sv* of each agent. For scalability, it is notable that the behavioural research was undertaken with a previously-built ABM. As to implementation, this study uses data from a survey built to gather estimates of the various value parameters for the same population which was captured in the transportation model.

CSL

The authors of LNA also built a Cognitive Social Learners (CSL) model [12], which was intended as a general purpose normative model. CSL agents are constructed on the BDI approach. Each belief, desire, and intention is associated with a given uncertainty. Instead of using intentions alone to decide the action, agents use a social dilemma game, whose payoff matrix depends on the highest certainty intention. For example, if an agent needs to decide whether to litter or not, they play a social-dilemma game with a nearby agent using their most certain intention. The outcome from this game is the action the agent will undertake. The results of this action (e.g. changes in their happiness or reputation, or paying a fine) then allow the agent to update their certainty values for each of the norms.

Q-learning is implemented to manage how agents learn from their experience. Each time an agent acts, it then observes the rewards or sanction, which could come from the environment (e.g. a dirtier environment if the agent litters), an enforcing agent by means of a fine, or changes in social status through reputation. Beliefs, desires, and norms within the agent evolve through this learning process.

The underlying BDI logic provides cognition around which action to choose. Arguably, agents only have the deliberative frame. CSL agents have normative reasoning explicitly captured, as well as learning in that they use the effects of their actions to update their confidence in beliefs and norms. Affectivity is not described. Explicit social networks are not mentioned except as it carries the social learning game interactions, though agents are able to impact each other through observation and punishment. Habits are only viewed as highly certain beliefs, desires, and norms (as intentions rise from these). There is no special handling of addiction.

As with LNA, this work draws on survey data to assess normative influence and actions. In fact, this publication compares CSL, LNA, and another normative architecture in the ability to predict both the number of smokers from a data set, and the perceived social unacceptability of smoking. From a scalability perspective, CSL is presented in this work on two different ABMs, by changing the initial norms and the initial prescribed beliefs. One looks at park usage and littering, and the other at tobacco smoking.

2.2.4 Cognitive

These models attempt to account for the cognitive element of human decision making. Cognition can be described as anywhere from reactive to deliberative behaviour. As well, agents can take on various non-rational behaviours, such as norms or emotions. While other models in this report account for these behaviours, models in this section are inspired by cognitive ideas, seeking to explicitly define cognitive elements.

PECS

The primary motivation for PECS was to address the assumption in BDI-like approaches that humans are “purely rational decision makers” [100], that human decision making is not only influenced by rationality but other factors as well. It specifically encourages the use of thinking frameworks that allow for deliberative as well as reactive decision making [65, 7]. The core mental structures of PECS provides the acronym: Physis (physical or material properties), Emotion, Cognition and Social status [17].

It might be more correct to think of PECS as a behavioural design pattern rather than a behavioural model. Indeed, the authors themselves indicate that it is meant to integrate “a variety of theories at various levels of complexity”, and the intended use of PECS is to “support the design process” of models focused on human behaviour [100]. It was built specifically to capture human behaviour within individual-based simulation models.

However, it is not necessary to describe all the PECS components in a model [65]. For example, the prototype model built by PECS’s authors [100] in the framing of the approach only include cognition and a social/emotional measure. As the modeller needs to define the model scope, the PECS approach is intended to accommodate varying levels of behavioural complexity. This scalability is important to the framework, reinforcing the notion that PECS is not a behaviour model but a behavioural design pattern.

To accomplish this, the design of PECS follows a component-oriented hierarchical approach [100]. This assumes that complex agent behaviour can be described using smaller, nested, simpler, model components that interact. The design of a PECS model is in choosing which components are necessary at a given level of complexity, and defining the nature of the interactions between them. This is intended to give the framework a great deal of flexibility with different behavioural theories.

Figure 2.4 outlines the elements of PECS. The PECS agent is seen as observing the world through the 1st layer, which contains the sensor to the world, and an internal perception thereof. The perception component can contain information filtering mechanisms, perceptual processes, or observations of social interactions.

The 2nd layer contains the actual PECS components. Physis captures the physical reality of the agent, such as demographic, age, or disease state. Emotions can also be captured in the relevant component (which could include reactive features). The cognition component models the knowledge and thinking of the agent (i.e. deliberation), which could include learning and deliberative processes as needed. Any social roles, needs, and responsibilities are captured in the Social Status component. Each of the PECS components are

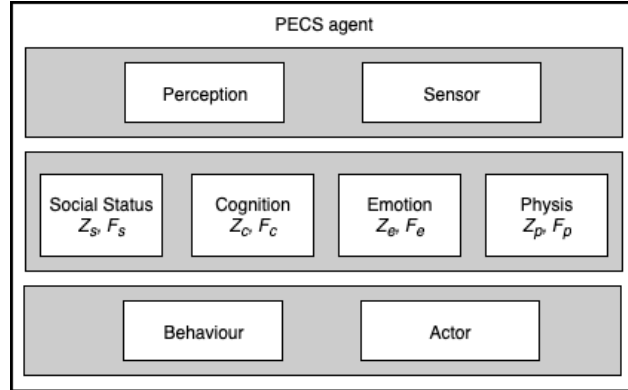


Figure 2.4: Internal representation of a PECS agent. There are 3 layers. The 1st layer, the sensor and the perception, observes the world. The 2nd layer is the internal representation of the agent, captured by the PECS components. Z and F indicate the internal state and transition functions for each PECS component, respectively. The 3rd layer, the behaviour and actor blocks, captures action with the world.

characterized by an internal state or memory (Z), and transition functions (F), as indicated in the figure.

The last layer captures the selection and implementation of actions with the world. An agent might have a list of possible actions, which are stored in the behaviour block, along with the relevant logic needed select any number of actions to take. These instructions for action are sent to the actor agent, which contains the knowledge of how to implement a given action.

All of the papers we reference which describe PECS present a diagram depicting the relationships between the various components and with the external environment. They include the distinction between causal dependencies and information sharing. There is no significant discussion, however, on this structure. For example, there is a proposed causal link from perception to both social status and cognition in [100], but this is not present in [17]. This could be a reproduction error, but it is not clear. There is also no discussion of why perception does not directly cause emotion changes. Even though the authors indicate that this approach is theory-ambivalent, there are clearly assumptions made.

As a tool to aid in designing a behavioural module for agents, it seems to be potentially valuable to map the space of necessary features. However, it has only rarely been implemented, and never, as far as we can find, with all features described. As mentioned above, the original authors describe a prototype model that uses only 2 PECS components.

The authors in [17] developed another partial PECS implementation. Here it is integrated with an implementation of the Health Belief Model, discussed below, and is aimed at predicting compliance with screening for diabetic retinopathy. The first layer (sensor and perception) is not included. The 4 elements of layer two (physical, emotional, cognitive, social) each influence different important blocks of the Health Belief Model (HBM - described below) but are static for a given individual. The PECS elements are then used in the final equation to determine probability of attending. The authors noted that the mathematical relationships were plausible but chosen arbitrarily.

A much more complete ABM was built by [65]. It describes burglary dynamics in a simulated city based on various parameters, such as demography, city layout, and wealth status. It also uses only a subset of the full PECS architecture through agent needs; the need for wealth and the need for sleep. Because agents who burgle have memory, and deliberate on where to burgle next (based on houses the agent has observed and on parameters of each potential target) this can be thought of as a cognitive component. The need to sleep might be considered a physical component. Intensity functions calculate the strength of each need over time. The need with the highest intensity drives the action that agent undertakes (if wealth, then either work or burgle, otherwise sleep). Agents who are not able to satisfy the need for wealth with employment do so by burgling homes. Agent heterogeneity and personality is captured by altering the rate of state variable change.

Given the few available implementations of PECS, when addressing the strengths along each of the dimensions, we are forced to be hypothetical. However, there is still much to value. PECS specifies a component for each of the cognitive, affective, and social dimensions. Schmidt [89], for example, describes in some detail relationships between cognitive and emotional factors. Along the addiction dimension, it is certainly plausible to assume that it depends at least on physical, social, and emotional dimensions. Addiction, then, might be treated as resulting from the combined states of all 4 core components, leaving the agent choices up to the behaviour block.

Balke and Gilbert [7] value PECS weakly in the normative and learning sections. Certainly, norms are not explicitly mentioned in the development of PECS. And no implementations take strides in adding norms or mature learning features. But since agents can reason, and are explicitly given memory, this would be readily accessible by a PECS architecture. Similarly for norms, PECS agents are able to take in environmental data relating to social network, update their social status, take action under different changes to their status, and again, retain memories. Information sharing is indicated in the PECS figures, which can be used to communicate between agents as well. Finally, the update functions need not be static. One could imagine placing an Hidden Markov Model or a Bayesian update function in place, though this is outside the initial spec.

For implementability, the modular approach of PECS could allow a very rich description of human behaviour. The design assumption of modularity requires that each of the 4 components be treated independently. Approaches that reduce human behaviour along different fundamental components would need to be translated into a PECS framework before they could be modelled in this way. PECS has the advantage of creating a common language and framework with which to compare and contrast human behaviour models and algorithms. It could also help in describing and sharing models, or even in their categorization. Further, PECS specifically encourages incremental scaling of complexity, and this can be of tremendous support when constructing a complex behavioural model.




Figure 2.5: Consumat heuristics, reproduced from [7], in relation to when they are invoked. Note that this depends not only on the problem, but on the agent. Details of the heuristics are not discussed here. See [7] for a more comprehensive discussion.

Consumat

Built initially for consumer and market behaviour (such as in technology diffusion studies [47]), and summarized well by Balke and Gilbert [7], Consumat focuses on 3 elements of human decision making; human needs, cognitive time required to make a decision, and uncertainty. Humans have many needs, some that might conflict, which are dynamic over time (needs are reduced with consumption, and increased with avoidance). It is argued [47] that a single utility value, for example, is too limited to capture this complexity. The behavioural theories used are Maslow’s pyramid of needs, and work by Max-Neef. From these, Consumat describes 3 main need types for agents - personal, social, and status. (Status describes how agents might get value possessing more of something than their neighbours.)

Agents optimize over available options which allows them the complexity of balancing these many, potentially conflicting needs. Consumat describes agents that attempt a further optimization step, but one to their decision making. Heuristic methods are used to reduce cognitive effort and time. Six different heuristics are proposed (Figure 2.5) which depends on the level of uncertainty and the amount of cognitive effort agents are willing to make (e.g. agents with a low level a need satisfaction will make a larger effort). Memory is captured with a mental map that stores information on abilities, opportunities and agent characteristics, and is updated based on experience.

Along the cognitive dimension, Consumat provides a very articulate process, allowing for varying levels of cognitive engagement and different cognitive pathways. Emotions and social norms are not directly mentioned. And while agents are not typically designed to see much beyond their own success, they are able in principle to reason about the social environment. No discussion of habits or addictions was seen. It has not been implemented in a model on health behaviour.

Consumat is described as a framework capable of generalizing over many different human behaviour domains [47]. This could provide a great deal of flexibility and potential scalability if the behavioural question can itself be scaled in complexity. It might be possible to include many behaviours not initially envisioned, but this comes at the cost of requiring more parameters, and therefore, more work in calibrating and validating the model.

2.3 Social-cognitive

In the construction of a dynamic simulation focusing on human behaviour, drawing on an independently validated behavioural model is useful. A number of behavioural theories, defined outside of simulation modelling, have been implemented in various ways within simulations. Any behavioural models that appear in the review literature are described below, including their strengths and limitations.

Brailsford and Schmidt [17] indicate two challenges to using these theories in dynamic simulations. Firstly, they are meant as explanatory tools rather than predictive ones, looking more for observable antecedents to behaviour rather than necessarily the causal roots. This focus on correlation over causation injects into these approaches a sensitivity to behaviour outside of the data. Dynamic simulations often involve intervention studies and sensitivity analyses, which can push agent behaviour to areas not captured by data.

Secondly, there is difficulty in translating these theories into the quantified simulation model. The traditional validation tool with these models is the survey, with each element being given a score based on specific questions. Saying, for example, that subjective norms (and therefore friendships) impacts a person's intention towards a given possible behaviour does not narrow down the possible equational relationships between the two, or provide insight into what features of the social network are most important. Therefore, we also provide discussion on the various ways each have been implemented.

The two most dominant theories that have been implemented or described are the Health Belief Model and the Theory of Planned Behaviour. A few theories were mentioned by some authors but not developed into simulation models. I mention them only briefly here, before we describe the two theories above in more detail.

The Health Locus of Control (described in [17] and [94]) depicts 2 psychological stances. With an internal locus of control, an individual believes they are in control of their own actions, with an external, factors outside their control are dominant. External locus of control further divides into "powerful others" and "fate". According to [17], it is a weak predictor of health behaviour.

The Elaboration Likelihood Model (ELM) (described in [90]) is a theory of attitude change. It proposes that individuals use one of 2 routes when changing an attitude: the central and the peripheral. The central route is marked by thoughtful consideration and is relatively permanent. The peripheral route requires only simple inference and results in relatively temporary changes. ELM describes the factors that lead an individual to taking one route or the other, such as motivation to pay attention to the message, and the ability to understand the message.

The Stages of Change Theory (as outlined in [90]) is a theory used by psychologists to examine patient progression through therapy. It depicts 5 stages of progression through which individuals move: precontemplation, contemplation, preparation for action, action, and maintenance. It describes features of each stage so psychologists can have a measure of where patients might lie. It is a descriptive theory without causal connections between the stages.

Feola and Binder [32] outline the *integrative agent-centred framework*, an integrative research approach designed for use with farmer behaviour in agricultural and social-ecological systems. It uses Giddens' Structuration Theory, which focuses on the mutual causal relationship between individuals and the societies to which they belong, and the Theory of Interpersonal Behaviour, a "psychological framework which aims at explaining individuals' 'interpersonal' behaviour". As a whole, it accommodates social norms, affect, habit, social roles, and intention. It also describes its features as modular, providing for possible incremental development and good scaling. While Feola and Binder explicitly mention that it could be used to inform an agent-based modelling program, its primary purpose is as a research approach, and no simulation studies using this framework were cited or have been found in the literature.

2.3.1 Health Belief Model

The Health Belief Model (HBM) is the oldest and most widely used of the social cognitive models [18]. Arguably, it was first used in 1958 in a study on compliance with tuberculosis X-ray screening. Abraham and Sheeran [24] provide a thorough overview of the theory, including a historical context for the HBM, saying that while connections between demographics and health behaviour were known, researchers and health professionals wanted to be able to use education to help modify those behaviours. Therefore, psychological models were sought which could connect antecedents to actual behaviour that were potentially modifiable through educational interventions while still accounting for differing socialization and demographics. Beliefs were thought to be these antecedents.

The HBM, shown in Figure 2.6, sees health behaviour as coming from 2 main drivers: perception of health threat, and evaluation of behaviours that might counteract those threats [94]. Each is further separated into 2 beliefs. Threat comes from perceived susceptibility to the health issue, and the anticipated severity. Behavioural evaluation consists of the 2 beliefs of perceived benefits of the behaviour and perceived barriers in doing the behaviour. Beyond these, the HBM recognizes the importance of cues to action, which are external triggers such as social influence or educational campaigns. Later versions of the model include a general "health motivation" or self-efficacy feature as well which captures the confidence in one's ability to take action.

Validation and operationalization of the HBM in non-simulation models is typically done using surveys, where questions are defined that tap each of the relevant constructs. Abraham and Sheeran [24] provide an in-depth overview of the validation process, including how to find questions, how to test them, and how to determine internal validity and reproducibility. Once a survey design is set, and a population is available for study, statistical tools are used to see how well the questions representing the relevant constructs are able to predict the variation in health behaviour.

Survey validation of the entire HBM has not been done [24]. This is in part because, while the terms used in HBM are easily understandable by non-specialists, they are difficult to quantify, and the relationships and interactions between each are not clearly defined. As a result, most of the work using the HBM has

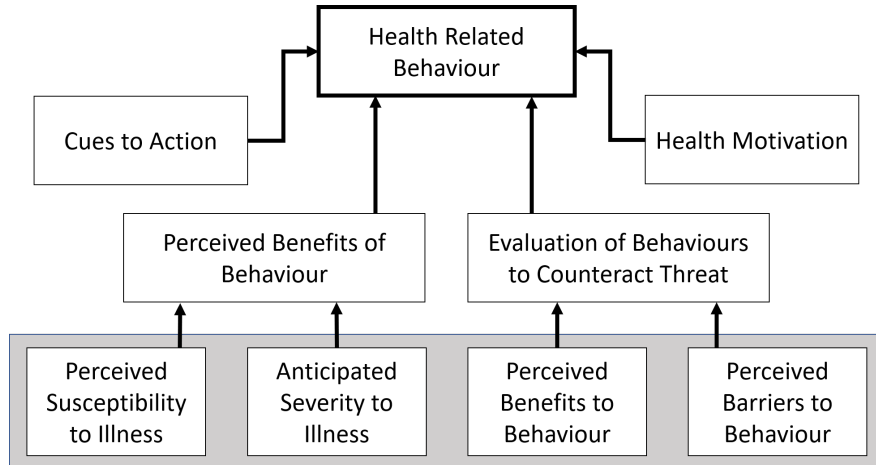


Figure 2.6: Formulation of the Health Belief Model. The 4 elements along the bottom represent the various perceptions. Cues to action represent external triggers to action. Health motivation captures confidence in one’s ability to act.

focused only components. The 4 main beliefs of HBM have been survey-validated independently, to varying degrees of success. Less often are cues to action empirically tested, and health motivation has only been simply captured with a single survey question [24]. And because of the lack of formal specifications for how to implement the HBC, different attempts might not be strictly comparable [24].

Beyond the difficulty in operationalization, there are a number of other limitations. It does not account for social networks, peer influence, or social norms, except possibly through the cues to action component [90, 71]. Economic and physical environment factors are only possible to capture in the HBM through the perceived barriers component [90]. More direct causal pathways are not described. Rarely is it used to study past behaviour or addiction [24]. Intention to act has also been identified as perhaps a more important direct driver of behaviour, captured in other theories (e.g. Theory of Planned Behaviour) but not well in the HBM [24]. Finally, according to [24], most of the studies have been cross-sectional, so there is uncertainty regarding the causal direction of belief and behaviour.

The HBM specifies types of beliefs and how they might affect health behaviour. Since it is a descriptive model of behaviour rather than a predictive one, it does not actually form a cognitive model. Implementations need to specify this component separately. It does not explicitly capture emotions, though one could imagine that they would alter the degree to which perceived susceptibility, for example, affects a given individual’s health behaviour. As mentioned above, the social aspect is only minimally captured, similarly with social norms. Cues to action and perceived benefits/barriers to behaviour might be able to carry minimal information about each. Learning is not mentioned, and neither is habitual or addictive behaviour.

Regarding implementation in simulation models, only 1 publication we read does so [18]. It is a hybrid between PECS and the HBM meant to predict probability of attending diabetic retinopathy screening. Each PECS element informed a subset of the HBM framework, which then influenced health behaviour with plausible, but arbitrary, causal equations. It also used artificial data. The conclusion from this work is that

the HBM is not ideal for use in simulation modelling, in large part because of the difficulty in quantifying relationships.

2.3.2 Theory of Planned Behaviour

The Theory of Planned Behaviour (TPB) was developed by Ajzen [1] as an extension to the Theory of Reason Action by Fishbein and Ajzen. The theory describes intention as the necessary precursor to action. This intention comes from three sources: attitudes towards the behaviour (BA), subjective norms about the behaviour (SN), and the individual's perceived behaviour control regarding the behaviour (PBC). Ajzen defined each of these three aspects to be a sum of products, as depicted in Figure 2.7. Attitudes are composed of beliefs and evaluations about the outcome of a behaviour. Subjective norms are not simply the normative beliefs, but also the motivation to comply with those beliefs. And finally, perceived behavioural control is composed of the control belief, defined by [18] as the likelihood of occurrence, and the perceived power of the facilitator or inhibitor (e.g. How much does the poor weather reduce my ability to jog outside?).

Even though the 3 main drivers of intent are defined, including a possible mathematical formulation (as visible in the Figure 2.7), authors in [51] critique the TPB for being difficult to define because it can admit theoretically infinite decisive factors and beliefs. Any number of multiplicative pairs can be included. This makes quantifying the TPB potentially challenging.

The TPB was not developed for use in health behaviour, so it does not capture health threat [17]. It does not capture all necessary features of adolescent behaviour as well. Authors in [23] describe how models with cognitive and non-cognitive components are necessary given that psycho-social maturation continues to develop after cognitive abilities are largely developed. Finally, the subject of addiction includes non-cognitive features, which are not captured by the TPB.

Of the literature covered by this report, 5 publications made use of the TPB in ABMs [53, 52, 106, 76, 49], with [18] representing the only discrete-event simulation (DES) model. Only 2 [49, 18] of these works relied on survey data to calculate TPB parameters, with the others relying on simplified ad-hoc approximations that lack behavioural validation.

In [49], the authors build an ABM to study the potential adoption of organic farming practices in two new EU member states. The TPB was used to describe the intention of the farmer agents, and was calculated with a weighted sum over the 3 attributes of the TPB. If an agent's intention was above a threshold, they adopted the new practice. When it was below, they reverted. Surveys were conducted with farmers from each state, and various regression analyses were done to determine weights and initial TPB attributes values. To accommodate social interactions, the authors used the Relative Agreement (RA) model [69]. The RA model allows agents to interact and influence each other's views. Importantly, it uses the idea of uncertainty. Someone with a strong belief that is very certain is likely to strongly influence those who consider that view a possibility. Someone with a less certain view will be less influential, and someone whose views are too distant will not be influenced at all.

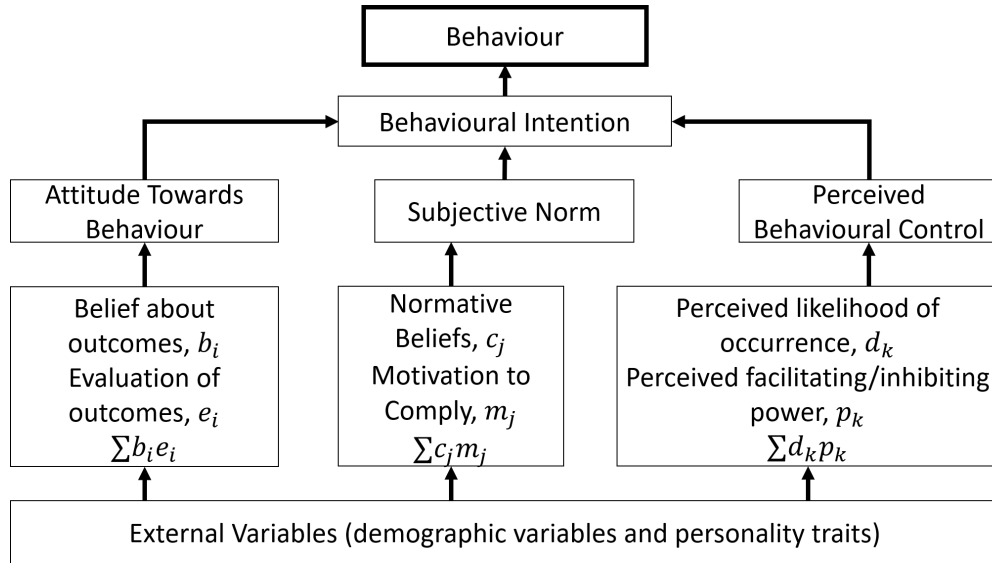


Figure 2.7: The Theory of Planned Behaviour, including the functional form of elements. Intention is the primary driver of behaviour.

In [106] and [76], the authors build an ABM to study smart-metering adoption. Behavioural Attitudes (BA) for a given energy option are calculated solely from pricing information for that option and pricing sensitivity of a given agent. Subjective Norms (SN) are determined through a weighted sum of opinions from all members of their social networks, though how this opinion or the respective weight is calculated is not clear. Perceived Behavioural Control (PBC) is defined as relating totally to energy and technology policies, and are seen to be universal to all agents in the model, and therefore appear to not be included explicitly.

In [53] behavioural intention is defined as the product of BA, SN, and PBC. BA is defined as the probability of an agent migrating, under various conditions (calculated as the fraction of individuals who did migrate). SN is defined as a function, fit from data in the study, taking the opinions of each of the peers in an agent’s social network. It is not clear if opinions are intentions or some other measure. PBC is set as 0 or 1 based on the probability calculated by summing an “asset rate” and an “experience rate” determined from survey data. It should be noted that this paper differs greatly from others, and from the definition of TPB in [1]. Importantly, BA is defined based on completed decisions rather than subjective attitudes. TPB makes the assumption that attitudes inform behaviours but this article appears to equate them during this calculation. The model in [52] appears very similar, though it is not described well enough to know for sure.

In [18], the authors describe work that uses survey data to inform the values for BA, SN, and PBC. These were used to generate a probability of attendance at breast screening visits. Being a DES, the model from this work is meant to study the process of breast screening programs, and so there is no other focus on social networks or behavioural prediction aside from attendance.

Along our dimensions, no mechanism for cognition is proposed, but since agents always act on their intentions, we can classify these agents as reactive. Emotions are not captured explicitly. Norms are described

as arising from externally defined normative beliefs (which could come from friends, or other sources), but there is no discussion of a norm adoption process or of the different types of norms. They are simply enumerated for a given question, with survey data serving to quantify the magnitude of each in a given situation. The social aspect can be captured if surveys ask questions about social networks and the resulting effects. The TPB is without feedbacks or dynamic behaviour, so no learning is captured. Habits could be captured from norms or agent beliefs, but nothing specific about addiction is defined. The TPB can be connected with survey data to validate behavioural models before being simulated, though this is not always done. Scaling the complexity of the TPB for different situations would likely include the level of survey complexity, but little discussion on this has been seen.

2.4 Observables

The six papers in this category are still based on individual actors, however the focus is on the non-latent (i.e. observable) aspects of their behaviour. They are all focused on tobacco control policies and use projections. Generally, this focus on observables means that there is little to no proposed internal mechanism for cognition, emotion, or norms. Any mechanism that is proposed (e.g. Markov chain) is simply meant to capture features of the data or other observed behaviour. For these reasons, it is also hard to discuss agent learning, since all learning comes explicitly from contributing data sources. Depending on the paper, these may include social networks informed by network data. Their use of data to calibrate means a wealth of experience in model validation, something not seen in the mechanistic models. No conversation on scaling model complexity is seen with most work here being a one-shot attempt at fitting model output to data.

