Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

VIRTUAL OPIOID USER: REPRODUCING OPIOID USE PHENOMENA WITH A CONTROL THEORY MODEL

Alexander Preiss
Anthony Berghammer
Georgiy Bobashev

Center for Data Science
RTI International
3040 E Cornwallis Rd
Durham, NC 27709, USA

ABSTRACT

Few attempts have been made to simulate the complex natural history of opioid use disorder. We developed a model to simulate an agent's opioid use over 15-minute time-steps. We followed the principles of control theory and opponent-process theory, formalizing representations of several processes (tolerance, effect, craving) as weighted integrations of opioid concentration, which was modeled with a pharmacokinetic equation. We calibrated our model to reproduce five qualitative opioid use trajectories commonly observed in the literature. We demonstrate how a relatively simple control theory approach can reproduce many of the key characteristics of real-world opioid use.

1. INTRODUCTION

This paper presents a prototype of Virtual Opioid User: a control-theoretic model that simulates an agent's opioid use. (We call our simulation Virtual Opioid User—rather than a virtual person who uses opioids—because the model is intended only to represent opioid use and does not cover the full spectrum of features that characterize a person.) Although many simulation models intend to help understand the epidemiology of opioid overdose, less attention has been paid to understanding and modeling the natural history of opioid use disorder. However, such models provide the basis for public health prevention and clinical treatment interventions. Because the natural history of opioid use disorder is complex, data is unavailable to parameterize and calibrate many components of such models. Therefore, we calibrated our model to reproduce five common, qualitative opioid use trajectories. These trajectories apply to both licit and illicit opioid use. In this paper, we describe the design of the model and the results of the qualitative calibration.

The opioid crisis in the United States reached National Public Health Emergency level as declared in 2017. In 2016, 91.8 million (34.1%) U.S. civilian, noninstitutionalized adults used prescription opioids; 11.5 million (4.3%) misused them. In 2021, over 100,000 Americans died from overdoses, 600% more than in 2010 (CDC, 2022). The mortality rate has been rising over the period of the COVID-19 pandemic.

Virtual patients have been used broadly in clinical and medical research to help understand and simulate effects that are difficult or impossible to obtain in real world (Kononowicz et al. 2019). In the substance use field, little has been done to simulate the natural history of the disease, largely because substance use is very complex. It involves neurobiology, psychosocial factors, behavioral economics, and even political and cultural aspects of society (Bobashev et al. 2017; Bobashev et al. 2020).

Many population-level models have been built to inform opioid policy. Typically, agents are considered to pass through discrete states (e.g., abstinence, experimental use, recreational use, dependence, disorder, recovery) (Bobashev et al. 2018; Chen et al. 2019). These models assume immediate transitions between

states like those in infectious disease models. However, in substance use, these transitions can take years and include movements back and forth between stages. For example, it is difficult to define how dependent a person is on a drug. Often a substance use disorder is recognized after it manifests strongly, but the disorder could develop gradually, and each point of time in an agent's state could be defined as being between recreational and problematic use. This motivated us to develop models in fine-grained, almost continuous time that describe agent behaviors rather than hard labeling them as discrete states.

Psychiatric and epidemiologic literature provides a rich basis of narrative pathways an individual could take as they use drugs (Goldstein 2001; Koob and Volkow 2010). A general narrative is that an individual starts with a small dose, likes the effect of the drug, and starts using it on a regular basis. After a while, they develop dependence (i.e., feel strong discomfort when not using the drug). A defining component of dependence is withdrawal, a physiological response to the sudden quitting or slowing of use. With time, an individual develops tolerance (i.e., a higher dose is needed to produce the same effect). Therefore, the dose tends to escalate, and regular use creates brain plasticity (i.e., the brain does not recover to the initial, unaffected state after cessation of drug use). These irreversible effects lead to substance use disorders and require treatment (Koob and Le Moal 2006; Koob and Volkow 2010). Another aspect of drug dependence is craving (i.e., a strong desire to use the drug). Craving is critical in the context of quitting and relapse. If a person is capable of withstanding craving, it will eventually decrease, but craving persists in long-term memory and can trigger relapse even after a period of cessation.

1.1. Control Theory Models

In this study, we use control theory formalism to simulate the use of opioids and illustrate potential pathways toward increased use and overdose. In previous work, we introduced a control theory approach to modeling opponent process (see Section 1.2) and applied it to smoking and cocaine (Ahmed et al. 2007; Bobashev et al. 2007, Bobashev et al. 2017; Newlin et al. 2012). In this study, we used a similar approach to develop a model to simulate opioid use. To make the problem tractable, we focused on a subset of aspects of opioid use, namely consumption and choice of doses. In Bobashev et al. (2017), we illustrated our approach with the following example:

> When an individual is walking, several complex processes happen simultaneously in the brain, but from the control theory perspective, maintaining balance is just a coordination of feedback loops that can be formally modeled and realized in self-stabilizing devices such as a walking robot. Although the real stabilizing mechanisms in a human brain and a robot are different, the stabilizing feedback principles remain the same.

