



UvA-DARE (Digital Academic Repository)

Social dynamics of substance use through minds and models

van den Ende, M.

Publication date
2025

[Link to publication](#)

Citation for published version (APA):

van den Ende, M. (2025). *Social dynamics of substance use through minds and models*. [Thesis, fully internal, Universiteit van Amsterdam].

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Summary and General Discussion

The crises in replication and theory signal the need for new methods to advance psychological understanding. The replication crisis is being addressed through efforts such as open science initiatives, meta-analytic techniques and efforts to enhance the robustness of experimental designs. The theory crisis however, remains a significant challenge, suggesting that we need new approaches to develop and refine psychological theories. In this dissertation, I have used a formal modelling approach to address these challenges and present my efforts to advance our understanding of the psychology of substance use and its interaction with the social environment.

Chapter 2

This thesis is the result of a journey in which I have become acquainted with the field of psychology, the state of research on drug abuse, and its overlap and interaction with the social sciences. The chapters described in this thesis follow this process, which began with a review of the literature on formal models of drug abuse. The aim was to find a model that could serve as a starting point for a formalisation of drug abuse behaviour that would be able to capture the social aspects of it – and during this research I found that while it was theoretically well recognised that there is a high degree of interaction between one’s drug abuse behaviour and one’s social environment, theories and models of this behaviour were highly disjointed. Many descriptions of theories did not directly mention how different social environments would affect individuals’ drug use behaviour. In addition, formalised theories were scarce, and those that were available were inadequate in capturing the social aspects of drug use. In contrast, the social science complexity approach faced its own unique challenges. This methodology has gained significant traction in recent years, following the success of the complexity approach in fields such as epidemiology, economics, and ecology. However, the available complex social models of substance use were very disconnected from contemporary psychological theories, almost as if there had been little dialogue with experts in the psychological field. This led to two key insights that formed the initial basis of my dissertation. First, there was a clear research gap in approaching substance use from a formalisation perspective that integrated contemporary psychological theory; and second, there was no simple starting point, no model of substance use and social behaviour that adequately represented both individual substance use behaviour and its interaction with one’s social environment.

Chapter 3

The second chapter describes a software package called DyNSimF, which has been developed to perform simulations of individual ‘complex’ agents in a social network. This was the result of a search for existing software that could meet the requirements of being able to solve differential equations for each individual, where each individual has information about some or all of the other individuals in the network, while also allowing each node to change its own connections, and doing this in an efficient way that could handle large networks. The fact that there was no software available that fully met these requirements for modelling complex networks is a sign of the relative infancy of the complexity field, and shows that while there is great interest in applying complexity science to the social sciences, the tools for doing so are still in their developmental stage. In response to this gap, we have developed the DyNSimF software package to meet these specific needs. It can be used to simulate drug abuse behaviour, but the software has been designed to be used for any kind of simulation of individual agents in a network. While DyNSimF is a clear improvement over other software packages, it could still use additional effort and maintenance to make it more adaptable and user-friendly. I am very keen to follow the level of interest in this software, which is directly linked to the wider interest in the network complexity approach, not only in the social and psychological sciences, but in many other fields such as ecology, economics and more.

Chapter 4

The epidemiological model of alcohol consumption presented in Chapter 4 stands as a fundamental chapter in my dissertation. It is based on the simple hypothesis that behaviours are transmitted through social ties. The model is thus a formalisation of alcohol consumption behaviour spreading on a social network. By adopting this approach, the model deliberately simplifies the psychological intricacies of individual actions, focusing instead on the effects of the social environment. The model takes advantage of the extensive knowledge available on the transmission of infectious diseases within networks; it is an adaptation of the established SIR-type model, rather than a novel framework constructed from scratch. This provides a familiar basis for those with a background in epidemiology, and provides rich resources for those from social or psychological disciplines to explore the system dynamics, without the need for extensive dynamics investigation within the paper itself.

However, the model also pushed the dataset to its limits. Although the data was exceptional in that it was one of the very few (if not the only) datasets with a large adult population that included information on the social networks of individuals where many of the other individuals were also sampled, it still had significant limitations. The long intervals between observations and the incompleteness of the social network severely limited the granularity of the model. A more complex model would have too many parameters and the dataset would not be large enough to estimate these parameters accurately. Yet despite its simplicity, the model successfully captured the general dynamics of social influences on drinking behaviour. The model showed that one’s social network has a significant impact on one’s drinking behaviour, and that treating both abstinence and heavy drinking as competing influences on moderate drinkers

accurately reflected the data. This facilitated the testing of hypotheses about large-scale social intervention strategies. Although some simulations were carried out, in retrospect more time and fewer constraints on the length of the paper would have allowed a deeper exploration of the possibilities and implications of the model.