2.4.1 Static Decision Tree

Authors in [70] use a static decision tree, in Figure 2.8, for a single agent. The goal is to model the use pattern from non-user to either cigarette, smokeless user, or eventually, dual user. Cigarette users are further categorized as stable, health concerned, in smokefree environments, and price sensitive. Each of these groups has a unique chance to quit or remain users.

Each of the transition probabilities are associated with a standard error, and are assumed to follow a normal distribution. Probabilities for initiation are gathered from survey data on the US population, asking what proportion had ever smoked more than 100 cigarettes in their lifetime. The size of each of the 4 initiator groups are estimated from survey data on why people had quit smoking. For each group, quitting probabilities are estimated by being adjusted to match observed tobacco use behaviours in the general population.

The research question is to understand how the use of a lower-risk tobacco option (smokeless) might change the health outcomes, and therefore each of the final smoking-state numbers on the right of Figure 2.8 are associated with a health outcome score. Multiple scenarios are chosen, which adjusts the transition probabilities assuming, for example, a very successful campaign to increase knowledge of the smokeless option.

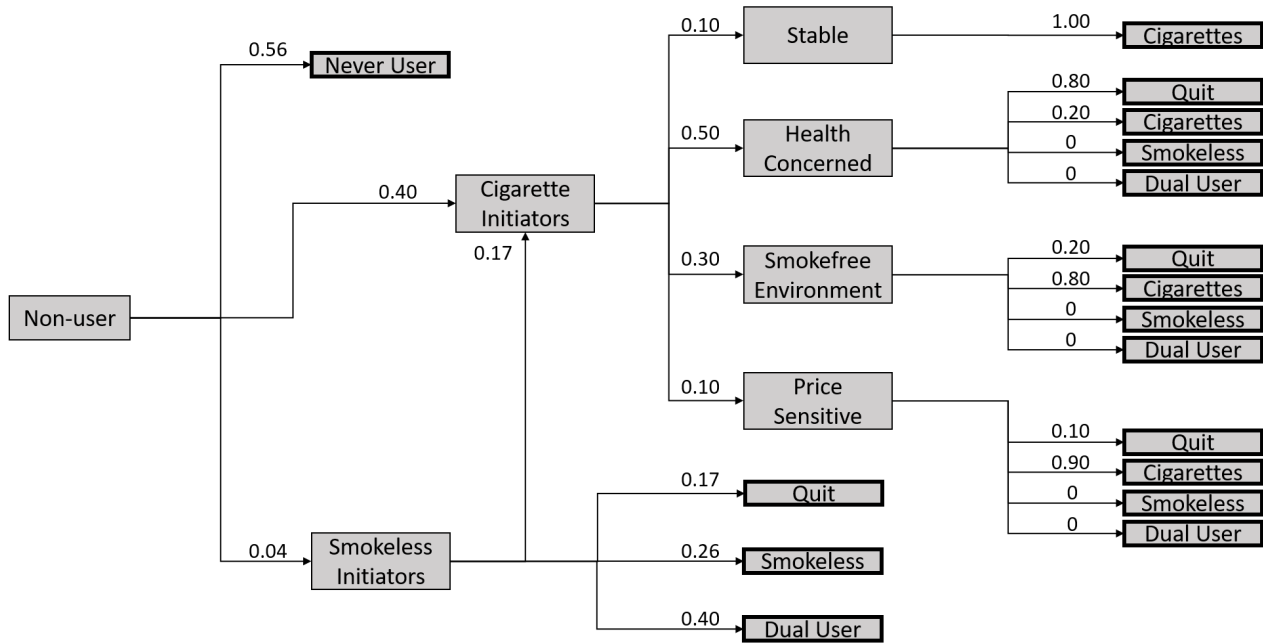


Figure 2.8: Static Decision Tree from [70]. Transition probabilities are shown on the arrows. There are 5 unique observable behaviour states: never user, cigarette user, smokeless user, dual user, and quit user.

And then, since probabilities are estimated and assumed to be drawn from a normal distribution, a Monte Carlo simulation is done to determine the size of each outcome category.

The decision tree approach requires the assumption of mutually exclusive groups [70]. While this can be overcome by adding states for each combination, this dramatically increases the number of states (e.g. all combinations of the 3 atomic quitting reasons generates 4 extra states). Further, this approach does not capture more nuanced combinations e.g. users whose primary concern is health, but who, to a lesser extent, accounted for the environment. For these reasons, this research assumes users only have a single reason to quit. Further, due to focus on observables, there is no independent focus on cognition, emotion, social, normative, learning, or habit formation. However, it is more easily implementable than models requiring specification of latent features since all the data that is necessary is estimated from population-level and survey-based data. The lack of complex systems features reduces the capacity of this model to describe feedbacks or emergence, but the ease of implementation might allow this approach to either inform a more complex modelling approach, serve to calibrate such a model, or be used to compare model outputs to see possible failure modes of the complex systems model.

2.4.2 Dynamic Behaviour

There are also several approaches that use a dynamic population while focusing on observable behaviour. These mostly use static Markov chains to describe individual histories and projections through several behaviour paths. A Markov chain is a set of mutually exclusive states (e.g. never smoker, current smoker,




Figure 2.9: Markov chain used by Killeen [50], where possible agent states are: u for use, w for withdrawal, or a for abstention. Transitions between states are fit to the data.

former smoker) with a set of transition probabilities. In a time-dynamic model, these probabilities are per a unit measure of time. Importantly, Markov chains are history-independent, meaning that the transition probabilities do not depend on previous states. A result of this is that statistics on a given agent’s state history need to be recorded separately (e.g. as counts of how often agents have occupied a given state, or how many times a certain transition has been used).

An example is [50], where agents are all assumed to be smokers. Figure 2.9 shows the Markov chain for each agent, either habitual user, recent quitter experiencing withdrawal, or long-term quitter without withdrawal symptoms. Agents in this model are not differentiated by age, sex, or severity of dependence. This approach can capture some aspect of agent psychology since the transition probabilities are defined per individual. As outlined by the authors of this study, the quit attempt rate “reflects personal and social forces, including public health or legislative intervention”. The rate of returning to regular use is associated with the physical severity of withdrawal. Transitioning to long-term abstention is impacted by “physiological renormalization”, another physical trait. The rate of returning from long-term to withdrawal symptoms depends on many social and psychological factors.

The habit dimension is captured explicitly in this approach through the Markov states. Focusing on observables only, it is hard to separate the cognitive, social, emotional, normative and learning dimensions. The benefit is that population level data, with demographical variation and some distributional assumptions would be sufficient to parameterize and calibrate a model.

Verzi et al. [102] describe the population structure model, a dynamic population model with net migration and births. It also uses two separate Markov chains, one for smoking status (never smoker, current smoker, former smoker), and one for health state (alive or dead). Age and gender of the agents are also individually defined, and impact health and smoking probabilities.

Data to validate and calibrate the model is mostly taken from census data, though other national survey data is used as well. The primary purpose of this research is to demonstrate the ability of the population structure model to forecast trends of smoking and smoking-related mortality over time. As with Killeen [50], psychological traits do impact the transition probabilities of the smoking Markov chain, however they are not examined independently. And while this work introduces a dynamic population, no social network is used.

Other work has been with variations on this theme. Bachand and Sulsk [6] developed a model that

is parameterizable by a user. It describes dynamics around alternative product use and counter-factual intervention scenarios. Agents are categorized into age groups, can switch to and from either of two products (or both). Users are able to specify relative-risk values for mortality estimates. While some transition probabilities are fixed, others (specifically those estimated from the literature) are normally distributed. Markov chain Monte Carlo techniques were used to estimate output variability.

2.4.3 Activity-Based Modelling

Beckman et al. [10] describes a method for generating a synthetic population, each individual with activities (such as shopping, socializing, and attending work or school). This study focuses on youth aged 12-18 and uses data from the National Longitudinal Study of Adolescent to Adult Health (Add Health), which is a series of surveys done in several US high schools that asks questions on smoking behaviour, social networks, and parental smoking. There are currently 4 longitudinal waves and a 5th one planned. In the model, agents have 6 smoking states, each corresponding to a range of daily smoking rates (i.e. 0, 1-5, 6-10, 11-20, 20-40, >40 cigarettes per day). Each agent is further characterized by age, gender, number of parents who smoke, and number of friends who smoke.

Initial states for each agent are determined by doing a Poisson regression model for each age, parameterized by age, gender and number of parents who smoke. When determining transition probabilities, the authors use the 1st 2 waves of the Add Health data to observe behaviour changes. Poisson regression models are used between each of the states, which generate yearly probability distribution based on age, gender, parents who smoke and number of friends who smoke. Then, each day a coin flip is made with a probability of 1.25%, calculated so that once every year with a probability of 99%, the yearly transition model is used.

What makes this project interesting from a behavioural perspective is the use of activity-based modelling. In the first of 4 steps, synthetic individuals are generated with “a complete range of demographic attributes collected from the census data”. Importantly, they have interactions with other people and with locations in the area (e.g. for shopping, work, school etc.). Individuals are then placed in households, maintaining statistical indistinguishability from the data.

In step 2, a set of activity templates is determined at the household level using surveys accessing activity patterns. Each activity template indicates what activities at what times each household member will carry out, such as eating, socializing, and shopping. Step 3 uses observed land-use patterns, tax data, and other sources to assign to each house a location. This is done with the *gravity model* which, in the example of assigning a work location, generates a weight for all work locations in the model relative to the house location. This process is iterative so as to avoid, for example, an agent needing to cross the city during rush hour in a few minutes.

Step 4 generates a social contact network using a labelled dynamic bipartite graph of people and locations, and a synthetic friendship networks using a learning algorithm over dendrograms, more fully described in the paper. This completes the process of generating a synthetic ground truth population.

While this research does not test intervention scenarios, it outlines how this could be done. If, for example, one wanted to know the effect of an anti-smoking protocol in a given school, one could change the probability of smoking for that particular school. Because the adolescent population is “embedded in a larger synthetic population”, the individuals might start smoking outside the school instead. As stated by the authors, “the system is designed to function in a policy planning loop”, where ideas are tested using the model population, plans are made from that learning, and the resulting outcomes are fed back into the model for the next stage of planning. This iterative approach is closely related to our scalability dimension.

This method uses census data, and longitudinal surveys to estimate individual activity patterns. When compared with even dynamic Markov models, it has the potential to provide more flexible and realistic responses to proposed intervention conditions. It does require longitudinal survey data, which is much harder to get, but it could be imagined that smartphone data, both with sensors and questionnaires, might improve the accuracy of this technique. And while it does not theorize about specific psychological features of human decision making, it might capture some of the behaviours. Very interestingly, this type of technique might be useful in validating a proposed behavioural model.

2.4.4 SIENA

Three papers use the SIENA model to study network effects of either obesity [105] or smoking [87, 56]. The primary purpose of SIENA is to study the separate effects of social selection and social influence. Social selection is the process by which individuals will select friends with certain observable features (i.e. choose your friends to match your behaviours). Social influence is the counter effect where people change their observable features based on their friends (e.g. choose your behaviours to match your friends). Since people do both, social network and behaviours are correlated. SIENA attempts to be able to pull apart these forces to see which is dominant.

To do so, the SIENA framework requires longitudinal data on social network and a behaviour measure (e.g. BMI). The 3 papers here all use Add Health data. While the SIENA model is a statistical approach, it uses an underlying, albeit extremely simple, ABM which contains many agents in a social network. At each time step, each agent examines the current network and either chooses to alter their behaviour value (e.g. BMI) by a small discrete step, up or down, or its social network by either breaking a current tie or forming a new one. How the agent updates their behaviour value is determined by a statistical function, characterized by the modeller, that can include influence and selection forces. A second function determines the predictors of possible network updates. These functions must include assumptions about how people update their state. For example, the average similarity assumption assumes that people change their behaviour value towards the average of their social network. This 2nd function also might include predictors for friendship such as whether befriending a fellow smoker is preferable. It can include any number of network features as well, such as out degree or transitivity. Rather than a mental theory, however, this is a statistical one that assumes that agents at some level make choices based on sometimes complex network characteristics.




Figure 2.10: Smoking states and transitions of the basic SimSmoke model [58]. Initiation and cessation rates are categorized by age, sex, and race/ethnicity. Relapse rates are categorized by age and years as ex-smoker. Years as ex-smoker is categorized in several year groupings: <1, 1-2, 3-5, 6-10, 11-15, and >15.

After the specification of the 2 update functions, SIENA attempts to predict network and behaviour states of each agent at later time steps, and matches it with the data from that time. By changing the update functions to assume more or less selection or influence, the output network from SIENA will have better or worse fits to the data, and this is used to infer the relative strengths of influence and selection in a given data set.

SIENA does qualify as a dynamic simulation model with a Markov chain inside each agent, albeit a simple one. It covers areas of public health including smoking. And it does have elements of human behaviour. As with the other data-driven methods above, it does not specify cognitive, or emotional elements explicitly. There is room for a social and a normative ability of SIENA agents, but it appears to be based less on psychological frameworks than statistical, network-based predictors of observed human behaviour.

2.5 Aggregate Models

At least seven models were found that treat populations in the aggregate, most of which used system dynamics, with the notable exception of SimSmoke [58] which appears to use difference equations. I outline some details of several, but owing to the similarity of some models along the dimensions covered in this report, we will neglect explicit discussion of several, including the Mendez-Warner SD model [72], and PRISM [45]. These 2 models also do not explicitly mention behavioural components.

2.5.1 SimSmoke

SimSmoke [58] is an aggregate compartmental model implemented in Microsoft Excel. Smoking states are characterized by the 3 aggregate categories of never smoker, current smoker, and former smoker, as seen in Figure 2.10. Each category stores counts of individuals, separated by age in years, sex, and race/ethnicity. Relapse rates are further characterized by years as ex-smoker. Birth and death are also described.

The transitions between smoking states are not informed by behavioural or psychological reasoning. They are population aggregate rate values, varying by age group, gender, and race, and are determined from relative population sizes in census data. For example, the number of current smokers in any age group below 24 is

found by comparing the count at the previous year, subtracting those that died this year, then multiplying it by the growth rate (i.e. initiation). (For ages above 24, it is assumed that there is no initiation, and below it is assumed there is no cessation.) Beyond census data, various health interview and tobacco habit surveys are used. Since SimSmoke is an aggregate model, summary statistics of the data sources are used to generate average transition rates.

SimSmoke uses modules to account for other factors. For example, [61] shows a module that looks at youth purchasing and vendor practices. They cite 4 general ways under-age youth obtain cigarettes, including by purchasing from vendors. To allow for possible interventions, they include merchants and their perception of risk, including both the penalty and probability of being caught and the social risk (e.g. loss of good will within the community). Law enforcement is included with an average checks per outlet rate variable.

Another interesting example is a module that allows SimSmoke to include smoke-free laws in public spaces, and account for an imperfect level of compliance [59]. Various cessation policies can be described as well, such as availability of pharmacotherapy, financial coverage for treatment, and quit line services. Another accounts for marketing bans on advertising and promotion of tobacco products [67].

It is recognized that cigarette taxes in industrial Western countries are well known, and there is an extensive literature around this area [60]. SimSmoke has the ability to be applied wherever relevant data can be sourced to study the effects of taxes on lower income countries. It has been applied in places like Taiwan [60], Russia [67], and Mexico [33]. Finally, it makes predictions of tobacco-related mortality.

Along our dimensions, SimSmoke does not explicitly address emotion or cognition. And sociality is captured essentially through the perfect-mixing assumptions of aggregate modelling. However, it does have the facility to describe normative forms, including both legal and social pressures. Also, calibration and validation is a regular feature of models (as examples, [23],[59]). Scalability is captured through the use of modules.

2.5.2 System Dynamics Models

The Tobacco Policy Model (TPM) [95] is a system dynamic model meant to study policy effects on health. The population is categorized into 6 age groups (10 year groups from 10 to 69 inclusive). The TPM compares the public health changes of 2 different kinds of behaviour change: a 10% change in initiation, cessation, or relapse for a given age category; and a one-time behaviour change in a fixed cohort size. Uniquely, this model provides a Quality-adjusted life years assessment of tobacco mortality/morbidity. It does not, however, incorporate assumptions about behavioural mechanism, instead, as with many population-level models, experiments with what population-level outcomes might be realized by proposed changes in behaviour effect, in this case assuming various situations of population initiation, cessation, or relapse.

The New Zealand Tobacco Policy Model [97] is an SD model meant to predict population cessation rates. What distinguishes it is the use of 2 feedback loops to incorporate social and normative pressures. Figure 2.11 shows the model in a stock-flow diagram. Loop R1 is the peer influence feedback, and loop R2 is the adult



Figure 2.11: This figure has been redacted for copyright reasons. It depicted the New Zealand Tobacco Policy Model with system dynamic elements, showing 5 smoking categories—never smokers, youth smokers, adult smokers, recent ex-smokers, and non-recent ex-smokers—and various interactions between them, such as how the adult smoker population influences the creation of youth smokers through role modelling. It can be found in [97].

role modelling. These role modelling forces are age-dependent in that they only affect adolescent smokers. This is another example of aggregate social effects. This model was calibrated against the New Zealand Tobacco Use Survey, the New Zealand Census-Mortality Study, and demographic data. Sensitivity tests were done by changing the base-case parameters by $\pm 10\%$ and observing output.

The ISIS model [14] is unique in that it is built more as an example of systems thinking than an applied model. As such, it does not place much emphasis on examining specific control strategies. What it does do is demonstrate how systems models, in this case SD models, can be successively built up as knowledge of the system is generated, relying on causal loop diagrams and workshop groups. The final model, shown in Figure 2.12 incorporates some behavioural mechanism. Smoking can become a social norm and increase the number of people starting. Public awareness of tobacco health risk can put pressure on tobacco companies to reduce marketing activities. There is also willingness of the government to legislate tobacco control. As a causal loop diagram, this description has not focused on the data required to parameterize or quantify this model, but it does display the ability of SD modelling to capture complex behavioural dynamics at a mechanistic, albeit aggregate, level.

Work done in [104] focuses on a model that incorporates many feedbacks (which we call the *Feedback-rich model*), and it describes the conditions during a proposed “societal lifecycle” of cigarette smoking, where each of the 5 loops, visible in Figure 2.13, has precedence. Loop 1 describes peer pressure, social norms, or role modelling. It takes precedence when smoking is low in the population. As more people start smoking, the population learns of negative health consequences (learning, and a competing social norm), which affect both cessation and initiation rates through loops 2 and 5. This serves to curb the unrestrained growth of the first stage. Loop 3 simply captures the population affect of a reduction of cessation numbers as prevalence decreases. This is seen when the decline in smoking after the health scare reduces in magnitude. Loop 4 describes how people forget (perhaps a cognitive element) the significance of adverse health effects, and so

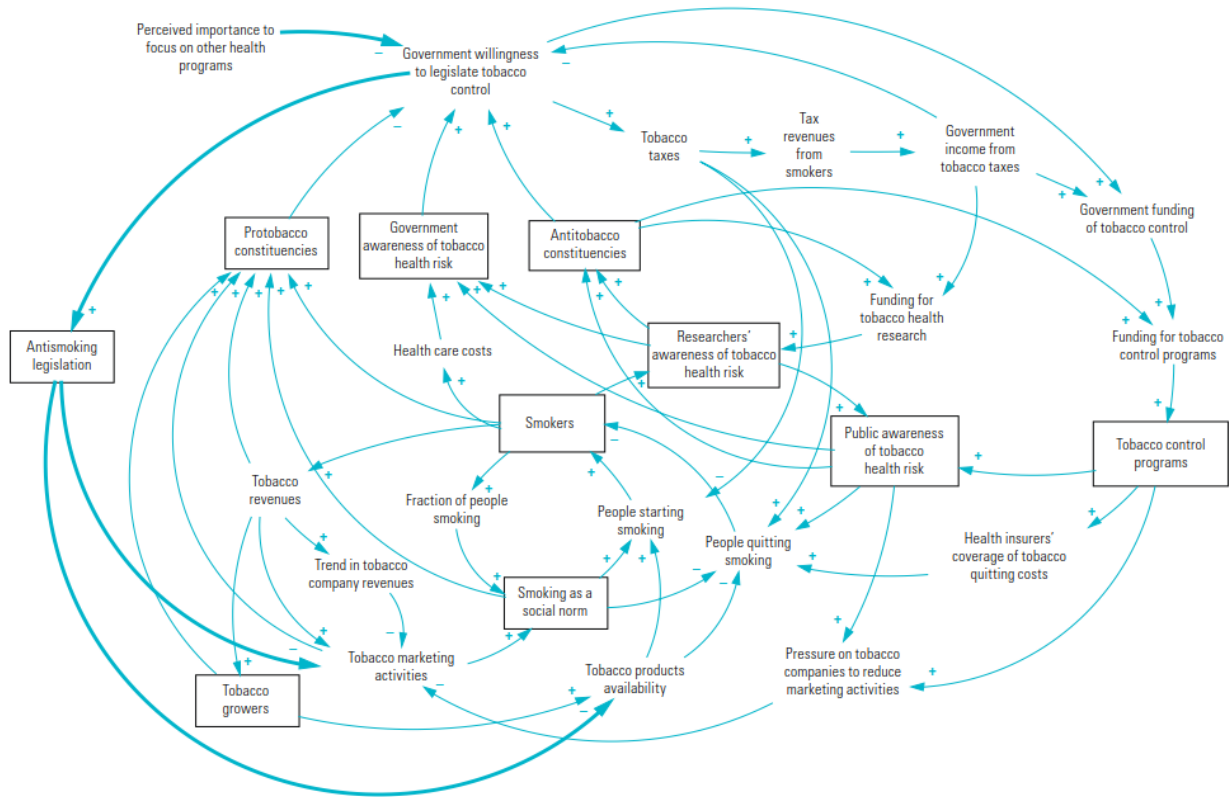


Figure 2.12: Final ISIS diagrammatic model. [14] Originally published as Figure 1 in A. Best et al. 2006. Systemic transformational change in tobacco control: An overview of the Initiative for the Study and Implementation of Systems (ISIS). In *Innovations in health care: A reality check*, ed. A. L. Casebeer, A. Harrison, and A. L. Mark, 189205. New York: Palgrave Macmillan. Reproduced with permission from Palgrave Macmillan.

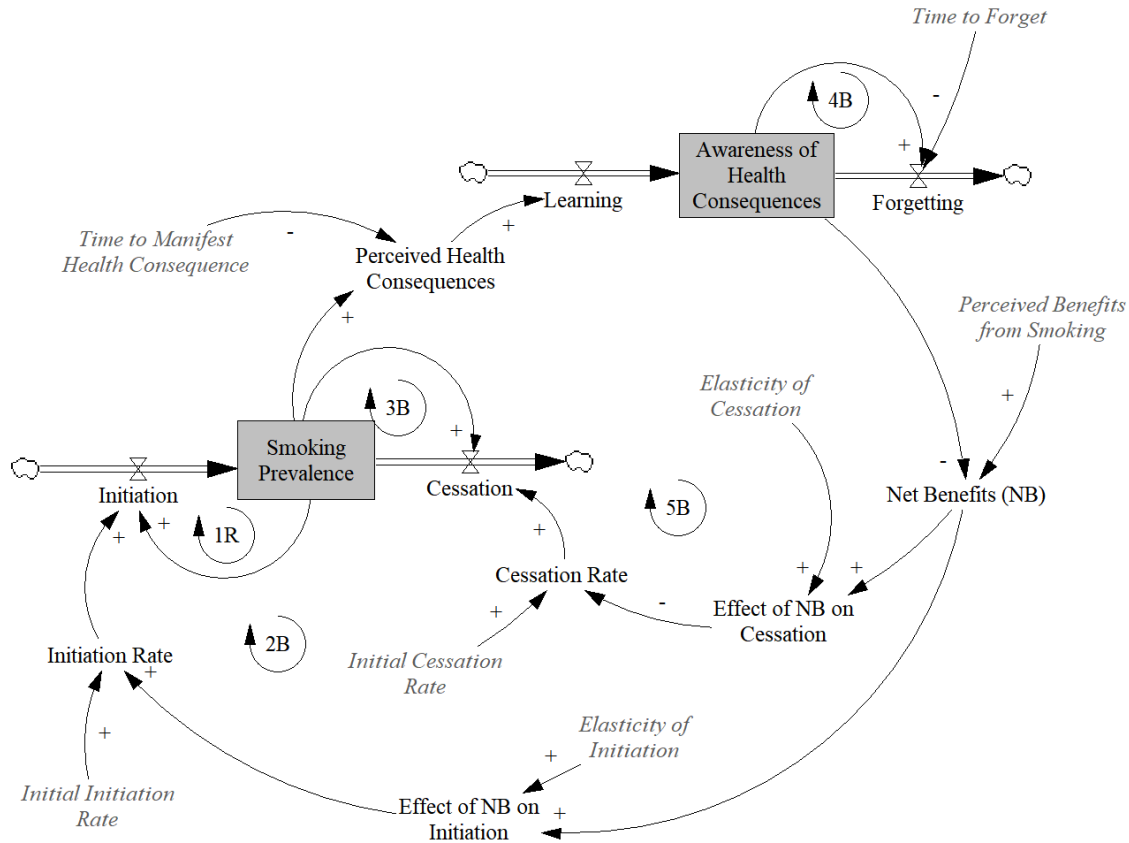


Figure 2.13: Stock-flow diagram from Feedback Rich Model [104].

the prevalence of smoking can start increasing again, starting the cycle anew. Elements of cognition, social norms, and learning are present in this model. As to applicability, this model is described as general enough to be applicable in multiple other problem areas. It appears calibrated to a much wider time horizon than other models, going back to the early 20th century, though the exact source of that data is not mentioned.

2.6 Discussion

One potentially unique difficulty with a literature survey in this field is its multi-disciplinarity. Relevant literature can be found in public health, artificial intelligence, sociology, psychology, and statistics, not to mention the areas of simulation and system science directly. As such, we do not expect that this survey serves as an exhaustive review of all relevant literature. Because of the frequency of cited literature already being in our list, however, we do expect that this survey can be considered a representative sample of the current state of the art.

What resulted were 4 very broad categories of modelling approaches: those from computational fields, such as artificial intelligence and game-theory; those from socio-cognitive and behaviourism fields; individual-based models that focused primarily on observable factors to human behaviour; and the aggregate approaches.

Some models have elements of several, such as the use of HBM and PECS together in [17], but we attempted to categorize them according to the primary purpose.