Here we developed an opioid-using agent whose behavior mimics patterns of opioid use of a human. The use behavior is controlled by formalized feedback mechanisms. Because several control theory models can lead to the same phenomenological description, we present "a" model rather than "the" model.

1.2. Opponent Process Models of Pathways to Dependence

Many interlocked neurobiological processes guide pathways to dependence and disorder, but one—opponent process—is key to the shifting setpoint characteristic of brain plasticity. In this study, we translate opponent process into mathematical terms and use it to guide agent behavior. Opponent process theory is based on the concept that the brain is designed to maintain a steady state (homeostasis). This state is characterized by a certain set point. When the system deviates from this point (e.g., because of an external stimulus), neurobiological processes bring the system back to the set point, thus maintaining a steady state. The process that restores homeostasis is thus called the opponent process, as introduced by Solomon and Corbit (1973). Applied to brain reward circuits, it means that when a strong reward (e.g., drug) creates a lasting effect, it must be countered by the opponent process to return the brain to the set point. Because the system does not need to respond to random noise, the opponent process is slower and responds to the signal

accumulated over an interval. Homeostasis, however, cannot explain a slow shift into drug dependence and disorder and must be modified by the concept of allostasis, where after prolonged use the set point shifts, and the system does not return to the original steady state. Opponent process thus leads to an important shift in individual motivation for consumption. When the opponent process accumulates, it overwhelms the system and starts driving substance use to counter the negative effect of the opponent process. After some time, substance use is motivated not by the pleasurable effects of the drugs, but the negative effects of the opponent process (Koob and Le Moal 2006; Koob and Volkow 2010). We therefore consider a cascade of processes. Each process in a sequence is slower than its predecessor, producing a delay between the peak values of these processes. The set point is a function of the combined effect of the slower processes, reflecting the accumulation of the deviation from the satiated state because of prolonged use.

1.3. Opioid Use Trajectories

In our study, we calibrated the model to reproduce several trajectories described in the literature as narratives. One is the pathway from the use of prescription opioids to gradual increase in dose and development of disorder, potentially leading to eventual overdose and death (Goldstein 2001). Another scenario is common for most pain patients, who are prescribed low doses of opioids and who use them without developing much tolerance, increasing dose, or developing problems (Hayden et al 2021). A similar scenario was described by Hoffer (2005) where opioid brokers who connect heroin buyers to dealers get rewarded with a small amount (tax) of heroin and maintain steady use of these small doses over time. Finally, we consider a scenario where an individual who has developed a high level of tolerance abruptly ceases opioid use (e.g., gets incarcerated) and experiences craving. We also show how taking opioids can lead to an overdose after such periods due to reduced tolerance. This scenario has been described in research showing that the rate of overdoses is up to 40 times higher among individuals who have been recently released from jail or prison (Frank et al 2017; Ranapurwala et al 2022; Victor et al 2022). Therefore, we calibrated our model to reproduce the following qualitative simulation scenarios:

- maintenance of low dose without craving or dose increase,
- gradual increase of dose at medium doses,
- rapid increase of dose and craving at high doses,
- increased craving during periods of cessation, and
- increased overdose risk when resuming use after a period of cessation.

2. METHODS

2.1. Overview

Virtual Opioid User is an open-source project. The code base is on GitHub (Preiss et al., n.d.) and an interactive app (Preiss, n.d.) is available to interact with the model.

The model core is a system of differential equations that represent an opponent process and control when the agent takes a dose. Other mechanisms such as tolerance, craving, dose increase, and overdose are formalized as functions of opioid use. External and individual factors such the drug type, starting dose, drug availability, and risk levels are inputs in the model. We use a pharmacokinetic decay function to translate opioids taken to the concentration of opioids in an agent's system at a given time. The model is a fixed-increment time simulation. There are 100 time-steps per day, with each time-step representing roughly 15 minutes. During calibration, we used a simulation length of 2 years. At each time-step t, a series of functions determines the agent's opioid use behavior. Broadly, the steps are the following:

1. Compute concentration of opioids at t using a pharmacokinetic model.
2. Compare concentration to the agent's threshold to determine if the agent will take a dose at t.
3. Compute tolerance and perceived effect of opioids in the agent's system at t.

4. Compute the agent's levels of craving at *t*.
5. If the agent takes a dose at *t*, check whether it caused an overdose.
6. If the agent takes a dose at *t*, compare its perceived effect to their preferred dose to determine if they will increase their preferred dose.
7. Recalculate the agent's use threshold based on opponent processes.

A schematic diagram of the model is presented in Figure 1.