A major limitation of this model is its assumption of Markovian dynamics: individuals in the same state behave identically, regardless of their history of drinking. This means that someone who has never drunk and feels no need to drink, will behave in the same way as someone who has just left a rehabilitation centre. As a result, the model fails to capture more nuanced dynamics of alcohol use over time. Integrating non-Markovian dynamics and taking into account the full history of an individual's alcohol use could provide a more accurate and comprehensive understanding of how drinking behaviour evolves and how interventions might affect long-term outcomes. This is the impetus for Chapter 6.

Chapter 5

One of the limitations I encountered in trying to formalise individual drinking behaviour was the lack of literature that both comprehensively described and empirically demonstrated certain phenomena. Take, for example, the widely accepted fact that individuals often consume more alcohol in social settings where others are drinking. Although this phenomenon has been clearly demonstrated and validated in adolescent groups, I have yet to find any relevant literature that definitively verifies or refutes its occurrence in older populations. However, when formalising the drinking behaviour of individuals, it is very important to know whether this phenomenon is unique to adolescents or whether it is a general phenomenon that can be applied to all age groups. Although there is a wealth of literature on alcohol use, many of these studies focus on specific populations or specific aspects of alcohol use. These more general phenomena are often overlooked, and the field of psychology could benefit from more studies focusing on these general phenomena, which would help to define constructs more clearly and thus advance psychological knowledge.

In the introduction to this dissertation, I describe this chapter as an *intermezzo* in the formalisation method. Unlike the rest of the dissertation, this empirical study involves an autoregression analysis, a technique widely used in psychology. Working with autoregressions has deepened my understanding and made me feel more like a psychological methodologist and psychometrician compared to the other more complexity-focused projects. By applying a panel GVAR (Graphical Vector Autoregression Model) to panel data that includes both individual internal variables and those inferred from one's social network, this project is at the forefront of innovative, empirical psychometric research. I believe that this methodology has significant potential for wider application to diverse datasets and complex phenomena. By generating graphical networks that articulate the nuances of dynamics both at the individual level and between individuals, and in particular by distinguishing temporal from cross-sectional dynamics, the use of such methods can be extremely valuable in generating hypotheses about the dynamics of complex systems. However, the set of assumptions made and the limitations of the approach and data need to be considered very carefully to avoid making false claims or over-interpreting the results; the presence of an edge is not the same as a causal relation-

ship. Therefore, the results of these models should be used for hypothesis generation and as a starting point for further research, rather than as a definitive answer to a research question. In future projects where I apply similar methods, I aim to delve deeper into the methodology to fully understand its capabilities, limitations, and assumptions, thus ensuring that I can utilise these methods to their full potential.

This study has achieved two things. First, it provides a proof of concept for the methodology, and I believe that applying these psychometric models to social environmental variables could help answer more questions that bridge sociology and psychology. In addition, the results strongly confirm that ‘older people’ – participants had an average age of 55 – are affected by both induction (the spread of behaviour, as also explored in Chapter 4) and homophily, which is the adjustment of one’s social network to favour individuals with similar drinking behaviours. This highlights the need for greater attention to the social environment in the study of individual drinking behaviour in the context of alcohol abuse interventions, particularly in relation to older populations where it appears to be under-researched. I hope that this study will act as a catalyst for further research into this critical issue, and that the methodology used here will be applied to other areas of psychology to further our understanding of complex systems.

Chapter 6

The development of the complex formal model of substance use described in the final chapter has been a four-year journey. Throughout this period, several early collaborations have highlighted the challenges in striking the right balance in modelling, as Smaldino (2017) says: the model needs to be as simple as possible, yet as complex as necessary. It was only after I acquired my own expertise in psychological theories of substance use that I was able to contribute effectively to the development of the mathematical description based on the ecological model of the interactions between budworms and predators within spruce tree forests. The result, I believe, is the most robust formal model of drug use currently available in the literature. In addition to addressing the main phenomena discussed in the paper, the model is capable of elegantly describing many other behaviours and phenomena with minimal adaptation or added complexity. I am excited to see future research applying this model framework to various specific cases and substances, and I sincerely hope that it will be considered in future modelling efforts. In particular, versions of the model that are validated and generate predictions and new hypotheses will have the potential to make interesting contributions to the field.