Additionally, we used a number of dimensions to facilitate comparing and contrasting the methods. These, outlined in 2.1.2, were behavioural (cognitive, affective, normative, learning, social, habitual), and design (implementability, scalability). The cognitive dimension is captured as both reactive and deliberative with BDI, and EMIL, and solely deliberative with CSL. LNA and CafeWilhelmina agents are reactive agents. Similarly, the HBM and TPB both only have the reactive mode of cognition, though this is no surprise since these latter two theories are explanatory rather than predictive and seek only precursors to behaviour. The individual-based observable models do not capture any explicit cognition. Only PECS and Consumat extend cognition beyond simply 1 of 2 cognitive processes. PECS provides a cognition component, envisioned with inner states and functions, though no implementations provide example mechanism. Consumat provides a very rich array of cognitive approaches, with varying cognitive loads and in various levels of uncertainty.

Very few models capture the affective dimension. The eBDI extension of BDI describes emotional dynamics, but no examples of an implementation were found. PECS contains the emotion component, so algorithms are capable of being adapted to that framework in principle. An example implementation of PECS [89] explores the within-agent dynamics between cognition and emotion that uses the concept of emotional intelligence and anger control. One could argue that the HBM could indirectly model emotions through the perception of illness.

The normative dimension is captured by many frameworks. While BDI does not directly, many of its extensions do to various degrees (e.g. BOID, BRIDGE). EMIL describes a norm internalization/learning process, but it does not distinguish between social and legal norms. CafeWilhelmina leaves out the learning process, but describes 3 types of norms (legal, social, and private). Both LNA and CSL describe the norm learning process, but only focus on social norms. PECS does not capture norms directly, though one could imagine using the internal states of cognition and social status to include normative dynamics. Consumat does not include norms directly, though mention was made of the designer of the approach discussing the use of legal norms. The TPB describes a subjective norm for each agent, but does not distinguish between norm types or account for the norm learning process.

From the sociality dimension, EMIL and CafeWilhelmina only use it for the norm learning process as social input. The only social interaction in LNA and CSL is done through game-theory interactions and physical observations. PECS contains a social status block, which allows, again in principle, social interaction. Consumat only mentions this dimension in relation to social comparison. The HBM and TPB do not provide social mechanism, though social input is likely necessary for the TPB subjective norm. Experimenters are left to develop their own approaches (such as opinion dynamics or follow-the-average). Of the observable models, the Beckman activity model [10] and the SIENA approach both rely on rich social network data to create their synthetic populations. The Beckman model also includes households and activity schedules for each agent.

Learning was captured by several approaches. EMIL uses dynamic memory while PECS maintains an internal state for each component, which are necessary precursors for learning. Consumat also provides agents with memory, and further specifies that that memory is updated with high-effort cognitive tasks. LNA and CSL offer the most descriptive method, both using a reinforcement learning technique to give their agents normative learning.

Habits and addiction were rarely mentioned explicitly. Most would use norms, memory, or as something that influences agent perception, but no mechanisms are described. Only the ad hoc below-the-skin methods explicitly capture addiction.

The ad hoc models are usefully considered in their entirety rather than on a per-model basis since they do not individually account for many of the behavioural dimensions, and then not very broadly. This is often due to the somewhat arbitrary algorithms used. For example, while the Reactance model relies on the psychological theory of reactance, the implementation is not argued to be generalizable, nor does it account for a variety of other cognitive elements (such as deliberation or reactive thinking). Taken as a whole, however, these approaches cover a wide range of human behaviour, capturing all of the behavioural dimensions, with the possible exception of the affective. Along the experimental dimensions, these models are noted for their scalability. The models are often very simply built, with little reliance on data, which means that the models can be applied in a very agile manner, and updated continuously. Indeed, one use of these models is to get initial insight into some of the key dynamics before seeking to build more complex, empirically grounded models. Reproduction of these models is usually quite straightforward, and without the need for much data, validation often need not be more than face validity. The utility of these models lies not in their numerical accuracy with real-world data, but in the insight they can provide. As part of a broader modelling research program, ad hoc models have much potential.

I also discuss aggregate models separately owing to their different approach to these questions, focusing on system-level accumulations and feedbacks. It is generally limited in the ability to describe individual human behaviour due to this aggregate focus, but it can elicit insight into what general forms of behaviour are most likely to give positive results. The New Zealand TPM includes the social effect of role-modelling; the Feedback Rich model describes learning elements of increase and decrease of health consequence awareness; and the ISIS model describes general motivational changes in many system players, such as tobacco growers, scientific researchers, and government legislators. Aggregate modelling promises to play a strong role, especially when considered within an iterative development environment. Further, the data used is generally easily accessible.

Both the ad hoc and system dynamics models excel in the two experimental dimensions. Both have low data requirements (little data required for ad hoc and more easily measurable population data for the aggregate models) and solid practices to study potentially key elements of human behaviour. Their uses in the iterative design pattern are complementary. Ad hoc models can be quickly built around a specific, though general, behavioural question (e.g. examining the impact of general types of social networks). SD models are able to describe the entirety of a system at varying levels of system complexity, as demonstrated by the

ISIS model. In the other frameworks, PECS, as a design pattern, is another excellent example of scalability.

Generalizing, we can see a dual-spectrum. At the one end, simulation-based frameworks tend to spend great effort on causal behavioural theories, excelling in the behavioural dimensions. However, they are rarely validated against public health data, so tangible predictions are not easily trusted. On the other, the observable models use many common data sources for public health and have rich experimentation with calibration, sensitivity, and validation procedures. These, however, make no direct causal assumptions of human behaviour, so are left without sure footing under intervention conditions not reflected in the data. The social-cognitive frameworks appear to sit in between these ends, with theories that are more descriptive at the agent level than the observable models, and more data focused than the simulation-based models. There is clearly a dichotomy between data-driven and theory-driven modelling, with the public health studies tending more towards the data-driven approach, leaving a clear opening for mature theory-driven success stories.

I believe this apparent dichotomy is resolved through the use of ubiquitous data, and with iterative and scalable methods of development. Given the recent and dramatic increases in sensor technology, human behaviour simulation will benefit from being agile to these new data sources. Many-dimensional individual-level micro-behaviour data is now possible with smartphone technology, which could provide extremely valuable means of grounding mechanistic, complex behavioural models in the real world, not to mention the increased benefits of using machine learning algorithms in behaviour models. Our own work has used the particle filter, a machine learning tool, to help in this grounding of a dynamic simulation with continuous real-world data, described in Chapter 3. Having much more grounded mechanistic models of human behaviour will not only help the simulation field, but provides another tool to develop psychologically and neurologically inspired frameworks as well.

Further, the benefits of scalability are hard to overstate. Iterating between model structure (and the necessitating theory) and the data available on the system can allow insight to be built successively which can lead to new experimental programs to gather specific data. There is no need to limit research to single techniques, relying instead on the aforementioned benefits of simple, ad hoc ABMs and system-wide SD models together. In Chapter 4, we describe our work that leveraged the benefits of so-called hybrid models, models which use both individual and aggregate approaches but yet find little expression in public health related modelling, in a framework that takes ABM and SD techniques together placing them in a standard relationship for easy scalability. Rather than being used for aggregate quantities, the SD tools are invoked for the dynamics of more abstract behavioural factors within the individual, such as “level of satisfaction”. I expect this to be a further area of intense future research.

Taken together, dynamic modelling provides many avenues of continued work. Research programs are currently invoking the system-mapping benefits of SD modelling; the insights and diversity of small, ad hoc models; the realism and flexibility of simulation-based models; the applicability of data-driven models; and the human-validated social-cognitive models. In the current burgeoning era of ubiquitous micro-data, we would do well to iterate over these techniques to avoid getting lost in the trees.

CHAPTER 3

PARTICLE FILTERING IN AGENT-BASED MODELS

This extends work that uses the particle filtering algorithm as a mechanism for re-grounding a dynamic simulation when repeated incoming measurements of the modelled ground-truth are available. This work is the first such application where the model associated with each particle is an ABM. This is partly due to the increased technical difficulty currently required to build an ABM particle filter using simulated ground-truth. The results were intriguing, reflecting the fact that the state-space of an ABM is in general much larger than a comparable SD due to the focus on agent heterogeneity, which poses a possible difficulty in an ABM particle filter. If so, the SD particle filter has advantages, not only in the relatively quick model development time, but in that the reduced computational complexity relative to an ABM gives SD a clear advantage when predicting short-term futures of the evolution of complex dynamics. However, the particle filter of the ABM still increased the predictive power of that ABM, and for some problems, an ABM might be more suited, for example if a strong focus on location or social networks is warranted.

3.1 Introduction

As described above, public health problems can reasonably be considered complex problems. The traditional use of time-series data and regression analysis techniques to infer possible causal relationships tends to scale poorly to complex systems. Dynamic modelling can be used as another tool to support learning about causal relationships. However, dynamic modelling has its own limitations. Models need to be sufficiently calibrated to real-world data to be able to make actionable predictions. While roughly calibrated models can often be employed to make short-term predictions, over time the discrepancy between model predictions and empirical observations will increase, requiring further calibration [79]. This growing discrepancy reflects inaccurate parameter values, omitted or oversimplified characterization of dynamic processes, and the vagaries of stochastic fluctuations in empirical processes. The resulting model can be unreliable and costly.

In the control of communicable illness outbreaks, there has traditionally not only been a heavy emphasis on use of continuous data from the field, but a strong recognition of the value of quickly built models that give actionable insight.

Predictor corrector models have been in common use for the better part of a century. One of the oldest and most widely applied is the Kalman filter (KF). In addition to extremely broad applications in real-time

control context, this approach has been applied successfully in the context of communicable diseases models [83]. However, Kalman filtering is associated with several notable shortcomings. Firstly, it relies on linear models, or linearized versions of non-linear models in the case of the Extended Kalman Filter [79], as well on maximum likelihood estimation, which provides only point estimates. In the context of the models typical in communicable disease epidemiology, both of these simplifications can pose tremendous challenges since these models are in general strongly non-linear. Secondly, KF presupposes Gaussian distributed errors for both measurement error and process noise - an assumption that can be highly problematic in the case of epidemiological count reporting processes, such as those with low numbers. Finally - and most importantly for this research - the KF linearization assumption cannot be readily applied to agent-based models, whose dynamics do not readily admit to representation using continuous-time differential equations.

Recent papers have leveraged particle filtering in the context of compartmental models [79, 78]. The idea is to use a compartmental model that incorporates broadly accurate dynamics for a given complex system. Then a particle filter algorithm is implemented over many iterations of the model. Using the weight-update and resample procedures of the particle filter, the collection of particles is hoped to approximate the posterior distribution over states and parameters over time. In practice, this can be used to enable the construction of quick and relatively simplified models which, when combined with real-world data, develop the ability to provide predictive value.

Indeed, in [78], the authors built a compartmental SIR model for a real-time H1N1 outbreak in Singapore, using data that came in from general practice and family doctor clinics throughout the country; they found that, in spite of some significant model over-predictions, and some significant model simplifications, the disease peak was well predicted. Also, the farther along the epidemic cycle, the more accurate the predictions of the particle filter became. The authors describe how particle filtering can be a powerful tool in building relatively simple models, which can be done in a more timely fashion, but still enabling those models to provide predictive ability for an ongoing and new epidemic outbreak.

In [79], the authors took a more theoretical approach, and built an agent-based model of a simple disease progression to serve as synthetic ground truth. A compartmental model was constructed using similar disease assumptions. The different natures of agent-based and compartmental models and systematic inaccuracies in a parameter value resulted in significantly different predictions of new disease infections over time. However, when using the particle-filtering technique, the compartmental model collection was able to strongly reduce the errors inherent in compartmental models.

In contrast to such previous contributions, we have not seen evidence before this work of using an agent-based particle model rather than the compartmental models. This is important for two reasons. Firstly, the particle filter - in contrast to the long-standing Kalman Filter [83] - does not assume a particular mathematical formulation of model dynamics. Secondly, there are areas where agent-based models have certain key advantages over compartmental models. Agent-based models are more able to represent the heterogeneous nature of agents - notably including heterogeneous histories - and their social networks [96]. *Our goal,*

therefore, is to implement and characterize the behaviour of an ABM particle filter.

3.2 Particle Filtering Algorithm

3.2.1 Intuition and Derivation

To understand how a particle filter is implemented here, it is helpful to understand the stochastic nature of the dynamic models being employed to understand an unknown ground-truth. In the communicable disease area, an agent-based model seeks to simulate processes such as contact, transmission over such a contact, incubation, recovery, and waning of immunity. Some sources of stochastics will be due to the need to capture temporal variability of a highly unknown parameter. How often or for how long does a person contact another person in their social network? That might best be modelled by drawing from a probability distribution.

A second type of stochastics comes from the need to generalize model predictions. If we could measure the social network of a population exactly, then we could make predictions about that particular population, but it would not be easily generalizable to others. So we would use a network type, which has basic structures in common but whose exact connections vary from simulation to simulation.

Thirdly, some characterizations of model predictions as stochastic reflect the fact that for certain processes, statistical regularities may be better understood than the mechanics of that process.

The implications of the above is that there are uncertainties about the model state from time-step to time-step. At the same time, there will be incoming empirical observations y_t (where, in general, y_t and other denoted quantities represent vectors). The goal, therefore, is to find a way of sampling from the distribution of model state x_t at a given time given all preceding observations $p(x_t|y_{1:t})$, taking into account both uncertainty in model predictions and that associated with observation error. By sampling over this distribution, we can - by extension - then sample over future trajectories of the model, sampling from the difference in gains between baseline and an intervention, or between two such interventions, and compute the probability of different events (e.g., the probability that a given intervention will save more than a certain count of lives, or will be cost saving). This challenge is particularly acute given that for non-linear models, we lack any closed-form characterization of the distribution.

The particle filter provides a means of accomplishing such sampling. Fully developed in [4], the approach is that of sampling from many weighted particles over model state, each associated with a specific state vector at a given time. More specifically, in accordance with the approach of importance sampling, we approximate sampling from such an unknown “target” distribution over model states by sampling from a readily available “proposal” distribution, but weighting the samples according to their relative probability (density) in the proposal compared to the target distribution. More formally, importance sampling states that sampling from a target density $p(x)$ can be approximated by sampling from another distribution, $q(x)$ - the proposal distribution, also termed the importance density - attaching a weight to that sample equal to $\frac{p(x)}{q(x)}$, and then sampling in turn from those samples, with a probability equal to their weight. In this case, each of those

samples from the distribution is termed a particle.

In the case of the particle filter, the sample weights are computed recursively, with the weights just prior to a given observation equal to the weight after the previous observation. Weights are only updated at the point of the observation. The weight update rule is given by

$$w_t^i \propto \frac{p(x_{0:t}^i | y_{1:t})}{q(x_{0:t}^i | y_{1:t})} \quad (3.1)$$

where the state of the model at time t for particle i is x_t^i , and w_t^i is normalized across all particles.

To obtain the full update rule and motivate the definition of $q(\cdot)$, we can decompose $p(x_{0:t}^i | y_{1:t})$ using Bayes rule into

$$p(x_{0:t}^i | y_{1:t}) \propto p(y_t | x_{0:t}^i, y_{1:t-1}) p(x_{0:t}^i | y_{1:t-1}) \quad (3.2)$$

Assuming that y_t only depends on x_t and x_t only depends on x_{t-1} , this can be simplified to

$$p(x_{0:t}^i | y_{1:t}) \propto p(y_t | x_t^i) p(x_t^i | x_{t-1}^i) p(x_{0:t-1}^i | y_{1:t-1}) \quad (3.3)$$

We will want to choose a form of $q(\cdot)$ that can be similarly recursively defined, such that

$$q(x_{0:t}^i | y_{1:t}) = q(x_t^i | x_{0:t-1}^i, y_{1:t}) q(x_{0:t-1}^i | y_{1:t-1}) \quad (3.4)$$

Combining these together and capturing the recursive character of the above through a recurrence relation gives us

$$w_t^i \propto \frac{p(y_t | x_t^i) p(x_t^i | x_{t-1}^i)}{q(x_t^i | x_{0:t-1}^i, y_{1:t})} w_{t-1}^i \quad (3.5)$$

The likelihood function $g(y_t | x_t^i) \equiv p(y_t | x_t^i)$ specifies the likelihood that we will observe the empirical data y_t at time t given the particle state x_t^i . This likelihood is specified explicitly by the modeler. In the context examined here, the probability density $p(x_t^i | x_{t-1}^i)$ is implicit in light of the emergent and stochastic behavior of the system under study.

Finally, we can greatly simplify this relationship with two choices. The first is to not consider the empirical data y_t in evolving the model between observations, only just after each such observation. The second is to require only a filtered estimate of $p(x_t | y_{1:t})$ instead of the full posterior. These allow us to choose $q(x_t^i | x_{0:t-1}^i, y_{1:t}) = p(x_t^i | x_{t-1}^i)$, resulting in the final weight update relationship

$$w_t^i \propto g(y_t | x_t^i) w_{t-1}^i \quad (3.6)$$

We only need to specify the likelihood $g(\cdot)$. We follow Osgood and Liu [79] and employ a negative binomial distribution.

$$g(y_t | x_t) = \binom{y_t + r - 1}{y_t} p^{y_t} (1 - p)^r \quad (3.7)$$

where $p = \frac{x_t}{x_t + r}$ and can be interpreted as the probability that a given reported case is a true case. r is the dispersion parameter. Because the distribution $g(y_t | x_t)$ will in general assign a non-zero likelihood to both the cases where the empirical data is greater than or equal to the actual model state and less than that model

state, this distributional assumption helps avoid a pathology in which all particles are assigned a likelihood of 0.

Every so often, the particle population will need to be resampled. This is due to the fact that, as the particle filter progresses, a few particles might have high weights while most others will have weights near 0. The advantage of simulating many particles will be diminished given that most predictive power will come from those few highly-weighted particles. In this case, therefore, a fresh batch of particles will be created by sampling from the particles according to their weights and creating copies. In a “survival of the fittest” phenomenon, most new particles will be copies of the fewer highly-weighted originals. Weights are then renormalized. While not explored here, sampling from the entire trajectory of states in the context of resampling requires tracing the lineage of a given particle from the first to the last time point, and then sampling from the particles at the final time according to their weight.

3.2.2 Final Algorithm

As can be seen from the section above, there are 2 basic steps: the first to set up the particle filter, and the next to iterate.

1. Initialization step: where time $t=1$, for $i=1..N$
 - (a) Sample x_1^i from $p(x_1^i|x_0^i)$
 - (b) Calculate an initial weight for each particle, assuming uniform weights $w_1^i = 1/N$.
2. Update step: where time $t \geq 2$
 - (a) Sample x_t^i based on $p(x_t^i|x_{t-1}^i)$ by running the model for each particle i from time $t - 1$ to t
 - (b) Set $x_{1:t}^i = (x_{1:t-1}^i, x_t^i)$
 - (c) Update weights using $w_t^i = g(y_t|x_t^i)w_{t-1}^i$
 - (d) Normalize weights to sum to 1
 - (e) Resample if necessary

3.3 Ground Truth Model

3.3.1 Introduction

We describe here the ground-truth model that generated the data for our particle filtering algorithm. The ground-truth model was constructed using the AnyLogic framework (based on Java), while, for several design reasons, the particle model was constructed in pure Java and implemented within the AnyLogic project. The entire modeling framework was thus a hybrid one from two different perspectives. Firstly, the framework

combined agent-based modeling methodology with a particle filtering methodology in a novel manner. Secondly, the model exhibited a hybrid handling of time, with the ground truth model running in continuous time using AnyLogics agent-based modeling mechanisms, and the particle filter versions of the models running in discrete time in custom Java classes.

3.3.2 Ground Truth Model

The structure for the ground truth model adheres to the extensively applied SEIR modeling paradigm [79]. Agents in themselves are categories as being in one of 4 states: (S)usceptible, (E)xposed, (I)nfective, (R)ecovered. [22] (This approximates many diseases of importance, including SARS, since infected individuals go through latent state before becoming infective themselves. See [62].) One clear benefit of this approximation is that it can be applied with agent-based models as undertaken here and in [79], or in a compartmental ground-truth model as employed in [78]. While not done here, this can facilitate the comparison of predictions between agent-based and compartmental models.

Within the ground-truth model, agents are placed inside of a social network, and agents have the ability to exchange messages only between other agents with whom they are connected. As discussed in the next section, the structure of this network varies from scenario to scenario.

Network Type

This work used four prominent network types: distance based, Poisson Random (“random”), Ring Lattice, and Watts-Strogatz small world (“small world”). The random network requires the specification of an average number of connections per agent; each pair of individuals in the population is connected with equal probability in light of the implied network density. In AnyLogic, the network choice was declaratively specified (and used AnyLogics built-in support for those explored network types), whereas it was constructed manually in the particle model.

Secondly, we employed a distance-based network. Here, we specified a connection distance threshold. A given pair of agents is connected if and only if they lie within a specified distance of each other. This network depends importantly on the spatial layout assigned to the agents. In both particle and ground-truth models, both the X and Y coordinates of agents were drawn independently from continuous uniform distributions between 0 and 500. A third network type used was a ring lattice network, in which all nodes are logically placed in a ring, connecting nodes with are within n nodes in either direction. Finally, the small world network represents a mixture between a ring lattice and a random network. The small world network can be generated by starting with a ring lattice network, and then a specified fraction of the connections are rewired to be connected with another node picked with uniform probability from across the population.

Several features distinguish the networks. While all serve as conduits for stochastic percolation processes for pathogen, the structure of a random network is less affected by these stochastics. In other words, many different random networks, with the same number of connections and the same number of nodes lead to

very similar connection patterns (because the variation of the number of connections for a given node is approximated by a Poisson distribution [30], and hence, similar behavior of disease spread.

On the other hand, distance based networks can be quite different between versions using the same construction parameters. If the connection distance is not so large as to make nearly everyone connected, distance networks can have disconnected subnetworks. The number and size of these disconnected sections can vary between parametrically similar networks. Even where the entire network consists of a single component, the highly overlapped nature of the paths from a given node (e.g., an infective) to other nodes means that the spread of contagion across such paths can be far more readily “blocked” than is the case in random networks.

By contrast, small world networks exhibit a mixture of the locality exhibited by ring lattice networks and the globally interconnected structure of random networks. As discussed below, this difference between networks under stochastics will affect how the particle models are able to adapt to different conditions.

Agent Behaviour

In the ground-truth model, agents decide their behavior according to the state-chart indicated in Figure 3.1. Here, agents all start off in the susceptible state. They can become exposed only when receiving an infecting message from an agent currently spreading the disease to whom they are connected (indicated by the icon over the transition). Once exposed, they will wait for a precisely defined incubation period, a parameter of the model, before becoming infective. For every time unit an agent spends in the infective state, it performs an action to expose a single neighbor to infection; if that neighbor is susceptible, the neighbor will enter the exposed state. Finally, after a precisely defined recovery period (another model parameter), that agent becomes recovered, which means that they do not infect others, and cannot become infected by others.

Actual and Reported Data

Following [79], within the ground truth model, we considered a difference between the actual daily incidence rate and the reported incidence rate. When someone gets sick, we assume a 50% chance of reporting that illness, with the reported count being drawn from a binomial distribution. The particle weights are updated according to the likelihood of observing that reported incidence (y_t).

Evaluation Metric

In order to assess the effectiveness of the particle filter in improving model accuracy, we compared the performance of the model absent particle filtering with that which included particle filtering. The comparisons were performed over 24 realizations and used a sum squared error metric to judge the discrepancy between the ground truth and estimates in sampled particles [79]. Specifically, equation 3.8 describes the discrepancy, Δ , where N is the count of particles, and n is the count of reported cases, subscripted for particle i , and the ground-truth g . The outer sum is over time.

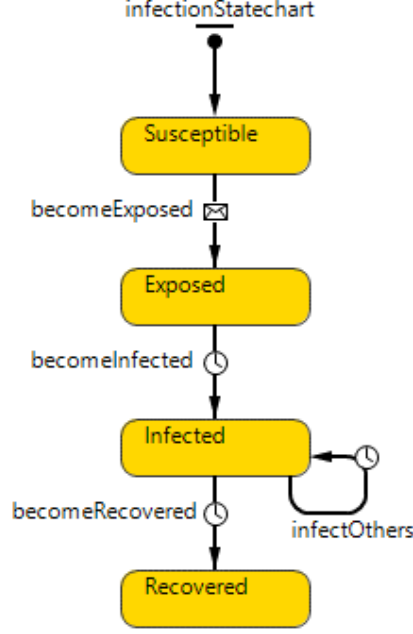


Figure 3.1: Agent state chart in the ground-truth model

$$\Delta = \sum_t \frac{\sum_i (n_i - n_g)^2}{N} \quad (3.8)$$

3.4 Methods and Results

Figure 3.2 displays the outputs from each of 9 Monte Carlo simulations; for each of the Random, Small World and Distance based network types, we ran a simulation with the particle filter off, on updated every time step, and on updated every 5 time steps. All of these networks were structured so as to have, on average, approximately 10 connections per agent.

What is shown in Figure 3.2 is that the particle filter actually results in poorer performance of the particles, most especially for the small world network. For this reason, we ran 6 more simulations, presented in Figure 3.3.

In Figure 3.3a we used a distance network, but characterized a situation in which the modeler was uncertain of the exact connection distance used in the ground truth model. As a result, while the ground truth model makes use of a specific connection threshold (41.5), the particles are created using a connection threshold drawn from a uniform distribution from 20 to 60 (runs are shown with “dist” appended to the name). As shown in Figure 3.3a, the particle filter without particle weight updating (D_PFOff_dist) returns an average error of 429 with standard deviation of 20.3. Updating each time step (D_I1_dist) reduces the mean to 344, but increases the deviation to 96.8. The Kolmogorov-Smirnov (KS) test between D_PFOff_dist and D_I1_dist is $r_{KS} < 0.0001$. Using an update interval of 5 (D_I5_dist) reduces the spread, but returns




Figure 3.2: Box plot of error metric over 9 simulations using 3 network types (random, small world, distance-based) with 3 particle filtering setups (no weight updating, updating every 1 time step, and updating every 5 time steps).

$r_{KS} = 0.14$ in relation to without weight update. Clearly, in the case of uncertain particle parameters, the particle filter results in a beneficial effect on the predictive ability of the particles.

In Figure 3.3b, we use a ring lattice network (implemented in the particles as a small world network with $\beta = 1.0$). Similar to the varied-parameter distance network, using an update period of 1 time step significantly drops the mean and increases the standard deviation. However, the update period of 5 time steps returns almost the same output as no update period at all.