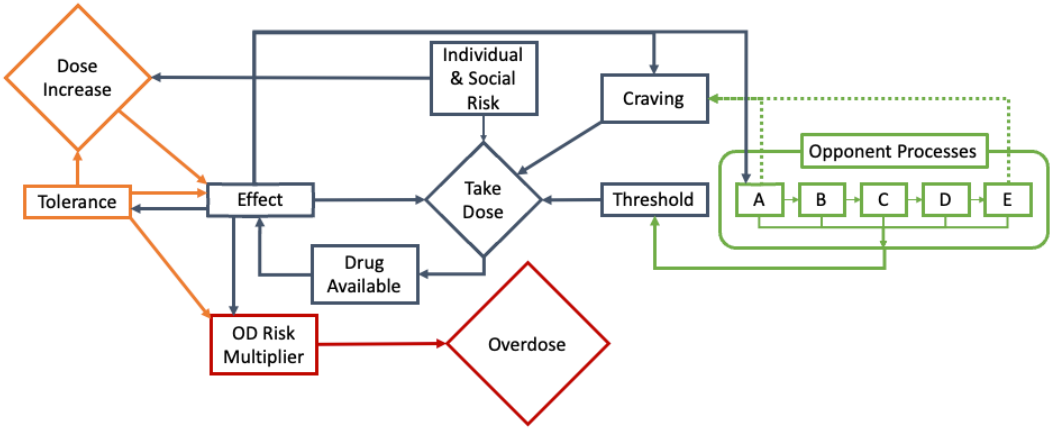

Figure 1: A schematic diagram of the model. Colors indicate linked components: green for the multiscale opponent process described in Section 2.2; orange for the tolerance and dose increase mechanism described in Section 2.5, red for the overdose mechanism described in Section 2.7, and blue for all other components described below. The model is flexible to incorporate other processes as necessary.

2.2. Multiscale Opponent Process Modeled as Weighted Integrations of Previous Processes

We consider a cascade of five continuous functions to represent a multiscale opponent process. The processes sequentially feed each other with linear accumulation and first-order extraction, which is consistent with models of many biological processes (Murray 2002). We used this approach in Newlin et al. (2012) and further developed the modeling framework in Bobashev et al. (2017), but the parameterization is different to reflect opioid use as opposed to smoking tobacco. Each process is characterized by a temporal scale associated with accumulation and extraction rates. Each process is constructed as a weighted integration of the previous process, thus, the scale of each is longer than the scale of the previous process. The equations are not designed to represent any *specific* biological process; they are designed to describe the observed phenomenon.

The first process, A, corresponds to the concentration of the drug, which is modeled with a pharmacokinetic equation. The second process, B, describes the accumulation of the drug and the body's processing of it. This process is modeled as a running weighted mean of process A. The third process, C, characterizes how much drug an agent consumes over a long period. This process thus reflects the recent history of opioid use and is defined on the scale of days. It was modeled as a running weighted mean of process B. The fourth process, D, is again a running weighted mean of process C. It does not have a clear biological interpretation and is added for consistency of opponent process (i.e., to prevent a large leap in scale from process C to process E). Finally, process E is a long-term hedonistic memory defined on the scale of years. After a long period of cessation when processes A through D are quite low or virtually zero, process E holds its slow-changing values. It is used primarily for calculating an agent's level of craving, which is discussed further below. Alphas and betas in the process equations are scaling parameters calibrated to reproduce our target trajectories. *P* is the concentration function discussed in Section 2.4.

$$\text{Process A: } \frac{dA}{dt} = P(t) - \beta_1 A$$

$$\text{Process B: } \frac{dB}{dt} = \alpha_1 A - \beta_2 B$$

$$\text{Process C: } \frac{dC}{dt} = \alpha_2 B - \beta_3 C$$

$$\text{Process D: } \frac{dD}{dt} = \alpha_3 C - \beta_4 D$$

$$\text{Process E: } \frac{dE}{dt} = \alpha_4 D - \beta_5 E$$

2.3. Use Threshold

Underlying the functionality of the model is a set point, or threshold, corresponding to a satiated state. When the actual state is below the threshold, the agent is motivated to use opioids. The threshold is dynamic and is a function of processes A through D, relying more on short- to medium-term memory of past use. When opioid use is high, the threshold gradually increases, and when opioid use is low, the threshold gradually decreases.

2.4. Pharmacokinetic Model of Opioid Concentration

We use a pharmacokinetic decay function calibrated to the half-life of morphine in plasma to measure how much of the opioid (in morphine milligram equivalents [MME]) was in the agent's system at any given time. Per Lötsch (2005), three studies identified the morphine plasma half-life as 2.8 hours, which we use in our calculation. We scale the concentration to MME based on the drug selected. Our model uses 100 time-steps per day (or 14.4 minutes per time-step). The half-life of opiates in our model's time units is 11.667 time-steps. This leads to a first-order decay constant of $k = \ln 2 / 11.667$, or a decrease of a factor of 0.0594 for each time-step.

The concentration of opioids in the agent's system is necessary to properly understand how an agent will interact with the drug. At every time-step in our model, we calculate the concentration of opioids and use this as input to calculate the agent's tolerance, perceived effect, and desire to increase their dose. This equation is modifiable to represent the usage of pills, injections, etc. To represent the rate at which opioids are reaching blood stream through digestion, the first-order decay equation is modified as follows, where k_a and k_e are coefficients of absorption and elimination respectively, F is an individual scaling coefficient, and d is the time since the last dose:

$$Concentration = F \times \left(\frac{k_a}{k_a - k_e}\right) \times \left[e^{(-k_e \times d)} - e^{(-k_a \times d)}\right].$$

2.5. Tolerance, Perceived Effect, and Dose Increase

To calculate the agent's opioid tolerance, we use a logistic function of the agent's recent opioid use. When an agent has been using opioids more often, tolerance will be higher, and when an agent has been using less, tolerance will be lower. Parameters of this function are relative to dose, resulting in the relationship between concentration and effect changing at different doses. In turn, this results in agents tending to maintain use of a low dose for a long time and increase their dose faster as the dose gets higher.