Rather than thinking in terms of constructs solely through words, this model has made thinking about substance use much clearer for me by thinking of these constructs in terms of the mathematics of the model. An example of this is tolerance and incentive sensitisation. Roughly speaking, tolerance represents the body’s adjustments to already counteract the effects of drugs, resulting in the need for a higher dose to feel a similar effect; incentive sensitisation is the brain’s additional ‘sensitivity’ to the positive effects of the drugs. Within the mathematics of the model, they both affect the appetitive motivation part:

$$\frac{dN}{dt} = r\left(1 - \frac{N}{K}\right)$$

However, sensitisation is located in the numerator (r) and tolerance in the denominator of the negative factor (K). Consequently, both result in an increased appetitive motivation and thus additional use. This perfectly aligns with the concept of describing sensitisation as ‘reverse tolerance’, as suggested by Robinson and Berridge (1993) when the notion was first introduced. Furthermore, the model provides a more detailed understanding of what this ‘reverse’ entails. I view this unintentional agreement as evidence that the model accurately reflects the psychological constructs. Specifically, the notion of tolerance was not deliberately included in the model, yet its presence corroborates the theoretical constructs and underlying dynamics. I see this unplanned concordance as another sign that supports the model’s validity as a representation of the underlying psychological processes involved.

On data

The ability to represent the theoretical relationships between constructs and to implement different situations and drugs without significant adjustment to the model is a strong indication of the theoretical validity of the model. However, unlike verbal theories, formal models have the advantage of being able to make empirical predictions that can be validated. With an appropriate model, the complex behaviours associated with individual substance use—which may initially appear erratic or nonsensical—can explain the patterns within the data in a comprehensible way. While I have obtained sufficient data to calibrate and somewhat validate the epidemiological model described in Chapter 4, the more sophisticated substance use model in Chapter 6 would require more granular data for adequate validation.

In principle, obtaining more detailed psychological data is like needing a telescope with a higher resolution: a larger sample size, more observables, and shorter time spans will go a long way. However, this more complex model does not only contain directly observable constructs: certain constructs can be seen as partially latent; for example, it is difficult to measure the control delay A directly. Nevertheless, with long-term observations of consumption, possibly combined with information on biological limits to consumption and availability K , it becomes possible to estimate other parameters and calibrate the model. The risk of overfitting, which increases with model complexity, is a potential problem but a surmountable challenge. The primary objective for accurate validation remains to improve the resolution of our ‘behavioural telescope’. Therefore, in my view, the mutually beneficial interaction of formal model predictions and data would benefit greatly from data collection efforts that take a much larger perspective, using modern data collection technologies, rather than the current trend of many smaller, targeted data collection efforts.

Path forward and conclusion

My initial view of the formalisation approach in psychology and my PhD was that it was a relatively unexplored route into the field — one of many different routes that could be taken to try to move the discipline forward. I saw my PhD as an experiment in trying this method on substance use; a necessary part of doing science; a potentially successful method with many possible limitations and pitfalls in its execution. However, as my

theoretical knowledge of psychology has grown, and as I have completed the chapters of this dissertation, I have come to believe that the formalisation of psychological theories is indeed a worthwhile endeavour in the advancement of psychological knowledge. It should not be expected that our descriptions of human behaviour, the product of billions of neurons taking in almost infinite amounts of information and history, can achieve the pinpoint accuracy that general relativity does in predicting the orbit of an inanimate round object orbiting the sun.

While we are quite advanced in terms of verbal theories in psychology, the formalisation of psychological theories is still at a stage more akin to the age of Newtonian dynamics; formal descriptions that can capture broader patterns of behaviour already represent a substantial progress in our ability to describe and understand the underlying dynamics. Moreover, the very act of formalisation brings clarity to the theoretical underpinnings and requires the unambiguous articulation of psychological constructs. Thus, while future work should take advantage of the primary goal of formalisation: to aid theory development by generating empirical predictions and accurately representing the real world, for which highly granular data are needed to validate and calibrate, I believe that it is not only the results of the formal models that advance the field, but also the process of formalisation itself.