Finally, Figure 3.4 displays the results from two simulations using a ring lattice network with 640 realizations (as opposed to the 24 realizations for all other results). With an $r_{KS} < 0.0001$, it is clear that the presence of the particle filter in Ring.II.640 provided a better fit to ground-truth output.

3.5 Discussion

The basic format of this experiment is mathematically similar to [79]. We initially expected to see the same type of particle filter output. That experiment produced Figure 3.5, which demonstrates two especially noteworthy features. The first is a wide range of particle trajectories often near the ground-truth data, expanding wider throughout the time between weight updates (indicated by the red and purple horizontal lines). A second is that the particle filter improves accuracy throughout the model. The range of the particle predictions (in blue) is very broad early in the model run. The discontinuities visible at $t = \{2, 3\}$ result from particles being reweighted and resampled. An example output from this particle filter is shown in Figure 3.6.

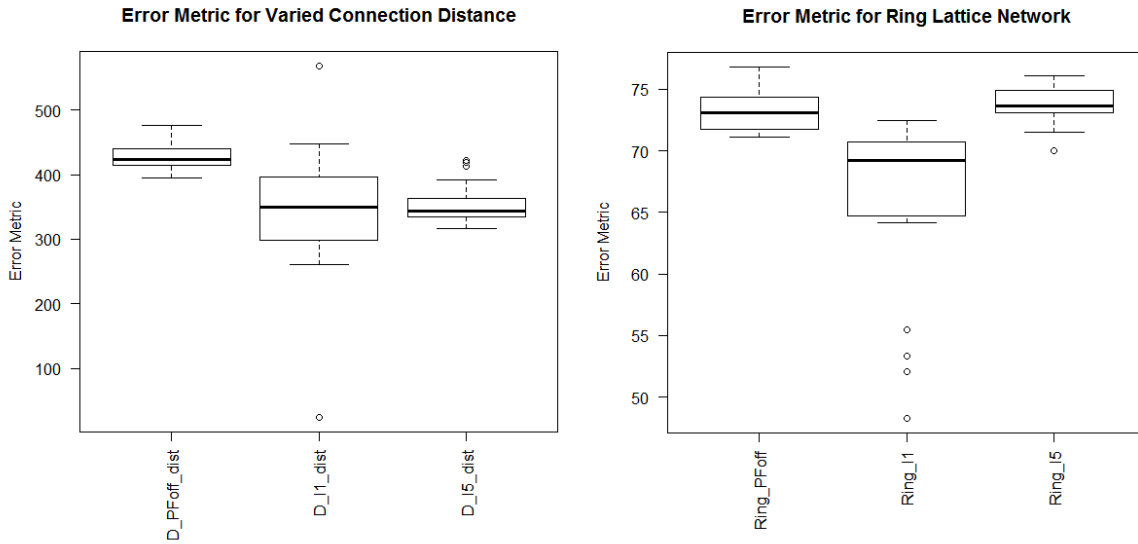


Figure 3.3: a) Box plot of error metric over 3 simulations using a distance-based network with ground truth connection distance of 41.5 and particle connection distance drawn from a uniform distribution from 20 to 60. Results shown for no weight update period, 1 time step updates, and 5 time step updates. b) Box plot of error metric over 3 simulations using a ring-lattice network with connection count of 2 per agent. Results shown for no weight update period, 1 time step updates, and 5 time step updates.




Figure 3.4: Box plot of error metric over 2 simulations using a ring-lattice network.

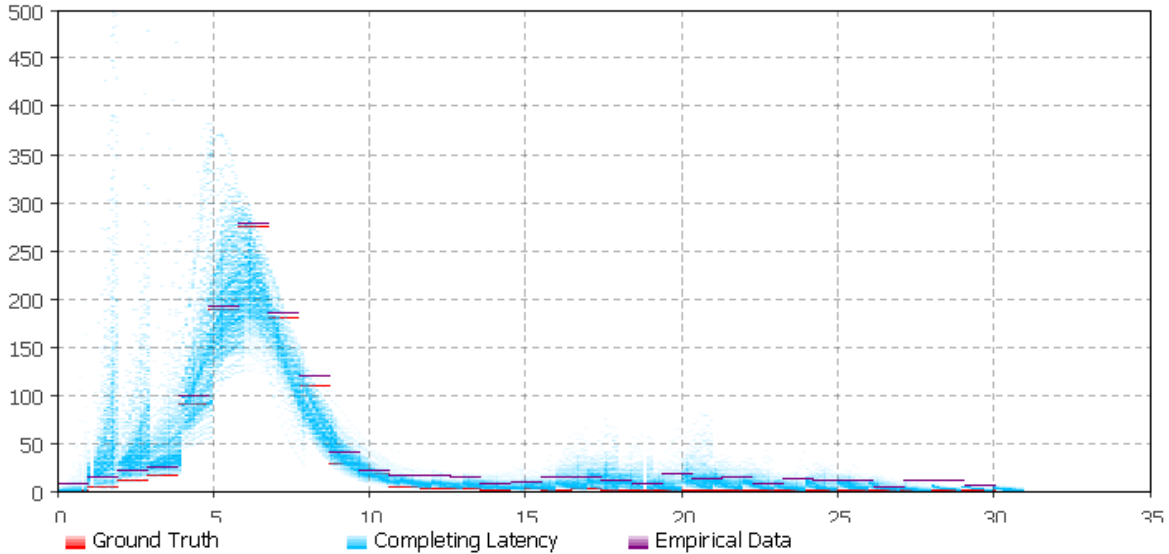


Figure 3.5: Similar particle filter using a compartmental particle model. This was obtained with permission from the original author’s source material.

These two features – very prominent in Figure 3.5 – stand in contrast to the results here.

Our initial results in bear this out. The agent-based particle filter either increases the mean error, increases the deviation, or both. Figures 3.3 and 3.4, however, provide strong evidence that the particle filter logic is properly formulated in the model, and that the observed behavior differs strikingly between different scenarios. This was further borne out by an extensive examination of model logic that drew on much previous familiarity with successfully operating particle filters.

We can think of two possible explanations for the poor performance of the particle filter in Figure 3.2 and its improved performance in Figures 3.3 and 3.4. The first is that the ground-truth model and the particle model are logically identical. By contrast, in [79], the ground-truth model is an ABM whereas the particles employ an aggregate model. They are necessarily different from each other, with the particle models inevitably diverging from the ground truth.

By contrast, in the current investigation, the simulations where the weight-update cycles are turned off constitute the best possible modeling situation - a situation where we are employing a model that precisely captures the dynamics of the underlying system being studied. The failure of the particle filter to improve upon the results of the “open loop” model with the particle filter off may be simply be reflection of the difficulty of improving upon a highly accurate characterization of the underlying system.

To probe this situation, we constructed Figure 3.3, which assumes a modeling context in which there is imperfect knowledge of the connection length in the distance-based network. Introducing this source of uncertainty and this mismatch between the dynamics of the ground truth model and that resulting from the particle models demonstrates strong benefits of particle filtering. As to why the update cycle of 5 time steps offers a smaller dispersion, it is currently assumed that this is because it would be filtering particles based




Figure 3.6: Example output for small world network ($n = 10$, $\beta = 0.9$). Weight update occurs every 10 time steps starting at $t = 1$.

on a 5-time-step sum of incidence cases. Small variations in the ground-truth output (whether due to actual incidence cases or the reporting of cases) are smoothed out, leading to fewer good candidate particles being filtered out because of momentary disagreement with ground-truth output.

The second possible contributor to the failure of the particle filter to secure gains in accuracy for the scenarios depicted in Figure 3.2 is that the agent-based particles have a much higher dimensionality than the aggregate particles from [79]. This leads to two complexities. Firstly, increasing the state-space dimensionality increases the required particle count to obtain an adequate sampling density from the state-space. Secondly, the particle fitness metric (incidence cases), which is used to compare particles, is a highly aggregate model output that confers comparatively little information by which to select one particle over another. In essence, there could be many local minima in the state-space where models of vastly different internal structure might nevertheless compare favorably using the simple fitness metric of incident cases.

These hypotheses are lent credibility by Figure 3.4. By virtue of its use of a deterministic network, these scenarios result in the ground-truth and particle models all having the same network connections. There will still be stochastics involved (e.g., regarding the timing of agent infection), but such a model exhibits drastically smaller state-space dimensionality than do those examined in the earlier figures. Given the lower dimensionality of the state space, the 1,000 particles are thereby posited to occupy a much higher state-space density, with a correspondingly improved set of possible particles for selection. Figure 3.3 demonstrates that running the update every time step noticeably improves particle filter performance over the no-update “open loop” version. As to why the 5-time-step update does not improve performance, it could be simply that the update frequency is too low. More simulations need to be run to examine this particular issue. Connected with this increase in state-space dimensionality is the increased computational complexity of the agent-based particle filter. While the high dimensionality of the agent-based model would make highly desirable very large ensembles (e.g., 1,000,000) of particles, the per-particle computational cost makes that infeasible. By

contrast, aggregate particle filters require less large ensembles (due to lower dimensionality), and can more readily support them computationally. An open question is what is the gain from using agent-based particle filters with fewer particles vs using aggregate particle filters with many more particles?

There are several other notable complexities about this particle filtering algorithm in general. The ground-truth model uses a binomial draw to determine how many actual cases are reported. This reported number is what is used in the likelihood function of the particle in eq. 3.7. The dispersion parameter essentially determines how broad or narrow the filter is. Larger r values make the negative binomial more peaked, and therefore admits a narrower band of particles “theories” as to the current state. Hence, this parameter has an effect on the whole particle filter behavior. Finding an appropriate value for r might improve performance.

There is a question as to the very idea of using a negative binomial in the likelihood function. Indeed, a binomial distribution makes more intuitive sense (as a chance of “success” or “failure” reporting or not where each opportunity is independent of all others). However, the binomial distribution has the unfortunate quality where particles positing a count of infectives less than the reported incidence are associated with a likelihood and thus a weight of 0. In situations where all particles posit more than the reported incidence, no particle weight can be calculated. This is the primary reason for selecting a negative binomial likelihood formulation.

3.6 Summary

The goal of a particle filter is to allow a dynamic model to be automatically readjusted in the presence of incoming data from the field. This chapter describes the first known published work using an ABM particle model, in which we successfully implement and characterize its behaviour. It is demonstrated using a model of infectious disease.

We test the performance of this approach using a simulated ground truth ABM, and a separate particle ABM. The simulated ground truth ABM generates data at regular intervals, which is used by the ensemble of particle ABMs to update the weights of each different particle. The more data comes in from the simulated ground truth ABM, the more the particle ensemble is able to replicate its dynamics.

We found that the particle ensemble was unable to increase the predictive capacity of the particle ABM under initial experimental conditions. Previous work using a SD particle model showed much more significant improvements due to the particle filter. We developed two hypotheses for the poorer performance of the ABM particle filter. Initially the two ABMs were logically similar with the same parameter values, implying that the particle ensemble was already near optimal. When we estimated a particle parameter using a distribution, we observed the particle filter improve prediction of the ensemble. Secondly, the ABM is in a much larger state-space than the SD particle filters of other researchers, leading to the need of higher numbers of particles to adequately sample the state-space. To test this, we selected a social network with fewer degrees of freedom, and again, we observed an improvement in the particle filter. Our hypotheses were validated through these

experiments, but more work needs to be done to more fully understand what situations ABM particle filters might be better suited.

CHAPTER 4

MODULAR DESIGN PATTERN

In Chapter 2, I noted only one other individual-based behavioural framework, PECS [89], that placed emphasis on design scalability. However, PECS does not provide tools for actual implementations. The ability to build ABMs in a scalable fashion is currently tightly constrained by the development tools available, and no work in my literature review has advanced examples of actual construction patterns to achieve a scalable architecture. This contribution is the first such example for models of human behaviour.

4.1 Introduction

As described above, aggregate and individual-based dynamic simulations approach complex systems modelling in complementary ways. Depending on a variety of factors (e.g. available data, the scope of the question, the importance of population heterogeneity or location, the knowledge of system feedbacks, the understanding of stakeholders to an certain approach), a given project might benefit more from using one or another of these approaches. Another option is to build a model with elements of both. Much work has been done attempting to standardize these hybrid techniques [92]. And as with many projects studying complex systems, development paths may change as progress is made. *Therefore, we demonstrate a pattern for hybrid model construction that allows for easy connection between components of a model built with differing methodologies.* This pattern, demonstrated with an example hybrid model, focuses on model modularity to support incremental development, and continual and efficient testing.

The benefits this technique offers are threefold. Firstly, proper modularization of a model supports incremental development, a key feature of the agile programming paradigm [34]. This allows development to adapt more easily to changes in design requirements by enabling modules to be removed or added without requiring substantial changes to the model core. Given the role of models in driving insight that did not exist at the beginning of a modelling project, it can be seen how this might be a valuable feature of model design as well. Secondly, a modular design allows for improved testing and learning. Since modules are able to be simulated independently of each other, a given module's dynamics and response can be uniquely characterized. This improves investigation to see if unpredicted model behavior is due to fresh insight or code errors. Finally, for those software packages that allow multiple methods in the same single architecture (AnyLogic, being a prime example), modules themselves can be single-method even though the model as a

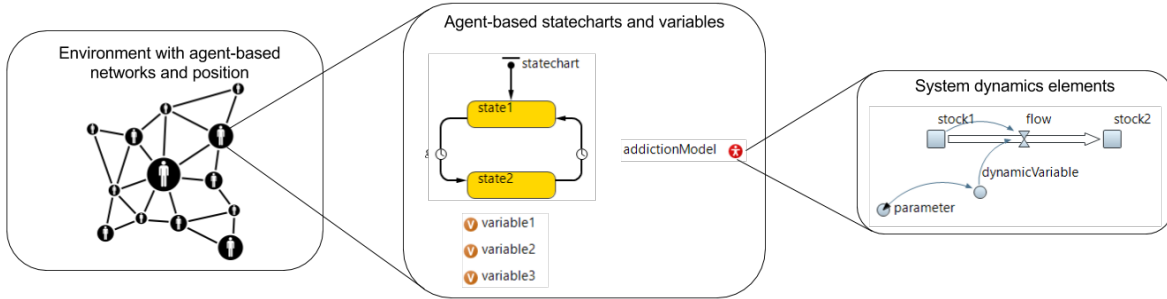


Figure 4.1: Our hybrid technique.

whole draws on several, enabling clearer thinking about model structure. This could allow a unified method of model characterization building on current examples, such as the ODD method of ABMs [35].

4.2 Technique Description

The example model, built in AnyLogic (Version 7.1.2), describes the evolution of addiction to tobacco cigarettes in a human agent population. It uses components of ABM and SD to leverage their unique advantages. As shown diagrammatically in Figure 4.1, an ABM is used to describe the social network, and some agent characteristics (statecharts and variables). It further captures the physical environment (allowing a proximal network), and agent location (work or home). SD is encapsulated within the agent (in the addictionModel object), serving as the formalism of the addiction-related mental processes. It deals with much less quantifiable measures (e.g. level of addiction). This SD module, along with a Markovian version used to contrast behavior, will be described in the following sections.

The model consists of a number, n , of agents. All agents move continuously between a home (h in number) and a school (s in number), spending 16 hours at home and 8 hours at school (determined by simple timeout transitions). All agents are placed in a scale-free social network (described in the AnyLogic documentation as being generated according to the Barabasi-Albert algorithm [8]). In this initial model, agents are not currently distinguished by demographic features, such as age or gender.

There can be multiple agents in each home and in each school. At model startup, agents are randomly assigned to a home. Schools and homes are distributed uniformly throughout a bounded 2D space. An agent is assigned a school based on the closest option from their home. Agents are aware of their social network (agents connected in the above-mentioned scale-free network), and their proximal network (defined as agents occupying the same space, either home or school, at a given point in time).

Agents also can occupy one of 3 different states in a smoking statechart, depicted in Figure 4.2, NeverSmoker, CurrentSmoker, and FormerSmoker. CurrentSmoker agents will either be in the NotSmoking or Smoking state. Four of the five state transitions are rate transitions (i.e. initiation, quit, relapse, and haveACigarette), whose transition times are sampled from an exponential distribution with the average transition time being specified by the respective hazard value, h_i , h_q , h_r , and h_c . The 5th (i.e. finishedCigarette) is

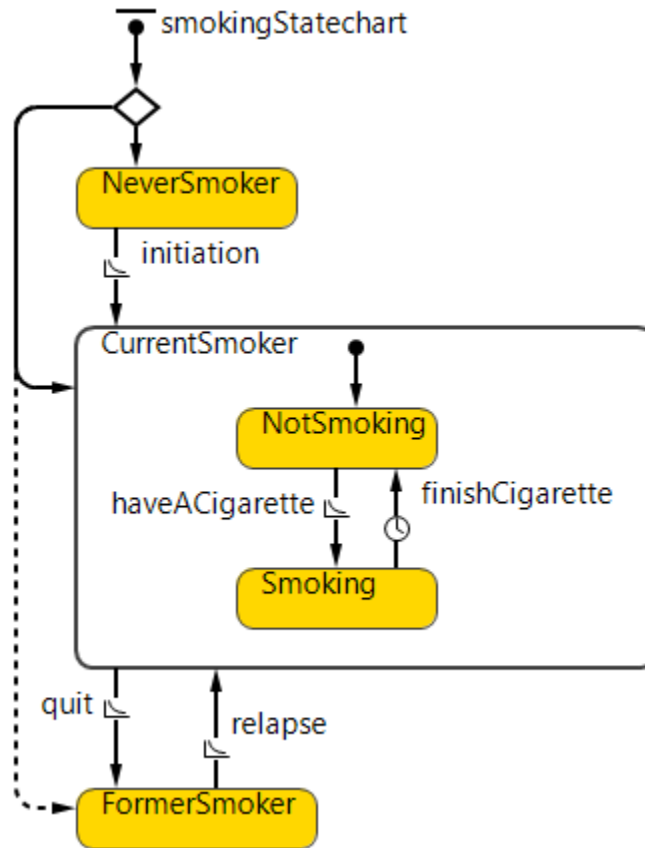


Figure 4.2: Agent smoking statechart.

simply a timeout transition that corresponds to the length of time spent in the Smoking state, set in this model to be 10 model-minutes.

To define the time evolution of these 4 rate transitions, each agent is given an AddictionModel module, which encapsulates the necessary logic. A Java Interface is used to define a contract that all individual implementations of the AddictionModel object must follow, enabling easy adaptability. The goal is to abstract over implementation details of a given addiction module. This serves two advantages. First, it reduces coupling between modules since modules do not need to know the details of implementation of any other module with which they interact, allowing those details to be changed at any time in the model development process. This will be used below to easily switch between different implementations of addiction without changing the rest of the model. Second, this enhances transparency by focusing attention on the contract-based interaction between modules, rather than on implementation details of any one. In other words, it helps to see the forest for the trees.

4.2.1 Modularity Through The Java Interface

As just discussed, the goal of this design pattern is to abstract over implementation details. This is accomplished here using the Java Interface, which specifies a contract defining the main methods, including inputs

```

public interface IAddictionModel
{
    public double getInitiationHazard();
    public double getSmokingHazard();
    public double getRelapseHazard();
    public double getQuittingHazard();
    public void updateModelValues(Person p);
}

```

Figure 4.3: IAddictionModel code.

and return values, of any Java class that implements the Interface, without specifying any implementation details. The Interface *IAddictionModel* defines the contract for this example: every *AddictionModel* module that implements the Interface will have at least 5 basic functions; 1 corresponding to each of the 4 rate transitions (taking no input and returning a double, the respective rate for that function), and 1 which takes as input the *Person* object, and outputs void (Java for “no output”). This input function provides the addiction module with access to all of the agent’s parameters enabling any addiction module to have access to any person parameters necessary for its normal functioning. The code is presented below in Figure 4.3. The next sections give examples.

Each agent has a setup parameter called *myAddictionModel*, which is of type *IAddictionModel*. Due to Java’s polymorphic structure, this allows *myAddictionModel* to contain any addiction module that implements that Interface regardless of its internal structure. A setup function is defined which populates *myAddictionModel* with one of the already-defined addiction modules.

This setup function takes input from a global parameter, specifiable from an *AnyLogic* experiment. An experiment in *AnyLogic* is a particular encapsulation of a given simulation. Among other things, it defines the global parameters. Multiple experiments can be defined that specify different initial parameter setups. Therefore comparing multiple different addiction modules is as simple as changing an input parameter.

In order to replace a given addiction module with another that has been independently constructed (e.g. with stocks and flows, or ABM statecharts) and initialized (i.e. with defined initial parameter values), only 3 steps are needed to plug it in to the greater model. Firstly, any inputs this new module requires to function must be defined and calculated by the *Person* in which it is embedded. Secondly, the setup function needs to be expanded to include the new module as an option (which is very simple given the setup function is basically a switch statement). Finally, available modules are named in a Java Enum class, which is simply a list of options. A new element in this list must be specified with the name of the new module. Then, whenever an experiment is defined that uses this new module, it will run with it. This presents a very modest amount of extra work.

We now give 2 examples of different addiction modules that can be easily implemented and swapped.

Markov-inspired Addiction Module

To make predictions of tobacco prevalence, Killeen [50] uses a 3 state Markov model (active smoking, recent cessation experiencing withdrawal, extended cessation), calibrating model parameters to various published research. This serves as an inspiration for our first addictionModel example. Our smoking statechart (Figure 4.2) is already essentially a Markov model, with slightly different states. As a result, each of the 4 output functions simply returns a daily hazard for each transition required. This addiction module uses a daily hazard for quitting and relapse that approximates some of the data reported by Killeen. Triangular distributions ($T\{\min, \max, \text{mode}\}$) are used. The quitting hazard, h_q , is a draw from $T\{0, 0.19, 0.095\}$, and relapse, h_r , a draw from $T\{0, 0.162, 0.081\}$. Initiation, not accounted for in the original research, is assumed to be half as likely as relapse, so h_i is drawn from $T\{0, 0.081, 0.0405\}$. Finally, for current smokers, the smoking or consumption hazard, h_c , is taken from $T\{0, 30, 10\}$, signifying that smokers will have up to 30 cigarettes per day, but averaging at around 10, or half a pack.

It should be mentioned that there are 2 main differences between this model and the usage by Killeen. Since Killeen is modelling the behavior of smokers exclusively, that project omits the NeverSmoker state, which we include. And while Killeen distinguishes between recent and long-term cessation, we do not.

This first addiction module does not contain the necessary logic to enable agents to change their addiction level over time based on external conditions. It is used as an example to demonstrate the flexibility of the modelling pattern employed here. However, it can be readily employed to validate other components of the model. By developing a simple standard, changes to network topology, for example, can be observed separately from the variations in agent decision making resulting from a more articulate addiction module.

Perceptual Control Theory Module

This module was designed following Powers' definition [81] of Perceptual Control Theory (PCT). It uses a layered set of two feedback control loops shown in Figure 4.4 displays the logic using stocks and flows, which clearly shows the 2 feedback loops; the *Action Loop* and the *Intent Loop*.

As a reminder, a stock is simply a variable that changes due to the flows - it is an accumulation, and therefore changes gradually over time. Increasing a flow does not instantly change the stock: it requires time for the flow to have its effect. The stocks in this figure represent the current state of either the *addictionAction* or *addictionIntent*. These determine the various hazards, using the equations below, where $h_{c,r,i,q}$ indicates the hazard for consumption, relapse, initiation, and quitting, respectively.

$$\begin{aligned}h_c &= \max(0, \text{addictionAction}) \\h_r &= \max(0, \text{addictionAction}/2) \\h_i &= \max(0, \text{addictionAction}/5) \\h_q &= \min(0, -\text{addictionIntent} * 100)\end{aligned}\tag{4.1}$$

Smoking or consumption, h_c , is defined in terms of the action stock. Relapse, h_r , is as well, but dividing

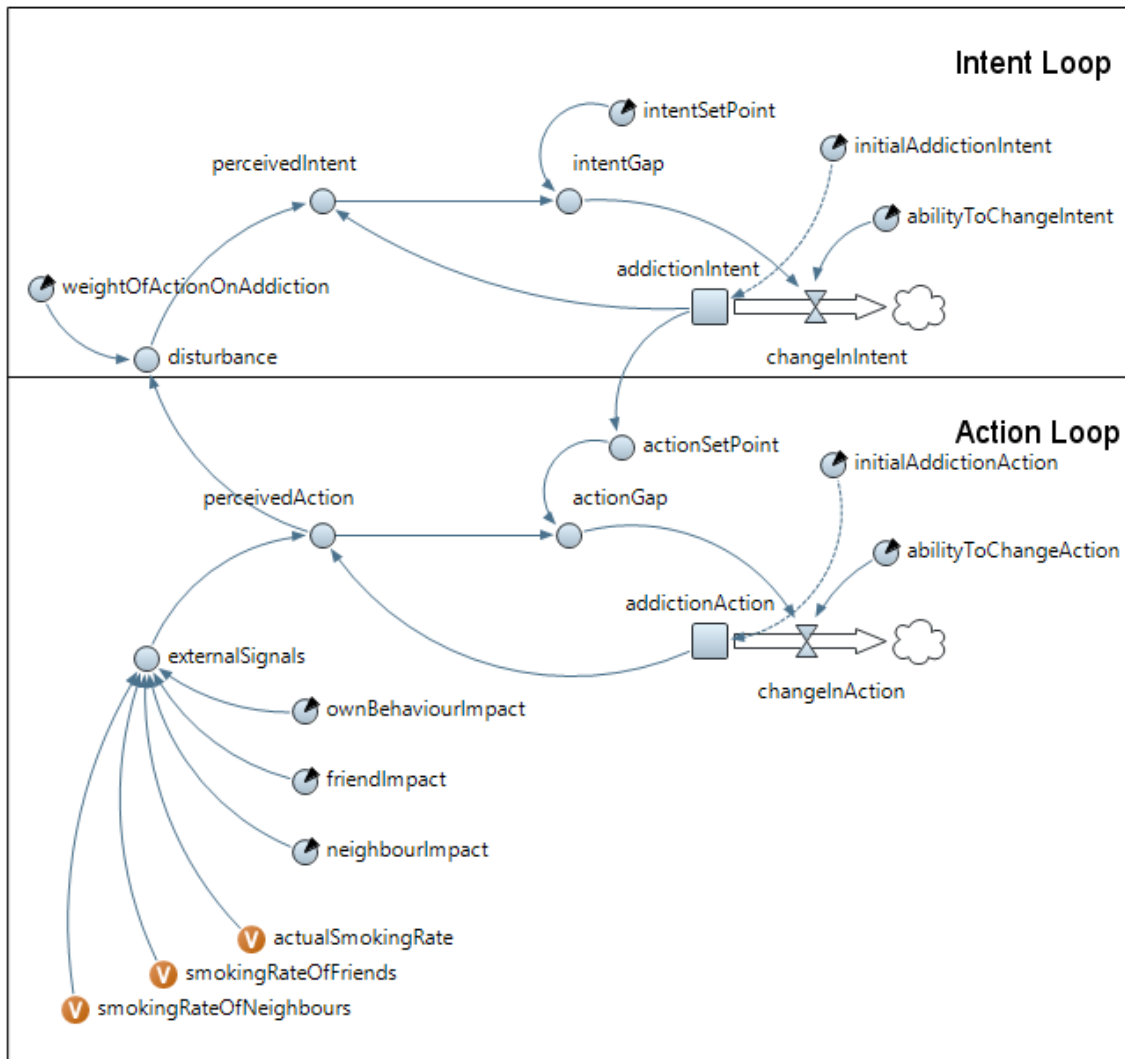


Figure 4.4: System dynamics diagram of our implemented PCT module.

by 2 reduces the strength, capturing the stylized fact that it is easier to continue smoking than to relapse. Initiation, h_i , is also defined using the action stock, and defined with the smallest hazard rate. Quitting, h_q , however, is defined using the intent stock. This is inspired by the idea that quitting involves a measure of conscious choice; quitting happens intentionally.