We use a standard logistic function to model tolerance: $Y = L / \left(1 + e^{(-k \times (x - x_0))}\right)$. The input to the logistic function (x) is a rolling mean of the agent's opioid concentration over a window of past time-steps, multiplied by a constant. The number of past time-steps and the multiplier are calibrated parameters. The parameters of the logistic function are functions of the agent's current preferred dose and several calibrated

parameters. The maximum of the logistic curve is defined as $L = dose^{l_1} \times l_2$. The growth rate is defined as $k = k_1 - dose \times k_2$. The midpoint of the logistic curve is defined as $x_0 = dose \times x$.

The agent's perceived effect is a function of opioid concentration and tolerance. When a weighted average of the effect of recent doses passes below their use threshold, the agent is motivated to increase their preferred dose.

2.6. Craving

Craving plays a key role in an individual's motivation to take opioids, their motivation to cease taking opioids, and their desire to increase/decrease dose. Craving is calculated from the longest-term memory of past use (Process E above) moderated by the agent's threshold, which decays more quickly after cessation. The overall effect of combining threshold and long-term memory is that craving peaks rapidly upon cessation, then gradually decays. Craving directly influences the availability of opioids in the model, since an agent experiencing stronger craving will try harder to seek opioids, thus increasing their likelihood of use. We define craving as $Cr = d_5 E(T - A)/(S + E)$, where T is a satiation threshold, S is a calibration coefficient, and A and E are processes.

2.7. Overdose

One of the primary outcomes of interest in our model is simulating the variables that can lead to opioid overdose. For each single dose taken, we calculate the agent's probability of overdose from that dose, based on their concentration of opiates, tolerance, and a baseline overdose risk function. The baseline overdose risk function is derived from Dasgupta et al. (2015), which presents population-level overdose risk statistics at various opioid prescription doses. We fitted a logistic model to their data, along with the assumption that a single dose of 2,000 MME has an overdose probability of 1. Because Dasgupta et al. (2015) used prescription data, we assume that the overdose risks they reported are for people who are tolerant to their prescribed dose. Therefore, we add an excess risk multiplier, based on the ratio of the dose to the agent's tolerance. We define excess as $(dose/tolerance - 1)^2$. Baseline risk (from the function derived from Dasgupta et al.) is multiplied by the excess risk multiplier (and then limited to 1) to create a tolerance-adjusted overdose probability. For each dose taken, a random draw is compared to the tolerance-adjusted overdose probability for that dose to determine whether the agent overdosed.

Once an agent has overdosed, we assess whether the overdose was fatal. Each overdose has the potential of having a fatal outcome. We modeled this using the statistic that 1 of every 8.5 overdoses is fatal, per Dunn et al. (2010).

Once an agent has overdosed, assuming it was not fatal, we adjust the agent's behavior moving forward. We calculate the amount of time the agent will cease using opioids after overdose using an exponential distribution, which is scaled by their social and individual risk (see below). This period can range from a few hours to 60 days. This stoppage can also be coupled with a reduction in dose, which is also based on the agent's social and individual risk. The lowest-risk agent will reduce their dose by half, while the highest-risk agent will maintain the same dose.

2.8. Inputs

The model is flexible to add additional inputs to conduct various experiments. Below we describe the inputs in our application and how they play a role in our simulation outcomes. For our initial calibration, the following inputs were most important. Other inputs are listed in Table 1.

- Starting dose: The agent starts the simulation taking their preferred dose consistently. This parameter controls their preferred dose at the start of the simulation.
- Composite risk factors for the agent and their environment:
    - Social risk: a composite of external/environmental factors (e.g., social determinants of health) motivating the agent to use opioids and seek increased effects from them

  - Individual risk: a composite of psychological/biological factors (e.g., risk tolerance) motivating the agent to use opioids and seek increased effects from them
- Use pattern: to simulate the effects of resuming opioid use after a period of cessation, each model run can either include or exclude a 6-month period where opioids are unavailable in the middle of the simulation.

Table 1: Other inputs that can be varied across simulation runs.

Input	Description
Dose Variability	The agent always has a preferred dose. If this parameter is 0, they always take that exact dose. If this parameter is > 0, each dose varies from the preferred dose.
Opioid Type	The agent takes one opioid type throughout the simulation. They can increase their dose over time, but always continue to take the same opioid. The chosen opioid is converted to morphine milligram equivalents within the simulation.
Opioid Availability	The probability that, each time an agent intends to take a dose, opioids will be available for them to take.
Dose Increase Amount	When the agent is no longer satisfied with the effect of their preferred dose, they may increase their preferred dose. This parameter controls the amount by which they will increase their preferred dose.