Each loop follows the same structure. It describes a target-following process. The task is to keep the value of the stock as close as possible to the set point - for addiction action, this is the *action set point*; for addiction intent, it is the *intent set point*. Therefore, the task is identical to minimizing the difference: *action gap* or *intent gap*, respectively. This is done by changing the flow: when the gap is positive, it means that the target is greater than the value, resulting in a flow that increases the value; and the reverse when the gap is negative.

What makes this a control of perception is the fact that the variable doing the target seeking is not the stock itself - the action or the intent - but rather the perception of each. The person is only able to influence their action, and only able to perceive their perception. Looking at the action loop of Figure 4.4, the *perceived action* is determined from both the “true” action and the *external signals*. I have defined the external signals to be the amount of smoking that the agent observes among their friends, other agents in their location, or even of their own smoking. The motivation is that the agent’s perception is skewed by the behaviour of other agents in their social network or physical space. The 3 impact parameters (*own behaviour impact*, *friend impact*, and *neighbour impact*) determine the weight given to its respective external signal.

The agent also has a *perceived intent*, but this is not influenced directly by the external world. The analogue of the external signals for action is the *disturbance* for intent. It is the perceived action weighted by the *weight of addiction on action*, and sums with the addiction intent to form the perceived intent. This communicates the hierarchical nature of the PCT; the intent loop does not interact with the external world, neither through controlling the agent’s smoking action, nor by observing. It interacts through changing the action loop’s set point and observing the action loop’s perception.

4.3 Results

For each of the simulation runs conducted, the starting population was the same: 100 agents, 5 schools, 50 homes. Each person had a 20% chance of starting as a CurrentSmoker, and an 80% chance as a NeverSmoker. All agents were connected in a scale-free network ($M = 10$). For each run, 3 outputs are reported: the final proportions of the 3 different smoking statuses, the daily smoking rates of all current smokers at model end, and population-wide smoking-rate averages over the time of the model. These were chosen to show the different model outputs, both at model end and over time, of smoking status and use. All simulations were run for 1,000 model days.

Given that the PCT model allows incorporation of network effects, we also ran this model with no network effects (i.e. by setting `friendImpact` and `neighbourImpact` to 0), so as to enable a better comparison with




Figure 4.5: Histogram of final smoking status.

the no-network Markov version. Both network and non-network version of the PCT model were initialized with an addictionLevel=uniform(0,10), and addictionAction=0. These initialization parameters can be, in general, quite important for the time-evolution of model parameters. We have not attempted to characterize this sensitivity, nor did we allow the stabilization of addiction stocks (through a warm-up period), in part because the dynamics themselves are part of the difference between the 2 addiction module approaches. Future research in using the PCT model will better characterize this behavior.

A word should be mentioned on the stochastic nature of these models. A proper analysis of output for such models would include a measure of the variability of the output given the same starting conditions. However, the focus of this research is on the nature of this particular hybrid design approach. The reported outputs are representative of a typical simulation run, and focus is on comparing the patterns that appear rather than the specific values.

4.3.1 Markov addiction model

Figure 4.5a shows the final smoking status distribution. With the Markov chain addiction module, the count of never smokers would stay relatively high, around 70 to 80. And current smokers usually outnumber former smokers. This was a generally consistent pattern over multiple simulation runs. For smokers, the smoking frequency usually had a mean between 10 and 20 cigarettes per day, and was a two-tailed distribution, as seen in Figure 4.6a. Finally, from Figure 4.6b, the average smoking rate of the whole population (including non and former smokers) did not display much variation, and stayed relatively consistent between 2 and 4 cigarettes per day.

4.3.2 PCT addiction model

We can see from Figures Figure 4.5b and Figure 4.5c that with the SD addiction module, nearly the whole population became and remained active smokers, whether networks were used or not. This, obviously, is not

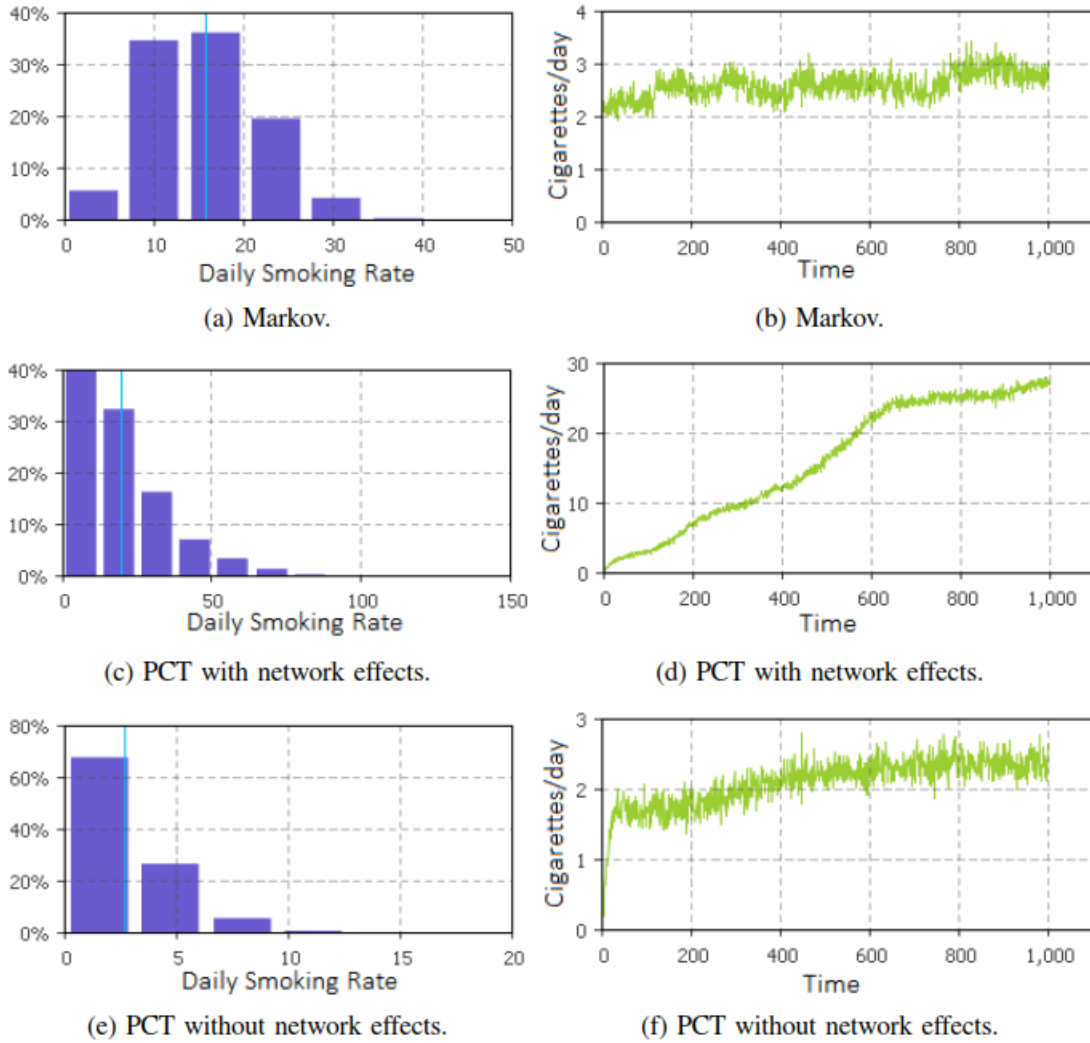


Figure 4.6: Histogram of daily smoking rates for smokers and population-wide average of daily smoking rates over time.

realistic for most populations, but displays a dramatic departure from the Markov addiction module.

Figure 4.6c shows notable differences in the smoking rates for active smokers. Rather than a two-tailed distribution, most agents here smoke relatively few cigarettes, with a single tail reaching much farther up to more than 3 packs per day. Without a social network, Figure 4.6e, the single tail structure remains with much smaller smoking rates among smokers (most under 3 cigarettes per day). There is a large apparent effect of the social network on smoking intensity, but less so on the shape of the distribution.

Finally, Figure 4.6d shows a gradual increase in population smoking rate, likely due mostly to the increase in smoker initiation. Reinforcing this point, removing network effects Figure 4.6f shows a largely flat population smoking rate curve. While initiation is not caused by the social network, it is clearly accelerated by it.

4.4 Discussion

Clearly the two addiction modules make vastly different assumptions. One uses unchanging, individual parameters and neglects network or internal feedback effects. The other relies on networks and feedbacks, and incorporates a framework for individual addiction dynamics. The Markov model outputs consumption plots with 2 tails and a clear central maximum. The PCT method predicts a single tail decline. Further, since the PCT method takes into account social network, the outputs are drastically different.

While neither of the demonstrated addiction modules are validated to real-world data, they serve as a very useful test case to display the benefits of this modelling pattern. As mentioned in the introduction, one benefit is the ability to adapt to changing model requirements through an incremental approach, an important part of agile software development. If initial requirements assume that agent heterogeneity and location are key interests to examine in the model, the Markov module might be a good choice. If throughout the modelling process it becomes clear that agent behavior needs to be more complex, without resorting to building a new model or substantial recoding of core model behavior, a more complex PCT module can be implemented. The model design is relatively ambivalent to which modelling approach is used, allowing many possible hybrid combinations.

There is a need for comparative testing of behavior theories. This is currently done with human subject research, which is costly and time-consuming, e.g. [29, 91]. Using this basic architecture, it is quite easy to imagine the development of any number of differing addiction modules based on, for example, the Theory of Planned Behavior or the Health Belief Model, and using any combination of aggregate or individual modelling techniques. The ABM framework is indifferent to the technique, so long as the contract is matched. The many benefits of dynamic simulation can be extended to this area of comparative research.

Increased learning and testability was also mentioned. Since the module has been reified as a model parameter itself, it can be treated as such in sensitivity analyses. Here, it is clear that there are substantial differences in output between modules, indicating that it is a highly sensitive parameter. This could suggest more work go into the internal mental model rather than other features (e.g. GIS). And, module behavior can be tested with very simple agent populations, allowing selective and precise testing plans.

Finally, this pattern has specific benefits to hybrid models for those packages that allow hybrid techniques (including purely code-based approaches). Formalizing a set of design patterns has helped standardize software engineering practices [42]. This pattern, then, can serve as another step in the same process for hybrid dynamics models.

Compartmentalizing each technique into a distinct module allows for extensions unique to that approach. For an SD model, due to an often aggregated approach to system parameters, adding a new causal pathway (for example, adding the influence of law makers on tobacco prices) might be a relatively straightforward addition, whereas for an ABM this might add substantial extra mechanism (for example, requiring the introduction of law maker behavior under a variety of model conditions). ABMs, on the other hand, can easily

add new agent-level characteristics (for example, including a new demographic), which would require complex subscribing in an SD model. Separating the details of module implementation through an abstraction of the Java interface allows each module to be extended relatively freely from each other.

4.5 Summary

We defined and demonstrated a novel application of a design pattern in a dynamic model, using an ABM of tobacco addiction with modular addiction logic. Three modules were built as examples: one inspired by a static Markov chain and the other two by the Perceptual Control Theory.

Within the agent is a statechart that determines the smoking state over time. There are four state transitions that need to be determined to completely define the smoking state: initiation, quitting, relapse, and consumption. Using a contract specified by a Java Interface, we require any module to provide values for each of these four transitions. This allows us to maintain flexibility while reducing unforeseen behaviour.

The results demonstrate significant differences across addiction modules in both the smoking rate distribution and smoking prevalence. As well, only with the PCT module that used the social network was a change over time in smoking rate seen: the population was becoming more addicted.

Using the modular pattern, the choice of addiction module can be made independently from other model constraints. It also allows for determination at runtime which module to implement. It even allows multiple modules to be enabled simultaneously for different subpopulations, though this was not done here.

This flexibility in implementation and design provides three key advantages. Firstly, testing of the model as a whole is more straightforward. Since the behaviour of modules is defined by a contract, they can be tested independently. Relatedly, the source of any discrepancy between expectations and model behaviour can more quickly be discerned, leading to faster insight building. And finally, it provides a very principled method for building hybrid models. In this case, the PCT module was built with SD elements.

CHAPTER 5

TOBACCO PRICE-MINIMIZING STRATEGY MODEL

This chapter describes an ABM that is built to study tobacco-related purchasing behaviour in the light of excise tax. Using the modular design pattern discussed in Chapter 4, four separate modules that describe the dynamics of agent addiction are created. The ABM contains homes, schools, and tobacco retailers, as well as people who consume and purchase tobacco, and are distinguished not only by their location in the city, but by their education level. Parameter estimates and spatial distribution is roughly matched to Minneapolis/St. Paul and surrounding area. We performed calibration for each of the different addiction modules, comparing and contrasting their outputs. Results suggest that different sub-populations are driven by different dominant behavioural mechanism, that externally validated quantitative models can improve accuracy, and that an iterative development process can help to improve model design.

5.1 Introduction and Motivation

Tobacco remains, worldwide, one of the leading causes of preventable mortality. One key public health intervention used to reduce consumption of cigarettes, and increase health outcomes, has been the excise tax. While much has been studied and learned regarding price elasticities and consumption rates in the light of excise taxes [20], in 2016 the National Cancer Institute published a monograph [101] which describes the current level of understanding in this regard. Two areas not well understood were stated. Firstly, price elasticity is not known over time or for different levels of tax and price changes. Secondly, price elasticity for non-cigarette products or cross-price elasticities, are not understood. They state explicitly, “these studies will be very useful.”

Additionally, tobacco companies are not resting on their laurels. Chaloupka [20] describes that they have engaged in significant marketing of price-minimizing behaviours (such as buy-one-get-one free deals, coupons, or free-samples). This opens a variety of other possible policy options, such as restricting types of sales, or locations of stores. Chaloupka indicates that understanding the effect of policies under these conditions is an important area of future research.

The use of the Discrete Choice Experiment (DCE) is motivated by the focus on price elasticities and purchasing behaviour, as well as being a very standard technique of studying human preferences. One of the assumptions of DCE is the rational agent but addiction is a non-rational process. As discussed in Chapter

2, addiction is rarely treated in dynamic models beyond statistical assessments, but it is no doubt a primary driver in tobacco purchasing. *Our goal, then, is to compare and contrast several different addiction models in an effort to select a single model for use in further research.* This is done with the construction of a dynamic simulation model that incorporates purchase decisions with a wide variety of tobacco options, and physical location of agents within homes, work, and tobacco retailers, while also providing explicit mechanism for consumption and addiction patterns.

5.2 The Model

I constructed the Tobacco Price Minimizing Strategy (TPMS) ABM, which explores the relationship between tobacco taxation strategies, price-minimizing purchasing behaviours, and consumption habits. The model accounts for population heterogeneity across education levels, income levels, and gender; spatial locations of stores; and pricing and availability of several product types (including non-cigarette options). To inform this model with empirical data, we are using results from a well defined and independently validated survey tool, the Discrete Choice Experiment to provide the purchasing behaviour of agents. Addiction behaviour is captured through an addiction module.

The TPMS model starts with the modular ABM described in Chapter 4 where agent consumption patterns are captured with a statechart and reference to an addiction module. As a reminder, consumption, initiation, quitting, and relapse for each agent are controlled by an addiction module, defined using a Java Interface, which can be designed according to a variety of different addiction theories. It was demonstrated using 2 different modules: one using a simple Markov chain, and another using a complex design that relies on social information and past use. However, the TPMS model is a significant extension in two ways. Firstly, agents are placed within a stylized city meant to simulate a large US metropolitan area. Secondly, and most importantly, agents in this model are given the ability to purchase tobacco products. This behaviour is parameterized by a Discrete Choice survey study recently completed with over 1,800 survey responses. Both aspect are discussed below.

This work was done in AnyLogic (version 8.1.0). It was chosen for several reasons. In spite of being closed-source and proprietary, and therefore requiring a licensing fee, there is are many options for experimentation. As discussed below, calibration and stochasticity experiments are a key component of this work. Many simulations of this model are required, and using AnyLogic's mechanisms, we exported the model file to run on a combination of in-house desktop machines and online cloud servers.

AnyLogic also possesses a rich set of visual tools to communicate model structure, which facilitated strong communication with our stakeholders. A more transparent model structure also reduces model bugs and improves development time. This is especially true since many of the building blocks of a dynamic model are, in AnyLogic, available with simple drag-and-drop methodology. This also makes it very simple to construct hybrid models, using elements of both ABMs and SD, as well as Discrete Event Simulation.




Figure 5.1: Example of stylized city. There are 3 radial sectors. From the inside out these represent urban, suburban, and rural areas respectively.

5.2.1 Environment

A visual representation of the stylized city is depicted in Figure 5.1 where the red lines indicate the radial zone boundaries, and the various building images represent either homes, schools, or tobacco retailers. This city has 3 zones which, working from the inside out, correspond to the dense urban city core, a suburban region, and the rural area, respectively. The model allows for homes, stores, and schools to be placed with densities in each zone matching aggregated data from US cities. The home objects are parameterized with an education and income level of their occupants, which are defined based on population data and their respective city zone. Also from population data, we set a home size for the home, though only 1 agent is actually placed in each home. This home size is used because the Discrete Choice Experiment accounts for that variable. Finally, the prices and stock availability of each store are determined based on data for stores in matching zones in real cities. We set the location and prices per city zone based on data for the Minneapolis Metropolitan Statistical Area.

5.2.2 Discrete Choice Purchasing Behaviour

Purchasing behaviour was determined using a Discrete Choice Experiment survey study that we conducted over the course of this research. Details of the DCE framework are below, but central to the DCE approach is the *decision context*. Participants are instructed to place themselves in an imagined scenario, of which there are two: one for planned decisions (e.g., “You are at home and have decided to buy cigarettes”), and one for impulse decisions (e.g. ”You are in a store for another purpose”). The question presents store options with prices, informs participants of any possible taxes, and asks them to choose what they would purchase in this situation. To allow for DCE results to be used within the TPMS model, we recreate the same decision contexts for the agents in our model. Therefore, we assume that current smokers in the TPMS model make a planned decision 3 times per week, while they make an impulse decision once per week. The output from the DCE is a set of probabilities for a given decision context depending on the deciding agent’s education level, income level, gender, and their current cigarette inventory size. This should result in agents that behave in a manner statistically similar to the respondent pool.

Background

In health policy and economics, knowing the preferences, health and otherwise, of patients or health professionals, or indeed the general public, is of great value. When data about the actual choices people make, i.e., *revealed* preference data, is not available, researchers require methods of eliciting specific preferences through tools such as surveys. This is called *stated* preference data. This is common, for example, when preferences are sought for products or services that are not yet available. They are also able to measure what factors of a given product, beyond price, influence preference [66]. However, they come with the burden of not being actual choices, meaning that any discrepancy between what people think they would do and what they actually do is not easily captured. Another limitation is that it is difficult in general to capture behaviour long after a potential intervention, since this becomes hard for survey respondents to imagine.

DCE is one such preference elicitation method. It has a broad application base, and is very common for studies in health economics and policy [26]. It is based on random utility theory and econometric analysis. As outlined in the textbook by Train [98], DCE theory begins by assuming that people make decisions based on both observed (x) and unobserved (ϵ) factors, as captured in the decision function, $h(x, \epsilon)$. For a given binary decision of action or no action, we assume a utility, either positive or negative, for taking the decision action. This utility is the weighted sum of each observed factor plus the overall unobserved factor.

$$U = \beta x + \epsilon \tag{5.1}$$

where β is a vector of parameters and x is a vector of variables.

The probability of choosing this option, therefore, is the probability that $U > 0$. Since ϵ is the only

unobserved piece, it is weighted by the distribution of the unobserved factor, $f(\epsilon)$.

$$\begin{aligned} P(y|x) &= P(I[U > 0] = 1) \\ P(y|x) &= \int I[U > 0] f(\epsilon) d\epsilon \end{aligned} \tag{5.2}$$

where $I[a]$ is the indicator function (returning 1 if a is true and 0 otherwise). Specifying a DCE involves specifying observed factors, assuming a form of f for the unobserved, and solving $P(y|x)$ for each decision, either analytically or numerically. When decisions are not binary, the probability of each choice relative to a common alternative (in our case, this is the no purchase option) can be found and normalized across the choice set for a given decision context.

As is commonly done, and in order to allow for a closed-form expression, we assume that $f(\epsilon)$ is distributed logistically, where $F(\epsilon) = \int f(\epsilon) d\epsilon = 1/(1 + e^{-\epsilon})$. This enables an analytic solution for the probability of making the decision, y .

$$P(y|x) = \frac{\exp(\beta_i x_i)}{(1 + \exp(\beta_i x_i))} \tag{5.3}$$

When dealing with more than two options, the denominator is represented as a sum over the available options.

$$P(y = 1|x = \{x_0, x_1, x_2, \dots, x_n\}) = \frac{\exp(\beta_1 x_1)}{\sum_i \exp(\beta_i x_i)} \tag{5.4}$$

Study Implementation

The DCE survey consisted of 12 DCE questions presenting various purchase options and asking for decisions. An example of a survey question used in the DCE used here to elicit preferences with respect to tobacco purchasing is depicted in Figure 5.2. The top section provides the respondent with the decision context, the number of cigarettes currently in their possession, whether any tax is planned on packs or cartons, and the tier of their chosen brand (answered in an earlier question of the survey). The decision context is either meant to elicit a planned purchase (*at-home*) or an impulse purchase (*in-store*). Below that is a table that displays anywhere from 1 to 3 stores for the at-home context, or just a single store for in-store. Each store has different products and prices, and, for the at-home context, the distance from the person to the store. This distance option was calibrated with an earlier question in the survey. The stores in the model and the DCE survey carry many different possible products, each with a different price: cigarettes in packs or cartons and in low, medium or high tiers; e-cigarettes; roll-your-own tobacco; and chewing tobacco. The survey also asks demographic and smoking use questions.

We designed the DCE survey itself in consultation with Dr. Joffre Swait of the Institute for Choice at the University of South Australia. Our study population was defined as current US smokers between the ages of 18 and 30. We also categorized each person by 1 of 3 education categories - high school or below, technical school or some college, and bachelor's degree or above - and 1 of 3 location residences - rural, suburban, and urban. Based on the required sample size for each sub-population, and in consultation with Dr. Swait, we decided that we needed a minimum of 1,500 participants.

You are **at home** and have decided to buy tobacco. The stores below are your options.

You have **1.5** cigarettes left.

Planned tax increases:

- Packs have **no tax** planned.
- Cartons will increase by **\$1.00 in 3 months**.

Your chosen regular cigarette brand is considered **MEDIUM** tier.

	Store 1			Store 2			Store 3			Opt Out
Distance	4 minute ride			4 minute ride			1 minute walk			No Purchase <input type="radio"/>
Cigarettes	Tier	Packs	Cartons	Tier	Packs	Cartons	Tier	Packs	Cartons	
	Low	\$6.46 \$5.81 <input type="radio"/>	\$58.15 <input type="radio"/>	Low	\$5.17 \$4.13 <input type="radio"/>	\$54.92 <input type="radio"/>	Low	-	-	
	Medium	\$7.12 \$5.70 <input type="radio"/>	-	Medium	-	-	Medium	\$8.90 \$8.01 <input type="radio"/>	-	
	High	-	-	High	-	\$127.04 <input type="radio"/>	High	\$8.84 <input type="radio"/>	-	
Other tobacco products	Product	Price		Product	Price		Product	Price		
	Disposable e-cigarette	\$9.95 <input type="radio"/>		Disposable e-cigarette	\$9.95 <input type="radio"/>		Disposable e-cigarette	\$7.96 <input type="radio"/>		
	Chew	\$7.12 <input type="radio"/>		Chew	\$7.12 <input type="radio"/>		Chew	-		
	Roll-your-own with 50 rolling papers	-		Roll-your-own with 50 rolling papers	\$14.95 <input type="radio"/>		Roll-your-own with 50 rolling papers	\$16.82 <input type="radio"/>		

Question 9/12

Figure 5.2: Sample DCE survey question.

We then contracted a Qualtrics panel after extensive investigation into alternatives. As this DCE required formatting and customization beyond the capacities of the default Qualtrics survey, I developed a custom-built survey within the Qualtrics tool that dynamically generates DCE questions based on survey design and participant responses. We designed the survey for simplicity, brevity, and accuracy. Several experts were consulted, and we conducted user testing with smokers in Saskatoon. We used those results to improve the flow and wording of the survey. After obtaining approval from the Research Ethics Board for the study with the Qualtrics panel, we conducted an initial pilot, which enabled us to examine the data, compliance rates and times, and to make final changes for a full study. In the end, we received over 1,800 responses. A key advantage to using a panel company is that they can provide some assurance that the respondents are in the correct age category, that they are active smokers, and can control for education level. Nevertheless, the survey included 4 questions at the start meant to confirm these assumptions from the respondents and used to filter out unqualified participants.