2.9. Calibration

Our model used pharmacokinetic parameters from published literature representing human patients (Lötsch 2005). Dose-specific overdose risk parameters were derived from Dasgupta et al. (2015). Other parameters were drawn from expert opinion, such as the distribution of cessation length following overdose. However, many control theory parameters were not directly observable or did not have biological interpretation and therefore were calibrated to produce the five target opioid use trajectories described above. Because our calibration targets were qualitative, the typical calibration approach of minimizing a loss function was not feasible. Therefore, we used an iterative, human-in-the-loop calibration process, in which we calibrated individual functions to produce desired behavior, then combined functions, adjusting as needed to maintain desired behavior. For example, we calibrated the tolerance-building function independently as a starting point. Then, once it was integrated to the full model, we made small adjustments to the function's parameters to produce the desired trajectories. The trajectories were then reviewed by ethnographers and opioid treatment psychiatrists for feasibility.

3. RESULTS

Our model reproduced the five target opioid use trajectories.

3.1. Maintenance of Low Dose Without Craving or Dose Increase

When the simulation starts with a relatively low dose, later dose increases are unlikely, as seen in Figure 2 below. This is due to the design of the tolerance function. At lower doses, the logistic curve has a smaller growth rate, midpoint, and maximum. In other words, agents become less tolerant of lower doses and tolerance builds more slowly.

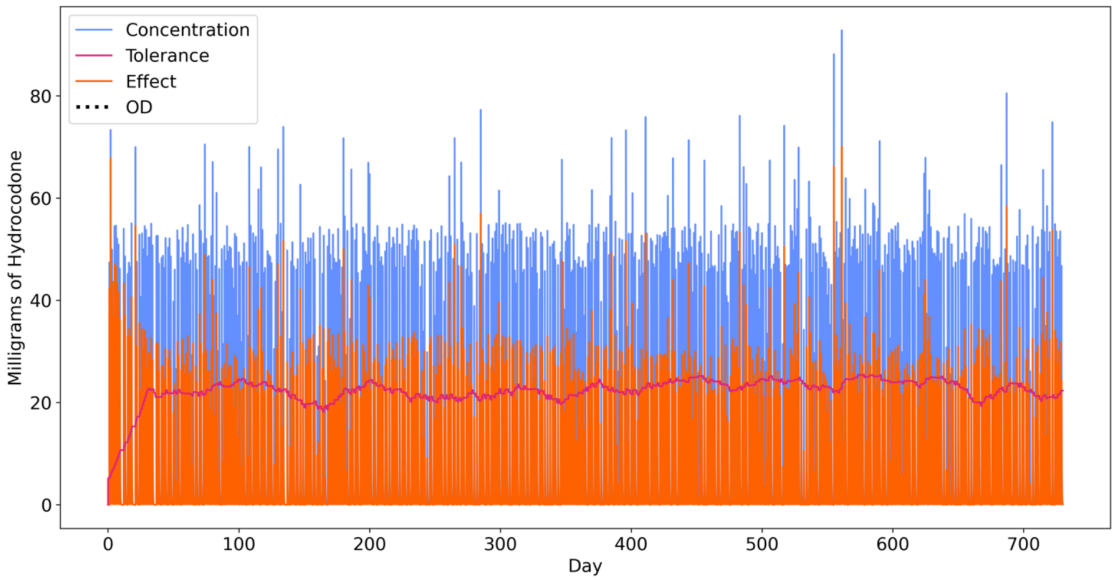

Figure 2: An average agent with a starting dose of 50 milligrams of hydrocodone maintains their starting dose throughout the simulation.

3.2. Gradual Dose Increase at Medium Starting Doses

At somewhat higher starting doses (around the Centers for Disease Control and Prevention's maximum recommended dose of 90 MMEs), agents are more likely to increase their dose during the simulation, as seen in Figure 3. At these doses, tolerance builds more quickly, reducing perceived effect relative to preferred dose and motivating dose increase.

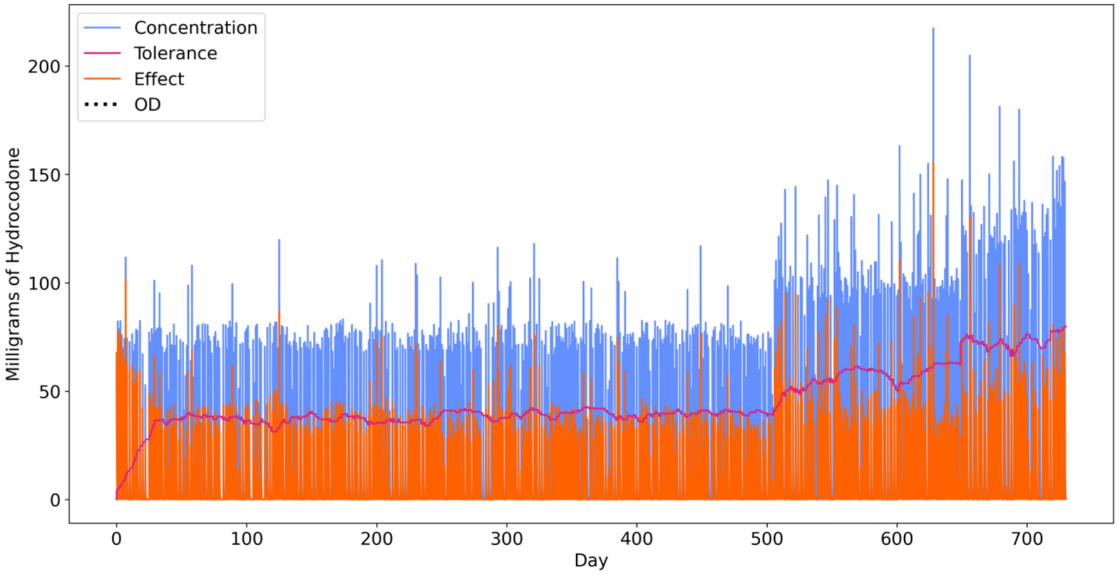

Figure 3: An average agent with a starting dose of 70 milligrams of hydrocodone begins to increase their dose after about 500 days.

At medium starting doses, other parameters in the model are especially important in determining the agent's trajectory. Parameters like opioid availability, social risk, and individual risk tend to be dominated at very low or high starting doses. However, at medium doses, they can have great impact. For example, at the same starting dose of 70 milligrams of hydrocodone, an agent with a higher individual risk level tends to increase their dose more than an average agent (not shown). However, at a starting dose of 50 MME, agents tend to maintain the same dose regardless of risk level. Conversely, at a starting dose of 100 MME, even agents with low risk levels tends to increase their dose over time.

3.3. Rapid Dose Increase at High Starting Doses

At starting doses above 90 MMEs, agents tend to increase their dose rapidly. As tolerance builds at a greater rate with ever higher doses, this can lead to extreme dose increases over the course of the simulation, which greatly increases the likelihood of overdose. No figure has been included for this scenario as the trajectory is similar to Figure 3, but with a faster and greater increase in dose.

3.4. Increased Craving During Subsequent Periods of Cessation

When an agent abruptly ceases opioid use for a long period (simulating a cold-turkey treatment program or incarceration), their levels of craving increase rapidly, peak within a few days, then begin to slowly decay, as shown in Figure 4 below. This represents the high desire to use opioids during withdrawal and the long-term desire to use opioids driven by hedonic memory. Increased history of opioid use prior to cessation leads to a higher peak and longer duration of craving.

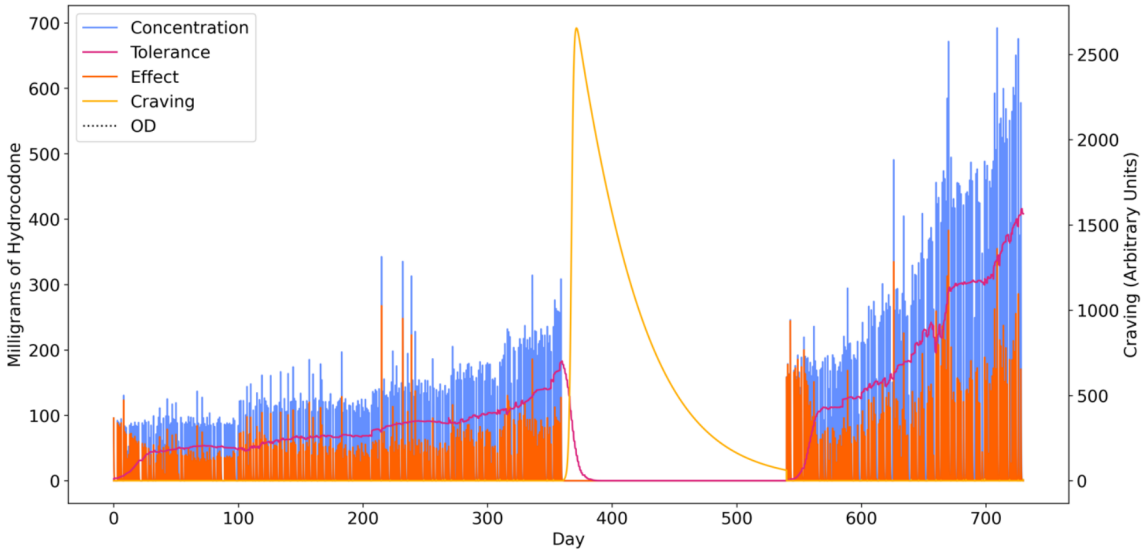

Figure 4: An agent with a dose of 225 milligrams of hydrocodone experiences peak craving within a week of cessation. After 6 months of cessation, craving nears zero. Craving is not shown in other figures for simplicity, but the craving mechanism was included in all model runs.