Formulation

We need a utility for each individual, i , using alternative j . The default alternative, when $j = 0$, is when the individual chooses no purchase. The observed utility, \hat{U}_{ij} (i.e. without the unobserved ϵ term) of individual i for alternative j , is composed of factors related to the individual, α_{ij} , and to the alternative, β_{ij} . The default alternative is defined to have a utility of 1.

$$\begin{aligned}\hat{U}_{ij} &= \alpha_{ij} + \beta_{ij} \\ \hat{U}_{i0} &= 1\end{aligned}\tag{5.5}$$

The individual preference of person i for alternative j is α_{ij} , and depends on the set of controlled individual characteristics, $c = \{\text{education, income, gender, inventory size, household size}\}$.

$$\alpha_{ij} = \alpha_{j0} + \sum_c \alpha_j X_{ic}\tag{5.6}$$

where α_{j0} is the intrinsic preference constant for that option (uniform across individuals), and α_j is the sensitivity of the individual's decision to their personal characteristics.

The alternative's factors are:

$$\beta_{ij} = \beta_p p_j + \beta_d d_j + \beta_r r_j + \beta_f f_j + \beta_t t_j + \beta_{ft} f_j t_j\tag{5.7}$$

where p_j is the price for alternative j , d is the price discount, r is the distance to the store in minutes of travel, f is the tax increase amount, and t is the tax increase time. The last term is an interaction between tax amount and tax time. The β terms describe the sensitivity of the agent to the respective factor.

The DCE analysis generates a table of values for α_{j0} , α_j , and $\beta_{p,d,r,f,t}$. Within the TPMS model, then, we recreate the individual and alternative characteristics. The agent population is created with an education, home location, and gender, and their cigarette inventory is tracked dynamically. Stores are created with

stock, prices, and potential discounts, and their distance from the agent is calculated by the agent at the moment of purchase. The intervention controls the present or absence of a tax on either cigarette packs or cartons. No tax is available for non-cigarette options. And in this version of the TPMS, no experiments have been defined that specify discount values. This was done for future work.

There are 9 alternatives for each store - 3 tier levels for packs and cartons, chew, loose tobacco, and e-cigarettes - plus the no-purchase option. Then, using eq. 5.4, the agent is able to determine a probability for each option in any given simulated decision.

DCE calibration of these parameters is done with the R package `mlogit` [25] and uses the standard maximum likelihood estimation [68].

5.2.3 Addiction Modules

The four addiction modules that were studied in the TPMS are described below. Each is based on different proposed mechanisms of addiction dynamics. Following the protocol from Chapter 4, we use a Java Interface to specify the desired outputs of all addiction modules. From that chapter, Figure 4.2 indicates the smoking statechart. Each of the modules is required to provide, at any time requested by the agent, 4 hazard values: the consumption hazard, h_c ; the quitting hazard, h_q ; the relapse hazard, h_r , and the initiation hazard, h_i .

Inventory Driven Markov-chain Module

The Inventory Driven Markov-chain (IDM) modules focuses on a very simple mechanism for each of the 4 required transitions. Each is treated independently within this module. The main idea is to allow the agent to alter their use hazard h_c based on their inventory. This is done with the exogenous function, defined below, and in Figure 5.3.

$$h_c(I) = \begin{cases} I + 0.5, & \text{if } I < 0.5 \text{ days} \\ h_{c,base}, & \text{if } 0.5 < I < 2 \text{ days} \\ (m_{inv} - 2)I + 1, & \text{if } I > 2 \text{ days} \end{cases} \quad (5.8)$$

where $h_{c,base}$ is the base use rate and I is the inventory size in days of use, calculated using $h_{c,base}$, at any given time.

If an agent has between 0.5 and 2 days worth of cigarettes in their inventory, they do not modify their consumption frequency. If, however, an agent drops below 0.5 days worth, they will begin to reduce their consumption hazard. Similarly, if an agent possesses enough inventory for more than 2 days, they will linearly increase their actual smoking hazard based on the parameter m_{inv} .

The quitting hazard, h_q , is:

$$h_q = \begin{cases} 0, & \text{if } n_{np} < T_{np} \\ 1000, & \text{if } n_{np} \geq T_{np} \end{cases} \quad (5.9)$$

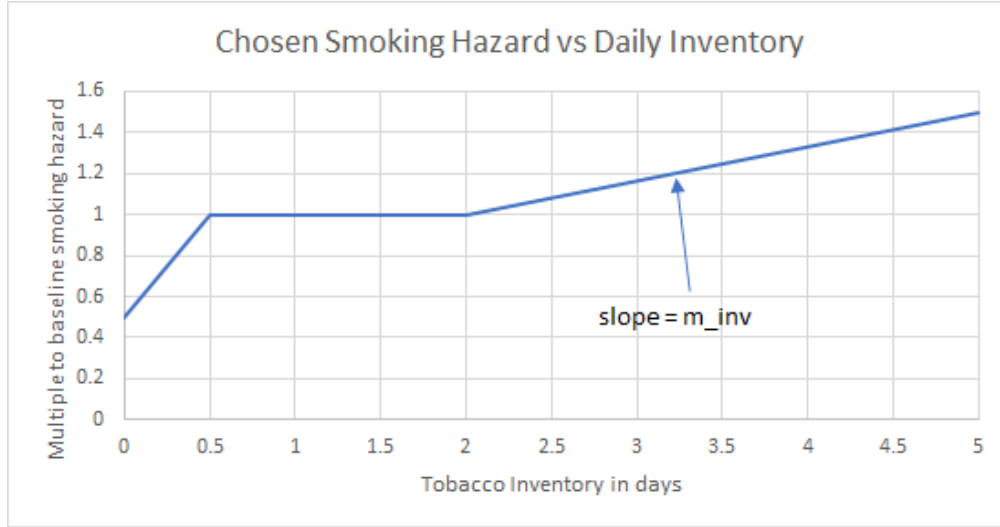


Figure 5.3: Impact of inventory on baseline consumption hazard

where n_{np} is the current number of consecutive “no purchase” decisions, and the threshold parameter is T_{np} . If the threshold is reached, the cessation hazard spikes to 1,000 (chosen simply to make quitting immediate once the threshold has been reached).

The relapse hazard, h_r , is calculated using a learning process adapted from the Temporal Difference Learning approach of Hammond et al. [36].

$$h_r(t) = C_r h_{c,base} (1 - r_{learn})^{t-t_c} \quad (5.10)$$

where C_r is a fitted constant parameter, r_{learn} is the learning rate and is another fitted parameter, t is the model time, and t_c is the time the agent ceased smoking. This captures the idea that relapse is more likely to occur in the time immediately after a quit attempt, and gets progressively less likely the longer that it is maintained.

Finally, for initiation, we expect this process to fundamentally depend on things outside the agent, in their social and physical environment. Since this module is focused solely on factors internal to the agent, we abstract over those external factors with a simple constant initiation hazard, h_i .

Overall, this module requires 6 parameters to be defined for a given agent: for consumption, we need m_{inv} and the baseline smoking hazard, $h_{c,base}$; for quitting, we need n_{np} ; for relapse we need C_r and r_{learn} ; and for initiation we need h_i .

Perceptual Control Theory Module

The PCT module is the same theory as was described in Chapter 4. As a reminder, it uses a double feedback-loop systems and tracks the intent and action of addiction over time. Figure 4.4 presents the loop system using SD stocks and flows. One advantage of the PCT over the IDM is that it accounts for social influence and gives agents an adaptive internal structure. This comes at the cost of increasing the number of calibration

parameters required to 9: the intent set point p , the initial stock values for action, A_0 and intent, I_0 , the ability to change either stock a_I and a_A , the weight of addiction on action w , and the impact values for external influences x_{own} , x_{friend} , and x_{neigh} .

However, in the TPMS model, the PCT was implemented differently. The earlier version was implemented using AnyLogic’s SD elements, but this proved to be prohibitively slow in the TPMS model. This is because numerical methods are generally required to solve an SD model. SD represents a system of continuous-time equations, but because SD models are commonly nonlinear, the analytic (closed form) solution is not in general available. Numerical solutions require frequent updating of model values using a relatively short event loop, resulting in the PCT parameters being updated perhaps 100 times or more each model day for each agent. However, the PCT is a system of first-order linear differential equations, enabling me to solve it analytically with Laplace transforms [55].

After bench-testing the analytic solution against the AnyLogic SD implementation to ensure agreement, the PCT was implemented with this analytic solution. However, because this model is an ABM/SD hybrid, the SD solution does not include the agent’s behaviour. The continuous PCT solution needs to be updated whenever one of its values is discontinuously altered by the ABM section of the model. When such an update takes place, the current stock values become the new initial values, and the time t is shifted by the update time. This means that we do not need to know the solution beyond the next event. This resulted in substantially improved run times over the numerical option, with the added benefit of being more accurate in principle than would result from traditional numerical integration of the AnyLogic SD model.

It is represented here in the form of differential equations, where $A(t)$ the addiction action and $I(t)$ the addiction intent, a_A is the ability to change action and a_I the ability to change intent, w is the weight of addiction on action, p_I is intent set point, and x is the external signals.

$$\begin{aligned}\dot{A}(t) &= -a_A A(t) + a_A I(t) - a_A d_1 \\ \dot{I}(t) &= -a_I w A(t) - a_I I(t) + a_I p_I - a_I w x\end{aligned}\tag{5.11}$$

FlexP

The FlexP module is built as an extension to the PCT module. A significant limitation to the PCT module is that the intent set point does not change over time. For example, when agents have a positive value for the set point, without consistent social or environment drivers resisting change, over time they will inexorably move towards the smoking state. Essentially, PCT agents are unable to change their fundamental underlying goal state. And since we cannot distinguish between agents beyond their sub population for our calibration data, we do not distinguish between them for the addiction parameters. Therefore, all agents within a specific sub population have the same set point, regardless of initial smoking state.

However, FlexP allows the process of addiction to change the internal set point over time. Whenever the module is updated, the set point steps towards the average of the addiction since the last update. To achieve this, I added a 10th parameter to the 9 already defined by the PCT module, p_{learn} , and integrated

the analytic equation for the addiction action to calculate the average action over an arbitrary time span. The update process is, therefore:

$$p_{t+\Delta t} = p_t + p_{learn}\Delta t(\overline{action}(t, \Delta t) - p_t) \quad (5.12)$$

where p_t is the intent set point, Δt is an arbitrary time between successive module updates, and $\overline{action}(\cdot)$ is the average action over that time span. In this formulation, p_{learn} controls the speed of learning, or the flexibility of p .

Bobashev Control Theory

The final module developed for this study is another control theory approach, developed by Bobashev et al. [15]. Bobashev et al. describe how the consumption mechanism has been validated in studies involving mice and cocaine use. We term this the Bobashev Control Theory (BCT) module.

The BCT module attempts to simplify the complex neurobiological dynamics of addiction using two simple processes. The first is homeostasis, maintained by an opponent process. Homeostasis is when the body has a “set point” and seeks to adapt its internal structure so as to maintain that set point. Under the experience of a strong negative or positive external stimulus, the opponent process will, over time, counteract this stimulus in order to maintain homeostasis. In the language of addiction, one opponent process would be the withdrawal state - when a regular-user stops using, the body seeks to return to the previous homeostasis through the uncomfortable feeling of withdrawal. In other words, the addiction process is more determined by wanting a drug than by liking a drug.

The second process is allostasis. This builds on homeostasis to allow for slow, long-term changes to the set point. The more an agent uses a particular reinforcer, the more difficult it is to achieve satiation, which requires larger doses and therefore the development of a dependence.

Originally defined in terms of a system of ordinary differential equations, presented below, the core of the model is also reproduced in Figure 5.4 using system dynamic stocks and flows. This approach seeks to capture the basic control-theory behaviour of an addictive process without needing to describe the complex neural processes of the human brain. The analogy used in [15] is that of a walking robot: when a person is walking, many complex processes are at work, but a walking robot can demonstrate similar behaviour in maintain balance using coordinated feedback loops. This model, therefore, is a “smoking robot” with feedback processes that can only be loosely mapped to a human smoker.

$$\begin{aligned} y_1 &= D(e^{-a\Delta t} - e^{-b\Delta t}) \\ \dot{y}_2 &= a_1y_1 - b_2y_2 \\ \dot{y}_3 &= a_2y_2 - b_3y_3 \\ \dot{y}_4 &= a_3y_3 - b_4y_4 \\ \dot{y}_5 &= a_4y_4 - b_5y_5 \end{aligned} \quad (5.13)$$

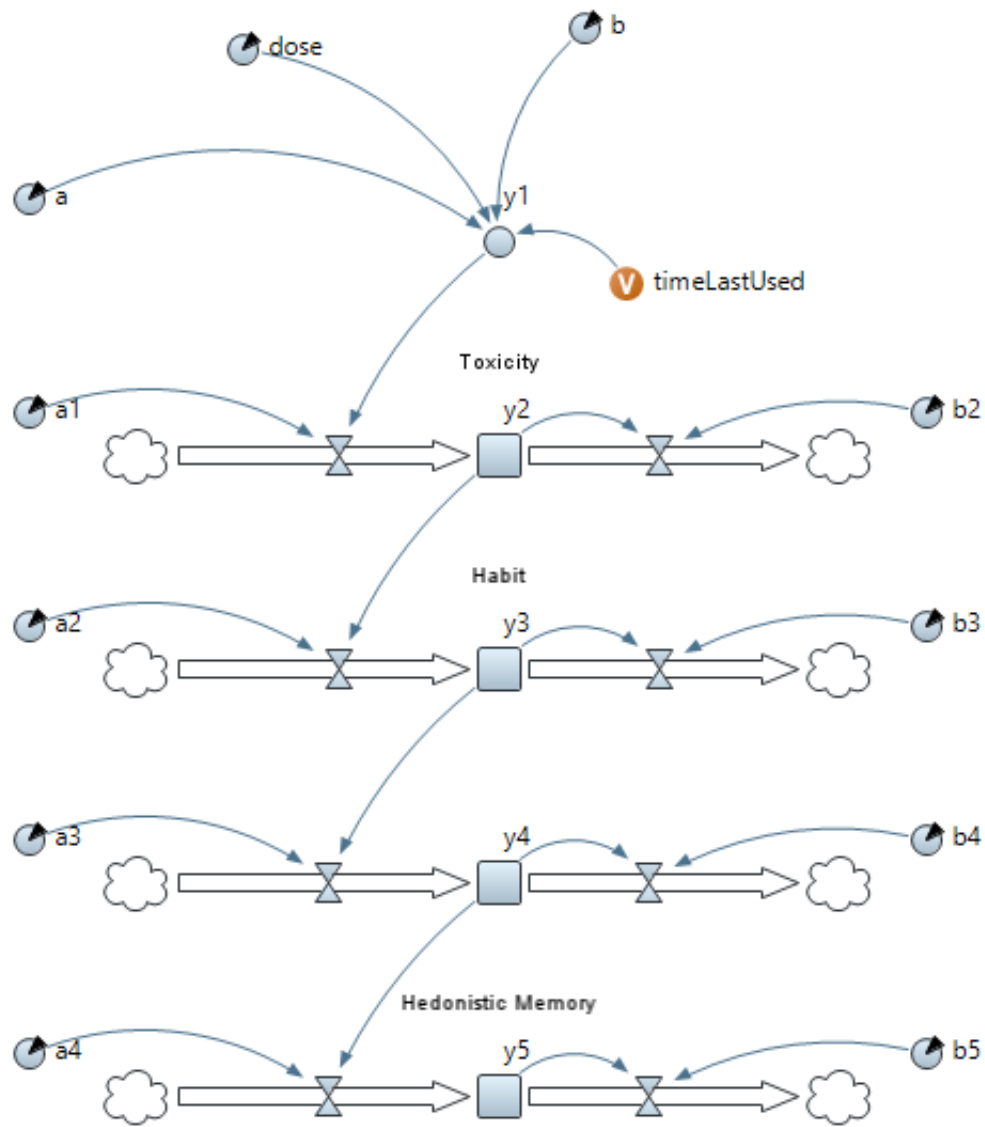


Figure 5.4: Stock and flow depiction of the Bobashev Control Theory approach.

Examining Figure 5.4 or eq. 5.13, there are 5 terms, y_{1-5} , that characterize the dynamically associated with several key types of feedback process. Each is placed in a successive running weight average of the former, with the first, y_1 , defined as something akin to the drug effect, or the amount of drug in the system. This term is modelled with a pharmacokinetic equation where D is the dose of each drug use, a and b are parameters that control the rate of drug absorption and dissipation in the body, and Δt is the time since last use. D , a and b are set at the values by Bobashev et al., which were defined for tobacco. I would note that after attempting to reproduce the results from the publication, in conversation with the original author, and examination of the code used to generate those results, it was found that the equations as given in the paper are not completely correct. I have corrected them here.

The next process, y_2 , is somewhat related to toxicity, which is the first opponent process. Habit is related to y_3 , defined on the scale of days. y_5 is understood as relating to long-term memory. It is defined on the year time scale, enabling the relapse of longer-term quitters. y_4 is not defined but is included in the framework for mathematical consistency, so as to provide another opponent process. a_{1-4} define the impact of a lower process on a higher, and b_{2-5} controls the time scale of the process decay. We used the values as defined by Bobashev et al.

In order to determine whether an agent smokes or not - or more precisely, how often they smoke - these processes are brought together to define a tobacco consumption threshold, T_c .

$$T_c = \frac{(\beta_3 y_3 + \beta_5 y_5)}{(1 + \beta_2 y_2)(1 + \beta_4 y_4)} \quad (5.14)$$

where $\beta_{2,3,4,5}$ are calibration parameters. In the BCT module, the tobacco consumption hazard, h_c , is defined as the fractional difference from the mean of T_c and y_1 , whenever $y_1 < T_c$.

$$h_c = \max\left(2K \frac{(T_c - y_1)}{(T_c + y_1)}, 0\right) \quad (5.15)$$

where K is a constant chosen simply to ensure that the agent consumes tobacco very quickly when $y_1 = 0$. This means that the lower the value of y_1 is below the threshold T_c , the higher the chance of smoking. But when y_1 is above the threshold T_c , the agent is satiated and will not consume. And with a higher threshold, the agent will spend more time consuming tobacco due to the natural decay of y_1 . Important to this definition of T_c , the opponent processes are on the bottom so as to reduce the threshold value, thereby reducing the use hazard.

For implementing this module, several approaches were explored. While AnyLogic uses an Runge-Kutta method to solve the differential equations of SD elements, the exact implementation leads to significant increases in run times when compared with an numerical solution. This is because AnyLogic updates the equation parameters everytime an event occurs anywhere in the model. When there are SD elements in each agent, any event in any agent results in updating all the stocks, flows, and variables across the whole agent population. Therefore, I implemented two numerical solutions: my own version of the Classical Runge-Kutta method, and using an adaptive-step method from the Apache Commons Math library. These resulted in significant speed increases, of an order of magnitude or more, however, it was still slow enough that

calibration was very difficult. Finally, I used the Maple software package to solve the system of ordinary differential equations analytically. To make the solution tractable, the BCT parameters were defined in Maple - the solution hard-codes those parameter values. We deemed this less flexible option suitable given the substantial time-complexity improvement.

The focus of the original BCT work was to examine the dynamics of current smokers. As a result, it did not develop mechanisms for initiation or quitting. Therefore, to match the module interface, we added static hazards for each, h_i and h_q , respectively. This brings the total number of BCT calibration parameters to 6: $\beta_2, \beta_3, \beta_4, \beta_5, h_i, h_q$.

A final note on implementing the BCT module also bears discussion. It returns hazard rates that vary significantly over the course of a single use - when y_1 crosses below T_c , the use hazard rate increases to the maximum, $2K$, within 2 hours (determined largely by a in equation 5.13), decreasing to 0 again within minutes after use (determined by b in equation 5.13). When the modular design pattern was implemented, this type of behaviour was not considered - the other modules return more stable hazard rates over time. Therefore, I had to make 2 small changes to the agent to accommodate this module, creating a version branch. These changes do not limit the use of the other modules, but there might be some very small changes to their behaviour. This branch has been merged in an updated version of the model. Since the experiments were already conducted for them, I did not implement these changes in those experiments.

5.3 Experiment

In 2009, the state of Minnesota introduced a tobacco excise tax of \$0.62 on all cigarette packs. Using data from the 2006 and 2010 waves of the Minnesota Adult Tobacco Surveys (MATS), we recreated this historical record in the TPMS model, including both pre- and post-tax statistics. As with the DCE, we used only results from ages 18-30, and stratified the population by the same education and home location categories as used with the DCE survey. We used these categories to define 9 sub-populations, indicating the home location (Rural, Suburban, and Urban) and education (high school - HS, technical school or some college - SC, and bachelor's degree or more - BA). For example, an agent living in the rural area with some college or technical training would belong to the RurSC sub-population.

For each survey year, we obtained data on the proportion of never smokers, current smokers, and former smokers. In accordance with widespread practice, never smokers are defined as those that have smoked under 100 cigarettes in their lifetime. Furthermore, we also used the self-reported average daily smoking rate for current smokers. This gives 3 data points for each of the sub-populations in each survey year, and 6 for a single experiment. The number of responses for each sub-population and survey year is shown in Table 5.1.

Several simplifying assumptions were used here. Firstly, the survey itself took place over several months, from May 2010 to January 2011; the data therefore does not represent a single snapshot of the population. In comparing the data from the TPMS model to that from the population, we set the time of the first survey

Table 5.1: Counts for survey responses per sub-population and year.

Year	RurHS	RurSC	RurBA	SubHS	SubSC	SubBA	UrbHS	UrbSC	UrbBA
2006	114	113	28	128	99	96	57	60	90
2010	29	54	52	183	97	126	87	38	68
Total	143	167	80	311	196	222	144	98	158

within the model to occur in August. The model runs for 4 years, but the model population is closed - there is no aging out or into the population over time. As a result, the agents in the simulation are effectively 4 years older by the end of the model. Therefore, the 2010 MATS data is for an age range of 22-34. However, the DCE data is for the age range of 18-30. Given that the age overlap is still 8 years, we are assuming that the DCE behaviour still provides a sufficiently accurate characterization of population behaviour to correspond to the simulation model.

5.3.1 Validation

Since the DCE is a central element of this model, the predictions of the DCE needed to be validated. Firstly, the DCE predictions were cross-validated. As described earlier, we garnered 1,800 DCE responses. We used 1,500 responses to fit the DCE model, and matched the predictions from that fitted model with the remaining 300 responses. The result showed a strong match for the purchasing behaviour of the chosen sub-populations.

Additionally, since the results of the DCE analysis include an estimate and a standard error for each parameter, the statistical significance can be assessed. While many of the DCE parameters were significant, this model had 168 parameters (this is due to the large number of tobacco options at up to 3 stores). Many of these 168 parameters were not statistically significant (i.e. the null hypothesis of being equal to 0 could not be rejected). However, we cannot simply avoid using these statistically insignificant parameters in the ABM. Therefore, to allow use of all parameters, we bootstrapped the DCE results, resulting in many different parameterizations within the values determined by the estimate and the standard error. Those parameters that are not statistically significant will have many values both greater than lesser than 0. This is incorporated into the ABM through sensitivity analysis that is described in the sections below.

Next, we examined the purchasing behaviour of current smokers, since only current smokers can make purchases. This model assumes that purchasing and consumption behaviour can be described separately, therefore, we implemented the full DCE, but only a placeholder addiction module was used. The placeholder module sets a static smoking hazard that matches the historical mean daily consumed cigarettes per day for each sub population, and keeps the initiation, quitting, and relapse hazards to 0. The whole population is set to be current smokers. Within the same survey as the DCE, we asked other tobacco-related questions, including whether the last purchase was a carton or a pack, with over 80% reporting having purchased a pack. We were able to run the model to make sure that the pack purchases were much higher than for cartons.

We also examined the relative purchasing of cartons, packs, and other forms by gender and sub-population, and found these to be sufficiently appropriate for this study. Finally, we observed the average inventory of current smokers.

During this purchasing validation, we observed that the average inventory for some agents was much higher than plausible. Normally, we expect that with increasing inventory, the chance to purchase should reduce, meaning that the relevant α from eq. 5.6 should be negative. But some of the inventory α terms were not statistically significant, which means that for some DCE bootstraps, they were positive. In this case, the more an agent had in their stock, the higher the utility to purchase. The inventory was one of the parameters in the DCE survey, but it was always kept to below 1.5 days worth of cigarettes. Therefore, we added a limitation to the ABM that did not allow agents to make a purchase decision if their inventory was higher than 2 days worth of cigarettes. This led to some changes in logic and assumptions which increased validity.

One point of learning that emerged from this validation process was that DCE agents cannot be considered as samples from the DCE survey population. This is because the DCE cannot distinguish them within a sub-population. In effect, each agent makes its decisions based on the behaviour of the whole sub-population rather than any given agent within. This is evident from further data from MATS, reported by [21]: under 50% of the population in 2010 indicated that they had purchased cartons at least once in the previous year. However, in the TPMS, every agent from every sub-population purchased cartons at least once in a single year. This indicates that there is more that is important to a person's smoking behaviour beyond their education and location; their purchasing history indicates something about the purchasing future. This idea was not included in the purchasing component of this model, but could be considered as possible future research. However, we expect this limitation to be somewhat moderated by the fact that the addiction modules do allow for habitual consumption.

With each addiction module, we observed sample agents over time to look for obvious illogical smoking or purchasing behaviour. For example, since agents are unable to consume tobacco with no inventory, with some modules we observed agents smoking all their inventory within a few hours after purchase. The monthly and population means were appropriate, but the micro-behaviour was erratic. This allowed quick exploration of different addiction module parameter values, the changing of module logic, exploration with different variations of a given module, or even the construction of entirely different modules. Over the course of this work, we designed and explored to varying degrees 12 separate addiction modules - variants as well as distinct designs - before selecting the 4 in this research.

5.3.2 Calibration

Calibration makes use of the AnyLogic *calibration experiment*, which repeatedly runs a given experiment, with each iteration varying the input parameters in order to optimize a user-created objective function. Our objective function is the square of the fractional difference between the historical and simulated values

relative to their mean, summed over all comparison values (i.e., pre- or post-tax, and consumption rate, and proportion of current and former smokers):

$$D(h, s) = \sum_i \left(\frac{h_i - s_i}{\frac{1}{2}(h_i + s_i)} \right)^2 \quad (5.16)$$

where h and s are vectors containing the historical and simulated values for each comparison, respectively. This function was chosen because it possesses a number of attractive qualities for objective functions. Firstly, it is dimensionless, which allows the consideration during within a single additive objective function of datapoints data with different dimensions. The function is also concave, meaning that two small discrepancies spread across matches to two data points are considered better than a single large one with respect to one of those datapoints alone. The function is further symmetric, meaning that, for some proportional error of α , the difference function returns the same whether it is an over-estimate or an under-estimate. The function is also non-negative, which protects individual error terms from cancelling each other out. Finally, the function is finite with non-negative historic and model estimates, except in the case where both the historical and the simulated values are 0, in which case we simply define the objective function to be 0.