3.5. Increased Overdose Risk When Resuming Use After a Period of Cessation

Tolerance decays during periods of cessation. When agent resumes use after cessation, the effect of a given dose is much higher, leading to higher risk of overdose. In Figure 4, the agent reduces their dose when they resume use after cessation. In Figure 5, the agent resumes use at the same dose. If an agent who is accustomed to a high dose resumes use at the same dose after cessation, overdose risk can be very high.

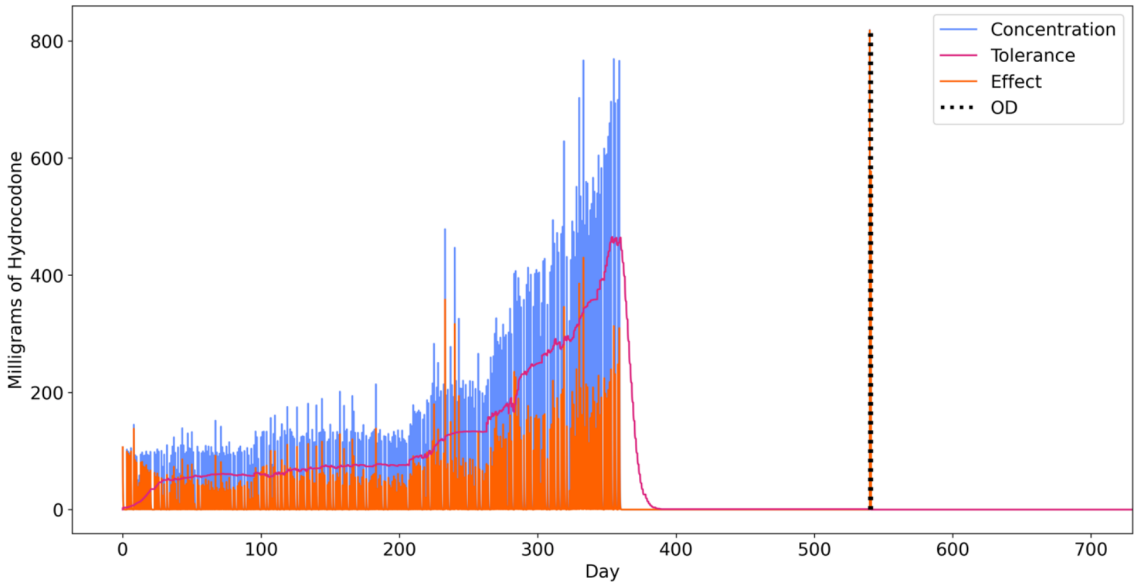

Figure 5: An agent resumes taking a dose of 700 milligrams of hydrocodone after 6 months of cessation and overdoses.

4. DISCUSSION

We developed a theoretical model that can qualitatively reproduce five key trajectories of opioid use. We showed that a simple control theory model with variability in the starting dose, use pattern, and risk factors can reproduce these patterns. We formalized a multiscale opponent process to control the allostatic threshold for opioid use, as well as mechanisms of tolerance, dose increase, and craving, which form feedback loops with opioid use. We also developed an overdose risk module, which determines whether each dose taken by an agent causes overdose.

Our theoretical model is an early step in a research program with various implications. From a public health perspective, improved simulation of individual pathways to dependence and disorder could help improve our understanding of these phenomena, ultimately leading to advances in prevention and treatment of substance use disorders. Methodologically, simulating substance use continuously rather than in discrete states offers theoretical and practical advantages. Agent-based models and microsimulations could use populations of Virtual Opioid Users to simulate emergent behavior more accurately. Finally, like other virtual patient models, our model could be used to help clinicians better understand their patients. For example, medical students could interact with our application [Preiss n.d.] to see how different dosing practices affect long-term outcomes.

4.1. Limitations

We acknowledge that true calibration to real-world data currently is not possible for many parameters due to the lack of corresponding data. Our model is conceptual and provides a theoretical base rather than being fully practical. Therefore, major limitations are related to the validity of many of the model assumptions. These key assumptions include the mechanism causing the increase in tolerance, the probability of overdosing given dose and tolerance, and the somewhat ad-hoc formulation of a withdrawal and craving relationship with the opponent processes. The functional forms of these processes were selected to represent narratives described by scientific literature, addiction psychiatrists, and people who discuss opioids on Reddit. Despite this limitation, such conceptual models can still have practical uses beyond theory

development and serving as a foundation for future models. For example, Abo-Tabik et al. (2020) combined our similar virtual smoker model (Bobashev et al. 2017) with external data to significantly improve the forecast of relapse episodes among smokers.

4.2. Future Work

We have two main goals for the future development of Virtual Opioid User. First, we plan to validate the model more rigorously. Second, we plan to use the model as the foundation for various microsimulations and agent-based models, adding features as necessary.

Validating the model entails the challenging task of identifying real-world data for many difficult-to-measure constructs. Nevertheless, we are optimistic that real-world calibration targets exist for many of the model's components. For example, parameters in the excess overdose risk multiplier could be calibrated to the overdose incidence rate in a population of interest. More challenging examples include the allostatic threshold for opioid use and the dose increase mechanism. Opponent process parameters could conceivably be calibrated to observed patterns of opioid use from clinical or ethnographic research. Similarly, tolerance function parameters could be calibrated to dose increases from longitudinal opioid prescription data. Although such calibrations would be imperfect, and no calibration target exists for the model's overall output, calibrating individual components in this fashion would increase the model's practical usefulness.