Given that each addiction module postulates different behavioural mechanisms, and the fact that these mechanisms might differ across each sub-population, we conducted calibration experiments with each of the 9 sub-populations independently. Two things varied between scenarios: the addiction module and the sub-population. We then ran optimization experiments for each combination of 4 addiction modules and 9 sub-populations. Each experiment consisted of a sufficiently large number of iterations for each optimization until there were no changes to the best cross-realization mean objective value for at least 500 runs. The parameters found for the best-fit result were recorded and used in the sensitivity analysis described below. To account for stochastic variation, we replicated each unique parameterization 10 times and averaged the output.

A few experiments, however, were run for longer. For example, the FlexP module initially returned a higher objective value than the simpler PCT when simulating the RurHS sub-population. Knowing details about model implementation allowed us to reason that likely the FlexP calibration had only found a local minimum. This was verified after about 3,000 iterations. Using the PCT with the same sub-population only required 1,220 iterations to achieve its maximum.

5.3.3 Bootstrapped DCE Sensitivity Analysis

All of the calibration experiments were conducted using only one of the many bootstrapped DCE parameter sets. This was done to reduce the computational time for calibration. However, to account for the variability due to DCE bootstrapping, we conducted a sensitivity analysis using the parameters reported from each calibration experiment. Using the best calibrated parameters for each scenario, we randomly selected 50 bootstrapped DCE parameter sets with which to run the TPMS model, resulting in a distribution of 50 values whose max and form the error bars in the plots below. As above, to account for stochastic variation, we




Figure 5.5: Best objective values for each module and sub-population.

replicated each parameterization 32 times, taking as output the average output value across all 32 replications.

5.4 Results

The fitted objective values for all sub-populations and modules are presented in Figure 5.5. The sub-populations are listed along the bottom. The coloured bars indicate the addiction module used.

The error bars in this figure, and all figures below, indicate the max and the min of the sensitivity analysis. The RurSC/FlexP scenario has a noticeable uncertainty. Further analysis into that scenario shows that the variation occurs within many of the bootstrap values rather than across them. In other words, there does not appear to be a privileged bootstrap index; all of them showed variation in the objective value. The remaining scenarios showed substantially less variation.

Figure 5.6 examines the most accurate module for each sub-population, showing the historical mean cigarettes per day (orange and yellow) and the simulated (blue and grey) for the best-fit parameterization, as measured by the best mean objective value. Figures 5.7 and 5.8 show the same, but with the proportion of current and former smokers, respectively. For each sub-population grouping, there are 4 bars. The first two are the simulated and historical values for pre-tax value, respectively, followed by the same for post-tax. For each grouping, the simulated and historical values are in pairs. The bar heights for all these figures are calculated from the mean of the 50 bootstrap sensitivity iterations, each iteration consisting of 32 replications to account for stochastic variation. The error bars are the max and min the same set of 50 results.




Figure 5.6: Comparison of the best-fit module by sub-population between the historical mean daily smoking rate and the simulated daily consumption hazard. Calculated from the mean of 50 bootstrap iterations. The error bars represent the max and min of all bootstrap sensitivity experiments.

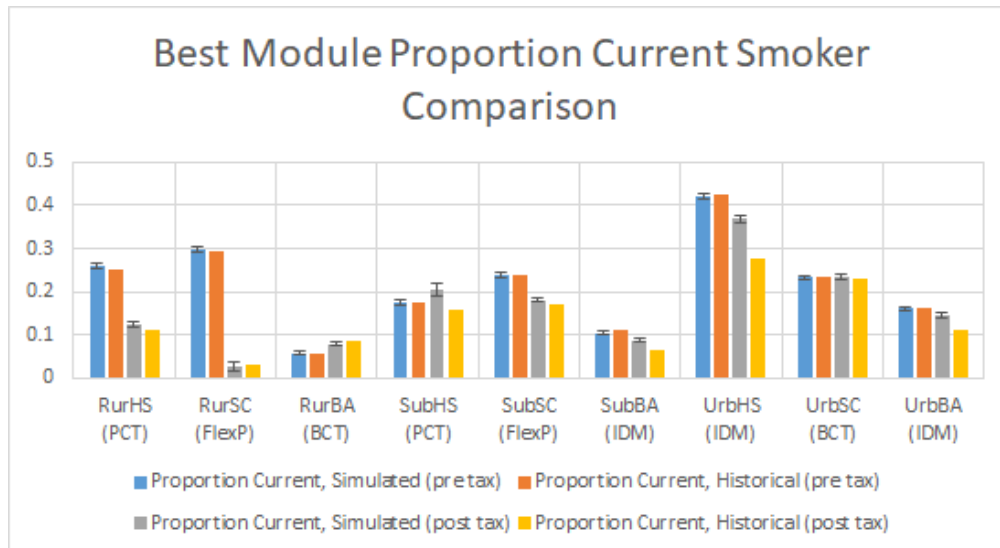


Figure 5.7: Comparison between historical and predicted values of the proportion of current smokers of the best-fit module, by sub-population. Calculated from the mean of 50 bootstrap iterations. The error bars represent the max and min of all bootstrap sensitivity experiments.




Figure 5.8: Comparison between historical and predicted values of the proportion of former smokers of the best-fit module, by sub-population. Calculated from the mean of 50 bootstrap iterations. The error bars represent the max and min of all bootstrap sensitivity experiments.

Figure 5.9 shows the fractional difference (using eq. 5.16) between the historical and simulated values, summing over all sub-populations and pre- and post-tax data. This is shown for each addiction module. Here, we can see how each addiction module performs for each of the 3 measures in the data (i.e., consumption hazard, fraction of current smokers, and fraction of former smokers), as well as the summed total of all predictions by the given module.

5.5 Discussion

The first observation from Figure 5.5 is that no module was the best-fit for all sub-populations. For most sub-populations, 1 or 2 different modules appear to achieve similar levels of accuracy. RurSC is unique in this regard since FlexP performs substantially better than the others, though it displays an especially large variation. From Figures 5.6, 5.7 and 5.8, we can see that IDM, the simplest mechanism of all the modules, performs quite well, achieving the best fit in 3 sub-populations. The other 3 each best fit in 2 other sub-populations. Relatedly, some sub-populations are more easily fit than others. For example, RurSC was the most poorly fit for 3 of the modules. Reference to the orange and yellow bars in Figure 5.6 shows that this sub-population also reported the highest change in smoking rates, and the highest change in current and former smokers. FlexP was the most adaptive module in this case. On the other hand, RurBA, SubSC, and UrbBA were fit well by all 4 modules. A possible explanation is that each sub-population is dominated by different underlying mechanisms of behaviour. A diversity of approaches would be warranted when seeking to examine whole-population behaviour.

More broadly, when examining the consumption rate accuracy in Figure 5.6, we can see that, while some




Figure 5.9: Fractional difference between historical and predicted values summed across all sub-populations for each module, including the total difference.

predictions are quite poor, 8 of the 9 best predictions capture the correct direction of change, either increase or decrease, after the tax. This is important because each sub-population has quite a difference response to the tax: in 5 sub-populations, we see an increase in smoking rate after the tax, in 2 we see the opposite, and in 2 we see no appreciable change. This diversity of response further supports the hypothesis that each sub-population makes decisions according to different dominant behavioural mechanisms. For example, most populations increased their consumption rate post-tax likely because of the normal addiction process. Since the agents are all 4 years older post-tax, they will have had more time to grow more addicted. Another possibility is that those who might more easily quit have the chance to do so leaving behind the more highly addicted to represent a larger proportion of the remaining smoker population. It could also be due to their personal wealth increasing, enabling more consumption. But 4 sub-populations did not experience a smoking increase, which implies that another mechanism is also at work there. There could be a growing awareness of health consequences with age. Or perhaps it is related to risk behaviour. It is known that risk behaviour reduces with age but is positively correlated with smoking, so older agents might also be more risk averse, and therefore more likely to quit. Whether the population as a whole increases, decreases, or remains the same would depend on how these diverse behaviours are distributed. More simulation and experimentation could be done to draw this out more fully, but it speaks to the importance of being able to not only reason about mechanism, but experiment with it.

Figure 5.7 compares the predictions of the most accurate modules by proportion of current smokers. The pre-tax values (blue and orange) show better agreement than post-tax (grey and yellow), which is unsurprising given that pre-tax values are near the start of the model run. The predictions for the proportion of current smokers are, in general, more accurate than for consumption rates, though for 4 populations, the post-tax

prediction over-estimates the historical value, which includes all 3 of the IDM experiments. Looking at all the IDM predictions, 8 of the 9 predictions for current smoker proportion post-tax are over-estimates. It is possible that, especially when the proportion of current smokers reduces post-tax, the IDM is not flexible enough to make up the difference. More experimentation with this module is warranted.

Figure 5.8 compares the predictions of the most accurate modules by proportion of former smokers. As with Figure 5.7, the pre-tax values (blue and orange) show better agreement than post-tax (grey and yellow). And again, the predictions are, in general, more accurate than for consumption rates. The largest differences appear as over-estimates in 3 sub-populations, which includes both of the FlexP experiments. Looking at all the FlexP predictions does not conclusively show consistent overestimation of this number, however.

Figure 5.9 shows the summed fractional difference from historical for each of the 3 metrics - consumption rate, fraction of current smokers, and fraction of former smokers - as well as a total for all objective scores. It was not surprising that BCT was the best match overall for consumption rate as it was developed specifically for that purpose. Our BCT implementation does not posit sophisticated theory for initiation or quitting, and uses instead static initiation or quitting rates, so it was not able to accurately reproduce those trends.

Our goal was to compare and contrast different addiction modules in order to inform module selection in the next steps of this research. Taken together, three of the modules have value in being pursued. IDM was overall the second most accurate and required the fewest number of parameters. BCT shows great promise since the consumption dynamics were very well captured, though it would require some work to improve the other aspects. FlexP was the most accurate, especially when it comes to predicting the fraction of current smokers. Though it required the largest number of parameters, with sufficient data a calibration process can help estimate their values.

One of the limitations of this work is that the TPMS model uses a closed population. This restricted our ability to perform cross-validation using data from another year. For example, with an open population, it would be possible to train the model on data from 2006 and 2010, and then to attempt prediction of 2014 data without resorting to aging the population. Using a closed population, from 2006 to 2014, our aged data would represent ages 26-38, which is significantly different from the 18-30 range used to parameterize the DCE. Beyond adding aging to the population, a conceptually simple option, though requiring computational time, would be to compare these predictions with a similar calibration process using the 2010-2014 data. This might shed further light on the behaviour of the addiction modules.

The TPMS model made restricted use of social networks, though the networks were not parameterized with real-world social network data. This is a clear limitation, but only 2 of our modules used social networks at all, and considering the other assumptions made, we judged this to be acceptable. Relatedly, we also constrained our scenarios to single sub-populations. This is another social network assumption, and made in this case both to focus our attention on the two aspects of addiction modules and sub-populations, and to improve the computational time. However, it is likely that there is in-group homophily, meaning that network connections are more likely within a sub-population than between them.

Overall, the most accurate module is FlexP which is not surprising since it carries more mechanism for all smoking transitions and has the greatest number of calibration parameters and thus requires more simulations to fit. However, both FlexP and PCT implement perceptual control theory and both account for social network and physical neighbours in the same way, and yet they are the most different in terms of ability to replicate the historical data. PCT possesses very nearly the same level of parameters and mechanism but it was the poorest performing module overall. Clearly, parameter count or mechanistic complexity does not, alone, lead to better performance. The two unique assumptions in FlexP - the flexible set point and the initiation of never smokers with a 0 set point - appear to be key in enabling the success of this module.

The design of FlexP was not included in the original scope of this project. However, given the surprisingly poor performance of PCT, we were able to use the outcomes from previous experiment to identify a key limitation in its design. This demonstrates a very clear strength with modular design approaches.

It should be recalled that the PCT module outperformed FlexP in 1 sub-population: UrbHS. We can think of 2 possible explanations. Firstly, FlexP has one extra calibration parameter than PCT. It is possible that, due to the larger FlexP parameter space, only a local minimum was found. More experimentation could draw this out. If this is not the case, however, the PCT assumption that everyone, including never smokers, has the same set point might be a better approximation in some cases. For example, it could be that some segments of the young adult population, in this case, urban youth with a high school or below education, might display higher risk behaviour. This behaviour could be seen as something of a set point. In other words, even never smokers might have an increased chance of initiating because of other high risk behaviour. If true, this means that non-tobacco behaviour could be a valuable predictor in smoking behaviour. Without more investigation, however, we are left with conjecture.

The data used to parameterize and calibrate this model came from a variety of sources. Smoking demographic data was taken from the MATS, which is prone to a number of potential biases. Recall bias, for example, might be a strong factor here since the survey was asking about smoking behaviour over the past month. A side-effect is that individuals are more likely to report smoking in multiples of 5 cigarettes per day. Another question is how representative this data is to the young adult population in the US, or even Minnesota. It has been adjusted using national population data to reduce this bias, but it is still present. Sample size is often an issue. Taking the population and dividing it into 9 sub-populations certainly reduces the sample size for each sub-population. Referring back to Table 5.1, we can see that most sub-populations have 50 or more responses, but 3 are below 40: RurBA in 2006, and RurHS and UrbSC in 2010. With greater sample sizes across sub-populations, we would be able to perform more rigorous validation techniques, such as k-fold cross validation. Finally, the MATS was conducted over a several-month period, but the model anchored the data as a specific time point for the whole population. Therefore, we chose an anchor point roughly in the center of the MATS survey period for both survey years.

Purchasing behaviour was driven by DCE results. The DCE resulted in a large number of parameters, some of which are not statistically significant. But, the bootstrap sensitivity analysis showed that, for most

conclusions drawn, the addiction behaviour of the agents was relatively insensitive to the particular DCE parameterization. Further work will need to be undertaken to examine why this is the case, but we expect that it is due to the primary outputs from the model being addiction and consumption related, rather than focusing on purchasing behaviour.

We also drew on geographic data for store locations, tobacco price, and non-tobacco demographics. We did not, however, have data to parameterize the social networks, adding a significant model assumption, though this data is difficult to obtain and rarely provides a complete map of a social network. Conducting the bootstrap sensitivity analysis does some to mitigate these limitations, but having access to more data would allow a more in-depth validation process, thereby increasing the confidence in model predictions.

A final word should be made about the nature of the inter-disciplinary team involved in this model. It involved close collaboration between several systems modellers and a tobacco expert. In order to facilitate the most useful discussion, and because the team was relatively small, we set up weekly meetings during which much of the modelling actually took place. This allowed real-time interaction using screen-sharing between the tobacco expert and the modellers, which allowed the modellers to quickly learn the tobacco research goals and relevant domain knowledge, and the tobacco expert to quickly see the types of mechanism that needed to be specified. This interaction led each group to a more complete understanding of the project as a whole. It also reduced model development time, since less model code needed to be abandoned when it was belatedly found to contradict with established domain knowledge. Important to this work was the use of a visual modelling package, improving the transparency of model logic to non-modellers.

Summing up, this work showed a number of valuable insights, but they can be grouped into two categories. Firstly, simple parameter count or mechanistic complexity does not guarantee an increase in predictive accuracy. Reasoning about the implications of actual causal relationships within addiction modules is invaluable to not only improving predictive accuracy, but also in the development of insight. Secondly, there is reason to believe that sub-populations are exposed to different dominant behavioural forces. Without accounting for this, it can confound results. There is great value in being able to model and study them at the sub-population level.

5.6 Summary

In this research, we constructed a dynamic model, the TPMS, with two key lines of research. First, we integrated two powerful, yet different, methods used to understand human preferences in the light of taxation interventions, Discrete Choice Theory and dynamic models, using the strengths of each to counteract individual limitations. Second, we used a modular framework to compare and contrast multiple formulations of addiction dynamics.

Along the first, Discrete Choice Theory is used across many disciplines, including public health and health policy studies. It uses stated preference data, which benefits from being more available than revealed

preference data. Discrete Choice Theory is able to statistically analyze hypothetical situations, controlling for a number of factors. A key benefit is its ability to examine counter-factuals. However, it is not easy to identify when hypothetical choices might diverge from actual choices. As well, DCE does not explicitly account for time, therefore it does not provide longer-term predictions, which are central to public health questions of behaviour change.

Dynamic models (including ABMs) are seeing increasing use in studies of public health. They are also able to investigate counter-factuals. Furthermore, they are able to incorporate data from diverse sources in a principled way. Unlike DCE, dynamic models can focus on the causative mechanism rather than simply statistical associations. This also requires that time is included in the model. However, the power of expression in such models requires the specification of many factors for which there might be no data or theory.

Bringing each of these different methods together allows us to benefit from strengths of each. The dynamic model places agents within a broader context, allowing us to reason through the implications of DCE-driven decisions beyond the narrow time horizon of DCE theory. The DCE provides independently validated purchasing behaviour, allowing us to focus our attention on the latent mechanisms of addiction. Indeed, since not all of the DCE parameters were statistically significant, we initially observed illogical purchasing behaviour: agents often maintained unreasonably large stockpiles of tobacco products. After including a few assumptions in the model to control this behaviour, we were able to validate the DCE element.

Along the second line, the modular framework was implemented in the TPMS to focus thinking around addiction dynamics. We examined how the logic of multiple addiction modules affected agent behaviour. We found that no single module proved to be more advantageous in all situations. The IDM module performed quite well, even though it was the simplest. Clearly, significant simplifications can still confer value. The BCT module shows the advantage of precise and clear logic characterizing behavioural mechanisms. It was generally able to match the consumption patterns across all sub-populations. However, it lacked an articulate theory for initiation and cessation, which reduced the fit of its predictions. Generally, FlexP was the most predictively accurate, but even though FlexP has the highest parameter count, it has a very similar structure to the poorest performing module, the PCT. This demonstrates two further insights: module complexity does not necessarily improve accuracy, and exploring variations in module design—a process made very simple using our modularization approach—is very valuable. We believe that such modularity can be easily implemented in other models of human behaviour and can aid in facilitating more rigorous and principled studies.

Beyond the mechanism of the addiction modules, we also argued that sub-population characteristics could drive the diversity of predictions across module types. Different sub-populations face different dominant forces on decisions, better matched by a diversity of proposed addiction mechanisms. This is supported by the data showing different sub-population responses to the same tax—some increased their consumption, some decreased, and some remained stable. While PCT produced the overall poorest fit predictions, it was the best module for two sub-populations. The modular approach allows the addiction module to be parameterized

to the level of the sub-population (or even the individual agent, should the data for calibration be present). Using the modular approach effectively makes the module a system parameter, which can be varied as the needs dictate.

Over and above these benefits of the modular approach, it provides two further strengths. Firstly, it provides a very general method of building hybrid models. While the model framework was agent-based, the modules were a mix of agent-based or SD. This allows models to be quickly built that benefit from the inherent benefits of multiple methods. Secondly, when maintain several addiction modules, it helps to reduces code errors. Without a modular construction, we would be required to create separate copies of the model for each addiction approach. Any change in one of the models would need to be propagated to each of the separate models. Developing with modules allows us to maintain a single model, which reduces copying code, and therefore, copying code errors.

A difficulty when working with ABMs, especially a highly modular one that requires significant calibration, is the high cost in computational time and resources. With AnyLogic's calibration scheme, which uses the proprietary global optimization algorithm called OptQuest, calibration experiments were not easily sent to clusters for distributed execution; they were instead run on single machines. We used a combination of virtual machines and lab desktops to accomplish all the calibrations required for this work.

CHAPTER 6

CONCLUSIONS

As described in Chapter 2, we identified a gap between theory-driven and data-driven dynamic modelling approaches when studying complex systems rich with latent characteristics, such as public health systems. We proposed a two-fold solution to filling this gap, benefiting from the strengths of both data and theory. The first was to develop methods that are flexible to the many and growing sources of data, including the personal micro-behaviour data from smartphones, so as to improve the possible points of calibration for theory-laden models. We proposed using a sequential Monte Carlo method called the particle filter. This work failed to generate the intended improvements in prediction, though learning still took place. The second was the use of models that are scalable in complexity, encouraging an iterative design methodology. Here we developed and implemented a modular design pattern, the results of which were much more successful.

Our work with the particle filter (PF), a sequential Monte Carlo algorithm, was the first known case using agent-based particle models, instead of the more common SD particle model [78, 79]. The results from our initial experimentation showed that the ABM-PF actually worsened ensemble predictions. We hypothesized that this is in large part due to the increased dimensionality of the ABM particle models, requiring a corresponding increase in the number of particles to adequately sample the state-space, and therefore, the computational resources required. This hypothesis was supported by preliminary experimentation. It remains an interesting research question as to what conditions might better indicate an ABM-PF, but for those who are focused on developing predictive models in the short term, an SD-PF would likely return more fruit.

A further learning was that the current agent-based modelling frameworks, including AnyLogic which was used here, do not provide mechanism or flexibility for integrating ABMs with the PF: our solution required implementing the particle ABM in native Java code thereby limiting access to the many powerful features of an ABM program. This inflexibility in modern frameworks significantly limits the ability to explore the usefulness of this type of dynamic simulation in complex public health systems.

The second piece of the solution outlined in Chapter 2 is the use of models that are scalable in design complexity, encouraging an iterative design methodology. To this end, we developed a modular architecture using design principles from software engineering, and implemented it in an ABM meant to study tobacco purchasing behaviour. We performed a calibration analysis on 4 addiction modules, each assuming different governing mechanisms. The results showed that different addiction modules fit better for different agent sub-populations, and we argued that this could be due to the variations in the addiction process between

the sub-populations. Using the modular approach, we were able to nimbly learn from experiments, and as a result, identified a key limitation in the mechanism of one module. Addressing that limitation resulted in the best overall module fit. We also applied a previously-developed model centered specifically around tobacco addiction, and it performed better than the other models in predicting consumption patterns. This was made possible with the use of our modular approach, allowing us to avoid the failure modes that come from unnecessary reproduction of models and code. This provides a powerful reason to support experimenting explicitly with mechanisms of latent factors, such as human behaviour, a prospect made much simpler with our modular design pattern.

Beyond our work already described in Chapters 3 and 5, other researchers are advancing this space. Dong [28] develops a method to use a Hierarchical Bayesian approach to combine stated and revealed preference data together allowing individual-level inferences. This is seen as “helpful to understand the underlying decision process” [28] taken by individuals. It is not about a dynamic model, but the technique could inform the types of theories implemented in them. The use of other numerical techniques as well could prove very useful in this iterative approach. Beheshti [11] develops a method to use Markov-chain Monte Carlo sampling and ABM modelling together to increase data-realism. While it does not directly address the behavioural model, it does bring new insight into quantitative methods that better unite sampling techniques and modelling. And Andrieu et al. [2] has developed the Particle MCMC algorithm, which can be leveraged to sample jointly from both latent state and parameter values in the study of high-dimensional and non-linear systems, such as those studied here. Clearly, more work needs to be conducted in learning how dynamic simulation methods can integrate and advance the study of complex adaptive systems using large and diverse sources of data.

This work also resulted in learning that was outside the original thesis statement, but related to supplementing standard survey-driven methods of understanding human behaviour. Firstly, to define agent purchasing behaviour within the tobacco ABM, we integrated within the agents the results of a discrete choice experiment (DCE) that we conducted for this work. DCEs are a standard tool of analysis used to understand possible changes to tobacco behaviours in the light of counter-factual tobacco taxes. However, it is limited in a number of ways, including that it assumes rational decision making, which is violated by non-rational behaviour such as addiction, and that it does not predict behaviour long after an intervention. Within our ABM, however, we moderated agent purchasing behaviour with a greater agent context, such as explicitly captured social networks and proposed addiction dynamics. It is conceivable that this could be used to counter-balance DCE limitations, though more work would need to be done to draw this out.

The second learning came about because the tobacco ABM required a means of describing agent addiction behaviour. A common method for validating behaviour models is using survey validation tools, though this does not explicitly describe the behaviour-generating process and thus struggle when used to understand counter-factual situations or capture population effects of behaviour adaptation. Though not a question we aimed to address in this work, we nonetheless demonstrated another potential method: using deep mathematical models of a narrower decision space, developed and calibrated at the individual level, within agents

of an ABM that capture a much broader set of decisions. The BCT addiction module was built from a behaviour model of addiction originally validated using animal experiments, where more in-depth neural analysis is possible. This model also invokes control theory to describe a complex adaptive process, a feature not present in social-cognitive theories such as the Theory of Planned Behaviour. This module was placed within a broader ABM that contained social networks and survey-validated tobacco purchasing. That it was the most accurate addiction module when predicting consumption patterns is a potential confirmation that this approach should be investigated further.

A final word about model validation is warranted. Model validation is a very important step in the model-building process, whose main goal is to allow researchers to build their confidence in model behaviour. Since all models approximate to varying degrees some underlying reality, validation measures similarly span a wide range of techniques, from qualitative comparison with stylized facts to cross-validation measures using multiple sources of time-series data. An important question one might ask, therefore, is, "When do we know that the model is good enough for our purposes?" An insight into this question from our tobacco ABM work is that the answer depends not only on the model behaviour, but on our understanding of the domain space. For example, our results from the tobacco ABM were initially surprising, that more complex and detailed addiction modules did not necessarily improve the fit to data. Using our knowledge of the causal structure of the ABM, we generated testable hypotheses, a goal of good science, relating to human addiction behaviour.

This work has potential impact on scientists across many differing fields. For example, social scientists might be able to supplement their behavioural theories with social context from ABMs; researchers studying the health implications of tobacco use might be able to develop and validate neurologically-inspired model of tobacco addiction; policy scientists could generate more insightful predictions of counter-factual interventions in the presence of regularly arriving public data; and computational modellers could improve the speed of development and quality of their models. In all of these cases, improvement is possible by focusing on both data and design.

6.1 Future Work

All of this work opens up a wide variety of valuable future research opportunities. Using the modular addiction model, we can very easily perform experimentation on addiction modules suited for specific agents or sub-populations. What mechanisms might be the key driving forces to addiction within certain demographic groups? What data would be useful in observing those forces? How do multiple sectors of a heterogeneous population respond to specific counter-factual tax situations? If we varied the spatial characteristics, how does the location of stores change the purchasing and consumption behaviour, especially given different addiction modules?

We can further test the DCE integration with dynamic models. An initial step would be to validate the TPMS model, not only for consumption behaviour, but purchasing behaviour as well. Expanding in this

direction, can the purchasing behaviour be validated using only a subset of the original DCE dataset? This might allow the dynamic model to supplement a poorly fit DCE, thereby improving its predictive capacity beyond what is currently possible with traditional DCE methods. And how much can the accuracy of the model be improved when also calibrating for other characteristics, such as quit attempts?

One key output of a DCE is short term behaviour post-intervention, in this case, a tax intervention. In this work, we did not study the agent behaviour immediately post-tax. This is a natural next step of the work, and could be used to compare and contrast DCE-only and ABM-DCE hybrid predictions of price elasticities. We also do have data from the MATS on quit attempts within the population. This could be studied in future versions of the TPMS model.

On the particle filtering side, we could explore the behaviour of the ABM particle filter with many more particles. More work would also be needed to understand clearly what conditions, if any, imply a ABM particle filter would perform better than an SD particle filter. Is there a threshold of population heterogeneity in the modelled population, for example, where ABM particles are more advantageous? Additionally, could a hybrid particle filter lend improvements, with some ABM and some SD particles mixed into the same ensemble? Going beyond the particle filter itself, the particle MCMC algorithm could be explored as well.

REFERENCES

- [1] AJZEN, I. The theory of planned behavior. *Organizational Behavior And Human Decision Processes* 50, 2 (1991), 179–211.
- [2] ANDRIEU, C., DOUCET, A., AND HOLENSTEIN, R. Particle markov chain monte carlo methods. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 72, 3 (2010), 269–342.
- [3] ANDRIGHETTO, G., CRANEFIELD, S., CONTE, R., PURVIS, M., PURVIS, M., SAVARIMUTHU, B. T. R., AND VILLATORO, D. (social) norms and agent-based simulation. In *Agreement Technologies*. Springer, 2013, pp. 181–189.
- [4] ARULAMPALAM, M. S., MASKELL, S., GORDON, N., AND CLAPP, T. A tutorial on particle filters for online nonlinear/non-GaussianBayesian tracking. *Signal Processing, IEEE Transactions on* 50, 2 (2002), 174–188.
- [5] AXTELL, R., DURLAUF, S., EPSTEIN, J. M., HAMMOND, R., KLEMENS, B., PARKER, J., SONG, Z., VALENTE, T., AND YOUNG, H. P. Social influences and smoking behavior final report to the american legacy foundation. Tech. rep., Brookings Institute, 2006.
- [6] BACHAND, A. M., AND SULSKY, S. I. A dynamic population model for estimating all-cause mortality due to lifetime exposure history. *Regulatory Toxicology And Pharmacology* 67, 2 (2013), 246–251.
- [7] BALKE, T., AND GILBERT, N. How do agents make decisions? a survey. *Journal Of Artificial Societies And Social Simulation* 17, 4 (2014), 13. Available at <http://jasss.soc.surrey.ac.uk/17/4/13.html>.
- [8] BARABÁSI, A.-L., AND ALBERT, R. Emergence of scaling in random networks. *Science* 286, 5439 (1999), 509–512.
- [9] BECKER, G. S. Habits, addictions, and traditions. *Kyklos* 45, 3 (1992), 327–345.
- [10] BECKMAN, R., KUHLMAN, C., MARATHE, A., NSOESIE, E., AND SWARUP, S. Modeling the spread of smoking in adolescent social networks. In *Proceedings Of The Fall Research Conference Of The Association For Public Policy Analysis And Management* (2011), Citeseer.
- [11] BEHESHTI, R. Modeling social norms in real-world agent-based simulations. *Ai Matters* 2, 1 (2015), 9–11.
- [12] BEHESHTI, R., ALI, A. M., AND SUKTHANKAR, G. R. Cognitive social learners: An architecture for modeling normative behavior. In *AAAI Conference on Artificial Intelligence* (2015), pp. 2017–2023.
- [13] BEHESHTI, R., AND SUKTHANKAR, G. A normative agent-based model for predicting smoking cessation trends. In *Proceedings Of The 2014 International Conference On Autonomous Agents And Multi-agent Systems* (2014), International Foundation for Autonomous Agents and Multiagent Systems, pp. 557–564.
- [14] BEST, A., CLARK, P. I., LEISCHOW, S. J., AND TROCHIM, W. M. K., Eds. *Greater Than The Sum: Systems Thinking In Tobacco Control*. Tobacco Control Monograph No. 18. U.S. Department of Health and Human Services, National Institutes of Health, National Cancer Institute., April 2007.
- [15] BOBASHEV, G., HOLLOWAY, J., SOLANO, E., AND GUTKIN, B. A control theory model of smoking. *Methods report (RTI Press) 2017* (2017).

- [16] BOBASHEV, G., SOLANO, E., AND HOFFER, L. An agent-based model to support measuring drug choice and switch between drug types in rural populations. In *The Sixth International Conference On Advances In System Simulation* (2014), Advances in System Simulation.
- [17] BRAILSFORD, S., AND SCHMIDT, B. Towards incorporating human behaviour in models of health care systems: An approach using discrete event simulation. *European Journal Of Operational Research* 150, 1 (2003), 19–31.
- [18] BRAILSFORD, S. C., SYKES, J., AND HARPER, P. R. Incorporating human behavior in healthcare simulation models. In *Proceedings Of The 38th Conference On Winter Simulation* (2006), Winter Simulation Conference, pp. 466–472.
- [19] CAREY, G., MALBON, E., CAREY, N., JOYCE, A., CRAMMOND, B., AND CAREY, A. Systems science and systems thinking for public health: a systematic review of the field. *BMJ Open* 5, 12 (2015), e009002.
- [20] CHALOUPKA, F. J., YUREKLI, A., AND FONG, G. T. Tobacco taxes as a tobacco control strategy. *Tobacco Control* 21, 2 (2012), 172–180.
- [21] CHOI, K., AND BOYLE, R. G. Changes in cigarette expenditure minimising strategies before and after a cigarette tax increase. *Tobacco Control* (2017), tobaccocontrol-2016-053415.
- [22] COLDITZ, G. A., BREWER, T. F., BERKEY, C. S., WILSON, M. E., BURDICK, E., FINEBERG, H. V., AND MOSTELLER, F. Efficacy of BCG vaccine in the prevention of tuberculosis. Meta-analysis of the published literature. *JAMA* 271, 9 (1994), 698–702.
- [23] COMMITTEE ON THE PUBLIC HEALTH IMPLICATIONS OF RAISING THE MINIMUM AGE FOR PURCHASING TOBACCO PRODUCTS, BOARD ON POPULATION HEALTH AND PUBLIC HEALTH PRACTICE, AND INSTITUTE OF MEDICINE. The Effect On Tobacco Use Of Raising The Minimum Age Of Legal Access To Tobacco Products. In *Public Health Implications of Raising the Minimum Age of Legal Access to Tobacco Products*, R. J. Bonnie, K. Stratton, and L. Y. Kwan, Eds. National Academies Press (US), 2015.
- [24] CONNER, M., AND NORMAN, P., Eds. *Predicting Health Behaviour*. McGraw-Hill Education (UK), 2005.
- [25] CROISSANT, Y. *mlogit: multinomial logit model*, 2013. R package version 0.2-4.
- [26] DE BEKKER-GROB, E. W., RYAN, M., AND GERARD, K. Discrete choice experiments in health economics: A review of the literature. *Health Economics* 21, 2 (2012), 145–172.
- [27] DECHESNE, F., DI TOSTO, G., DIGNUM, V., AND DIGNUM, F. No smoking here: Values, norms and culture in multi-agent systems. *Artificial Intelligence And Law* 21, 1 (2013), 79–107.
- [28] DONG, X. Applications of hierarchical bayesian method in analyzing individual level behavior. *ProQuest Dissertations and Theses* (2006), 124.
- [29] ELLIOTT, M. A., AND AINSWORTH, K. Predicting university undergraduates’ binge-drinking behavior: A comparative test of the one-and two-component theories of planned behavior. *Addictive Behaviors* 37, 1 (2012), 92–101.
- [30] ERDOS, P., AND RENYI, A. On the evolution of random graphs. *Publications of the Mathematical Institute of the Hungarian Academy of Sciences* 5 (1960), 17–61.
- [31] FARIAS, G. P., DIMURO, G. P., AND DA ROCHA COSTA, A. C. BDI agents with fuzzy perception for simulating decision making in environments with imperfect information. In *Mallow* (2010).
- [32] FEOLA, G., AND BINDER, C. R. Towards an improved understanding of farmers’ behaviour: The integrative agent-centred (iac) framework. *Ecological Economics* 69, 12 (2010), 2323–2333.

- [33] FLEISCHER, N. L., THRASHER, J. F., REYNALES-SHIGEMATSU, L. M., CUMMINGS, K. M., MEZA, R., ZHANG, Y., AND LEVY, D. T. Mexico simsmoke: how changes in tobacco control policies would impact smoking prevalence and smoking attributable deaths in mexico. *Global Public Health* (2016), 1–16.
- [34] FOWLER, M., AND HIGSMITH, J. The agile manifesto. *Software Development* 9, 8 (2001), 28–35.
- [35] GRIMM, V., BERGER, U., BASTIANSEN, F., ELIASSEN, S., GINOT, V., GISKE, J., GOSS-CUSTARD, J., GRAND, T., HEINZ, S. K., HUSE, G., ET AL. A standard protocol for describing individual-based and agent-based models. *Ecological Modelling* 198, 1 (2006), 115–126.
- [36] HAMMOND, R. A., ORNSTEIN, J. T., FELLOWS, L. K., DUBÉ, L., LEVITAN, R., AND DAGHER, A. A model of food reward learning with dynamic reward exposure. *Frontiers In Computational Neuroscience* 6 (2012), 82.
- [37] HANSETH, O., AND LYYTINEN, K. Design theory for dynamic complexity in information infrastructures: The case of building internet. *Enacting Research Methods in Information Systems: Volume 3* 25, 1 (2016), 104–142.
- [38] HIANCE, D., DOOGAN, N., WARREN, K., HAMILTON, I. M., AND LEWIS, M. An agent-based model of lifetime attendance and self-help program growth. *Journal Of Social Work Practice In The Addictions* 12, 2 (2012), 121–142.
- [39] HODGSON, G. M. Choice, habit and evolution. *Journal Of Evolutionary Economics* 20, 1 (2010), 1–18.
- [40] HOFFER, L. Unreal models of real behavior: The agent-based modeling experience. *Practicing Anthropology* 35, 1 (2012), 19–23.
- [41] HOFFER, L., BOBASHEV, G., AND MORRIS, R. J. Simulating patterns of heroin addiction within the social context of a local heroin market. In *Computational Neuroscience Of Drug Addiction*. Springer, 2012, pp. 313–331.
- [42] HOHPE, G., WIRFS-BROCK, R., YODER, J. W., AND ZIMMERMANN, O. Twenty years of patterns’ impact. *IEEE Software* 30, 6 (2013).
- [43] HOLLAND, J. H. Complex Adaptive Systems. *Daedalus* 121, 1 (1992), 17–30.
- [44] HOLLANDER, C. D., AND WU, A. S. The current state of normative agent-based systems. *Journal Of Artificial Societies And Social Simulation* 14, 2 (2011), 6.
- [45] HOMER, J., WILE, K., YARNOFF, B., TROGDON, J. G., HIRSCH, G., COOPER, L., SOLER, R., AND ORENSTEIN, D. Peer reviewed: Using simulation to compare established and emerging interventions to reduce cardiovascular disease risk in the United States. *Preventing Chronic Disease* 11 (2014).
- [46] HOMER, J. B., AND HIRSCH, G. B. System dynamics modeling for public health: Background and opportunities. *American Journal Of Public Health* 96, 3 (2006), 452–458.
- [47] JENSEN, T., AND CHAPPIN, É. J. Towards an agent-based model on co-diffusion of technology and behavior: A review. In *Ecms* (2014), pp. 782–788.
- [48] KARANFIL, O., MOORE, T., FINLEY, P., BROWN, T., ZAGONEL, A., AND GLASS, R. A multi-scale paradigm to design policy options for obesity prevention: Exploring the integration of individual-based modeling and system dynamics. *Sandia National Laboratories 202011* (2011).
- [49] KAUFMANN, P., STAGL, S., AND FRANKS, D. W. Simulating the diffusion of organic farming practices in two new eu member states. *Ecological Economics* 68, 10 (2009), 2580–2593.
- [50] KILLEEN, P. R. Markov model of smoking cessation. *Proceedings of the National Academy of Sciences* 108, Supplement 3 (2011), 15549–15556.

- [51] KLABUNDE, A., AND WILLEKENS, F. Decision-making in agent-based models of migration: State of the art and challenges. *European Journal Of Population* 32, 1 (2016), 73–97.
- [52] KNIVETON, D., SMITH, C., AND WOOD, S. Agent-based model simulations of future changes in migration flows for burkina faso. *Global Environmental Change* 21 (2011), S34–S40.
- [53] KNIVETON, D. R., SMITH, C. D., AND BLACK, R. Emerging migration flows in a changing climate in dryland africa. *Nature Climate Change* 2, 6 (2012), 444–447.
- [54] KOSTER, A., SCHORLEMMER, M., AND SABATER-MIR, J. Opening the black box of trust: Reasoning about trust models in a BDI agent. *Journal Of Logic And Computation* (2012), exs003.
- [55] KREYSZIG, E. *Advanced Engineering Mathematics*, 9th ed. Wiley, 2006.
- [56] LAKON, C. M., HIPPI, J. R., WANG, C., BUTTS, C. T., AND JOSE, R. Simulating dynamic network models and adolescent smoking: The impact of varying peer influence and peer selection. *American Journal Of Public Health* 105, 12 (2015), 2438–2448.
- [57] LEVIN, S., XEPAPADEAS, T., CRÉPIN, A.-S., NORBERG, J., DE ZEEUW, A., FOLKE, C., HUGHES, T., ARROW, K., BARRETT, S., DAILY, G., EHRLICH, P., KAUTSKY, N., MÄLER, K.-G., POLASKY, S., TROELL, M., VINCENT, J. R., AND WALKER, B. Social-ecological systems as complex adaptive systems: modeling and policy implications. *Environment and Development Economics* 18, 02 (2013), 111–132.
- [58] LEVY, D., CUMMINGS, K., AND HYLAND, A. A simulation of the effects of youth initiation policies on overall cigarette use. *American Journal Of Public Health* 90, 8 (2000), 1311.
- [59] LEVY, D., FERGUS, C., RUDOV, L., MCCORMICK-RICKET, I., AND CARTON, T. Tobacco policies in louisiana: Recommendations for future tobacco control investment from simsmoke, a policy simulation model. *Prevention Science* 17, 2 (2016), 199–207.
- [60] LEVY, D., WEN, C., CHEN, T., AND OBLAK, M. Increasing taxes to reduce smoking prevalence and smoking attributable mortality in taiwan: Results from a tobacco policy simulation model. *Tobacco Control* 14, suppl 1 (2005), i45–i50.
- [61] LEVY, D. T., FRIEND, K., HOLDER, H., AND CARMONA, M. Effect of policies directed at youth access to smoking: Results from the simsmoke computer simulation model. *Tobacco Control* 10, 2 (2001), 108–116.
- [62] LIPSITCH, M., COHEN, T., COOPER, B., ROBINS, J. M., MA, S., JAMES, L., GOPALAKRISHNA, G., CHEW, S. K., TAN, C. C., SAMORE, M. H., FISMAN, D., AND MURRAY, M. Transmission dynamics and control of severe acute respiratory syndrome. *Science* 300, 5627 (2003), 1966–70.
- [63] LUONG, B. V., THANGARAJAH, J., ZAMBETTA, F., AND HASAN, M. A BDI game master agent for computer role-playing games. In *Proceedings Of The 2013 International Conference On Autonomous Agents And Multi-agent Systems* (2013), International Foundation for Autonomous Agents and Multi-agent Systems, pp. 1187–1188.
- [64] MABRY, P. L., MARCUS, S. E., CLARK, P. I., LEISCHOW, S. J., AND MÉNDEZ, D. Systems science: A revolution in public health policy research. *American Journal of Public Health* 100, 7 (2010), 1161–1163.
- [65] MALLESON, N., HEPPENSTALL, A., AND SEE, L. Crime reduction through simulation: An agent-based model of burglary. *Computers, Environment And Urban Systems* 34, 3 (2010), 236–250.
- [66] MANGHAM, L. J., HANSON, K., AND MCPAKE, B. How to do (or not to do) designing a discrete choice experiment for application in a low-income country. *Health Policy and Planning* 24, 2 (2009), 151–158.

- [67] MASLENNIKOVA, G. Y., OGANOV, R. G., BOYTSOV, S. A., ROSS, H., HUANG, A.-T., NEAR, A., KOTOV, A., BEREZHNOVA, I., AND LEVY, D. T. Russia simsmoke: the long-term effects of tobacco control policies on smoking prevalence and smoking-attributable deaths in russia. *Tobacco control* 23, 6 (2014), 484–490.
- [68] MCFADDEN, D. Conditional logit analysis of qualitative choice behaviour. In *Frontiers in Econometrics*, P. Zarembka, Ed. Academic Press New York, New York, NY, USA, 1973, pp. 105–142.
- [69] MEADOWS, M., AND CLIFF, D. Reexamining the relative agreement model of opinion dynamics. *Journal Of Artificial Societies And Social Simulation* 15, 4 (2012), 4.
- [70] MEJIA, A. B., LING, P. M., AND GLANTZ, S. A. Quantifying the effects of promoting smokeless tobacco as a harm reduction strategy in the usa. *Tobacco Control* 19, 4 (2010), 297–305.
- [71] MÉNDEZ, D. Menthol cigarettes and public health: Review of the scientific evidence and recommendations. Tech. rep., Tobacco Products Scientific Advisory Committee, 2011.
- [72] MÉNDEZ, D., WARNER, K. E., AND COURANT, P. N. Has smoking cessation ceased? expected trends in the prevalence of smoking in the United States. *American Journal Of Epidemiology* 148, 3 (1998), 249–258.
- [73] MILLER, J. H., AND PAGE, S. E. *Complex Adaptive Systems: An Introduction to Computational Models of Social Life*. Princeton University Press, 2007.
- [74] MÜLLER, B., BOHN, F., DRESSLER, G., GROENEVELD, J., KLASSERT, C., MARTIN, R., SCHLÜTER, M., SCHULZE, J., WEISE, H., AND SCHWARZ, N. Describing human decisions in agent-based models—odd+ d, an extension of the odd protocol. *Environmental Modelling & Software* 48 (2013), 37–48.
- [75] NIANOGO, R. A., AND ARAH, O. A. Agent-based modeling of noncommunicable diseases: A systematic review. *American Journal Of Public Health* 105, 3 (2015), e20–e31.
- [76] NUTTALL, W., ZHANG, T., HAMILTON, D., AND ROQUES, F. Sociophysics simulations of technology adoption and consumer behavior. In *Second International Symposium On Engineering Systems, MIT* (2009).
- [77] OKAYA, M., AND TAKAHASHI, T. BDI agent model based evacuation simulation. In *The 10th International Conference On Autonomous Agents And Multiagent Systems-volume 3* (2011), International Foundation for Autonomous Agents and Multiagent Systems, pp. 1297–1298.
- [78] ONG, J. B. S., CHEN, M. I.-C., COOK, A. R., LEE, H. C., LEE, V. J., LIN, R. T. P., TAMBYAH, P. A., AND GOH, L. G. Real-time epidemic monitoring and forecasting of H1N1-2009 using influenza-like illness from general practice and family doctor clinics in singapore. *PLoS ONE* 5, 4 (2010).
- [79] OSGOOD, N., AND LIU, J. Towards closed loop modeling: evaluating the prospects for creating recurrently regrounded aggregate simulation models using particle filtering. In *Proceedings of the 2014 Winter Simulation Conference* (Piscataway, New Jersey, 2014), A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, Eds., Institute of Electrical and Electronics Engineers, Inc., pp. 829–941.
- [80] PADGETT, D. K., STANHOPE, V., HENWOOD, B. F., AND STEFANCIC, A. Substance use outcomes among homeless clients with serious mental illness: Comparing housing first with treatment first programs. *Community Mental Health Journal* 47, 2 (2011), 227–232.
- [81] POWERS, W. T. *Behavior: The Control of Perception*, 2nd ed. Benchmark Publishing New Canaan, 2005.
- [82] PUHL, R. M., AND HEUER, C. A. Obesity stigma: Important considerations for public health. *American Journal of Public Health* 100, 6 (2010), 1019–1028.

- [83] QIAN, W., OSGOOD, N. D., AND STANLEY, K. G. Integrating epidemiological modeling and surveillance data feeds : A kalman filter based approach. In *International Conference on Social Computing, Behavioral-Cultural Modeling and Prediction* (2014), Springer, pp. 145–152.
- [84] RAO, A. S., AND GEORGEFF, M. P. BDI agents: From theory to practice. In *ICMAS* (1995), vol. 95, pp. 312–319.
- [85] RIGOTTI, N. A., AND WALLACE, R. B. Using agent-based models to address 'wicked problems' like tobacco use: A report from the institute of medicine. *Annals Of Internal Medicine* 163, 6 (2015), 469–471.
- [86] SAVARIMUTHU, B. T. R., AND CRANFIELD, S. Norm creation, spreading and emergence: A survey of simulation models of norms in multi-agent systems. *Multiagent And Grid Systems* 7, 1 (2011), 21–54.
- [87] SCHAEFER, D. R., ADAMS, J., AND HAAS, S. A. Social networks and smoking exploring the effects of peer influence and smoker popularity through simulations. *Health Education & Behavior* 40, 1 suppl (2013), 24S–32S.
- [88] SCHELLING, T. C. Dynamic models of segregation. *The Journal of Mathematical Sociology* 1, 2 (1971), 143–186.
- [89] SCHMIDT, B. The modelling of human behaviour: The PECS reference models. In *Proceedings of the 14th European Simulation Symposium* (2002), SCS-Europe BVBA.
- [90] SHIMODA, T. An interactive software-agent smoking cessation program. In *Proceedings Of The 36th Annual Hawaii International Conference On System Sciences* (2003), IEEE, pp. 9–pp.
- [91] SUTTON, S., McVEY, D., AND GLANZ, A. A comparative test of the theory of reasoned action and the theory of planned behavior in the prediction of condom use intentions in a national sample of english young people. *Health Psychology* 18, 1 (1999), 72.
- [92] SWINERD, C., AND MCNAUGHT, K. R. Design classes for hybrid simulations involving agent-based and system dynamics models. *Simulation Modelling Practice And Theory* 25 (2012), 118–133.
- [93] TAILLANDIER, P., THEROND, O., AND GAUDOU, B. A new BDI agent architecture based on the belief theory. application to the modelling of cropping plan decision-making. In *International Environmental Modelling And Software Society, (iEMSs)* (2012).
- [94] TEHRANI, N. G., AGHDASI, M., LOUCA, L. S., CHRYSANTHOU, Y., OPLATKOVA, Z., AND ALBEGAIN, K. Proposing a framework for building human-based agents for discrete event simulation: a case study in screening programs. In *Proceedings of the 22nd European Conference on Modelling and Simulation* (2008), Citeseer, pp. 111–117.
- [95] TENGs, T. O., OSGOOD, N. D., AND LIN, T. H. Public health impact of changes in smoking behavior: Results from the tobacco policy model. *Medical Care* 39, 10 (2001), 1131–1141.
- [96] TIAN, Y., AND OSGOOD, N. Comparison between Individual-based and Aggregate Models in the context of Tuberculosis Transmission. In *Proceedings of the 29th International Conference of the System Dynamics Society* (2011), pp. 1–29.
- [97] TOBIAS, M. I., CAVANA, R. Y., AND BLOOMFIELD, A. Application of a system dynamics model to inform investment in smoking cessation services in New Zealand. *American Journal Of Public Health* 100, 7 (2010), 1274–1281.
- [98] TRAIN, K. E. *Discrete choice methods with simulation*. Cambridge university press, 2009.
- [99] TULLY, S., COJOCARU, M., AND BAUCH, C. T. Coevolution of risk perception, sexual behaviour, and hiv transmission in an agent-based model. *Journal Of Theoretical Biology* 337 (2013), 125–132.

- [100] URBAN, C., AND SCHMIDT, B. PECS-agent-based modelling of human behaviour. In *Proceedings of the AAAI Fall Symposium: Emotional and Intelligent—the Tangled Knot of Social Cognition*, North Falmouth, Ma. (2001).
- [101] U.S. NATIONAL CANCER INSTITUTE AND WORLD HEALTH ORGANIZATION. The Economics of Tobacco and Tobacco Control. National Cancer Institute Tobacco Control Monograph 21. *Bethesda, MD: U.S. Department of Health and Human Services, National Institutes of Health, National Cancer Institute; and Geneva, CH: World Health Organization NIH Public*, 1 (2016), 6242.
- [102] VERZI, S. J., BRODSKY, N., BROWN, T., APELBERG, B., AND ROSTRON, B. An agent-based approach for modeling population behavior and health with application to tobacco use. *Sandia National Laboratories* (2012).
- [103] WARNER, K. E., AND MÉNDEZ, D. Accuracy and importance of projections from a dynamic simulation model of smoking prevalence in the United States. *American Journal Of Public Health* 102, 11 (2012), 2045–2048.
- [104] ZAGONEL, A. A., MOJTAHEDZADEH, M., RICHARDSON, G. P., BRODSKY, N. S., BROWN, T. J., CONRAD, S. H., AND GLASS, R. J. Developing a theory of the societal lifecycle of cigarette smoking: Explaining and anticipating trends using information feedback. <https://pdfs.semanticscholar.org/518f/52a956b89f4d44e53a1f315569e7608a3422.pdf>, 2010. Accessed: 2017-01-25.
- [105] ZHANG, J., TONG, L., LAMBERSON, P., DURAZO-ARVIZU, R., LUKE, A., AND SHOHAM, D. Leveraging social influence to address overweight and obesity using agent-based models: The role of adolescent social networks. *Social Science & Medicine* 125 (2015), 203–213.
- [106] ZHANG, T., AND NUTTALL, W. J. An agent based simulation of smart metering technology adoption. *IDEAS Working Paper Series from RePEc* (2007).