We plan to use Virtual Opioid Users as agents to simulate the use of counterfeit pills (potentially laced with fentanyl); the mixture of pills and heroin; environmental factors (e.g., new prescription guidelines); the effects of medication-assisted treatments; and polysubstance use (e.g., heroin and benzodiazepines). Because of the model's modular structure, components can be added or replaced with more advanced models when needed. For example, the addition of dependence and disorder would require incorporation of psychosocial factors such as the failure to fulfil work and family obligations. This module could be added to the base simulation when the timing of use interferes with external environmental prompts. Individual characteristics such as risk scores and parameters guiding tolerance building could interact with environmental factors to create realistic psychosocial outcomes. So far, we model the environment only through the availability of drugs and a composite environmental risk factor. Even with this simple parameterization, we can simulate increased overdose risk after periods when drugs are not available.

REFERENCES

Abo-Tabik, M., N. Costen, J. Darby, and Y. Benn. 2020. "Towards a Smart Smoking Cessation App: A 1D-CNN Model Predicting Smoking Events." *Sensors 20*(4):1099.

Ahmed, S. H., G. Bobashev, and B. S. Gutkin. 2007. "The Simulation of Addiction: Pharmacological and Neurocomputational Models of Drug Self-Administration." *Drug and Alcohol Dependence 90*(2–3):304–311.

Bobashev, G., E. Costenbader, and B. Gutkin. 2007. "Comprehensive Mathematical Modeling in Drug Addiction Sciences." *Drug and Alcohol Dependence 89*(1):102–106.

Bobashev, G., S. Goree, J. Frank, and W. Zule. 2018. "Pain Town, an Agent-Based Model of Opioid Use Trajectories in a Small Community." In *13th International Conference of Social, Cultural, and Behavioral Modeling*, edited by R. Thomson,H. Bisgin, C. Dancy, A. Hyder, and H. Bisgin, 274–285). Cham: Springer International Publishing.

Bobashev, G., J. Holloway, E. Solano, and B. Gutkin. 2017. *A Control Theory Model of Smoking*. Methods Rep RTI Press: Author Manuscript. https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5578474/pdf/nihms895211.pdf accessed October 11, 2022.

Bobashev, G. V., L. D. Hoffer, and F. R. Lamy. 2020. "Agent-Based Modeling to Delineate Opioid and Other Drug Use Epidemics." In *Complex Systems and Population Health*. Edited by Y. Apostolopoulos, K. H. Lich, M. K. Lemke, 201-205, New York: Oxford University Press.

Chen, Q., M. R. Larochelle, D. T. Weaver, A. P. Lietz, P. P. Mueller, S. Mercaldo, S. E. Wakeman, K. A. Freedberg, T. J. Raphel, A. B. Knudsen, P. V. Pandharipande, and J. Chhatwal. 2019. "Prevention of Prescription Opioid Misuse and Projected Overdose Deaths in the United States." *JAMA Network Open 2*(2):e187621.

Dasgupta, N., M. J. Funk, S. Proescholdbell, A. Hirsch, K. M. Ribisl, and S. Marshall. 2015. "Cohort Study of the Impact of High-dose Opioid Analgesics on Overdose Mortality: Prescribed Opioid Dose and Overdose Mortality." *Pain Medicine* 17(1):85-98.

Dunn, K. M., K. W. Saunders, C. M. Rutter, C. J. Banta-Green, J. O. Merrill, M. D. Sullivan, C. M. Weisner, M. J. Silverberg, C. I. Campbell, B. M. Psaty, and M. Von Korff. 2010. "Opioid Prescriptions for Chronic Pain and Overdose: A Cohort Study." *Annals of Internal Medicine 152*(2):85–92.

Frank, J. W., T. I. Lovejoy, W. C. Becker, B. J. Morasco, C. J. Koenig, L. Hoffecker, H. R. Dischinger, S. K. Dobscha, and E. E. Krebs. 2017. "Patient Outcomes in Dose Reduction or Discontinuation of Long-Term Opioid Therapy: A Systematic Review." *Annals of Internal Medicine 167*(3):181–191.

Goldstein, A. 2001. *Addiction: From Biology to Drug Policy*. New York: Oxford University Press.

Hayden, J. A., J. Ellis, M. Asbridge, R. Ogilvie, R. Merdad, D. G. Grant, S. A. Stewart, and S. Campbell. (2021). "Prolonged Opioid Use Among Opioid-Naive Individuals After Prescription for Nonspecific Low Back Pain in the Emergency Department." *Pain 162*(3):740–748.

Hoffer, L. D. 2002. *Junkie Business: The Evolution and Operation of a Heroin Dealing Network*. Ph.D. thesis, Department of Health and Behavioral Sciences, University of Colorado at Denver, Denver, Colorado. https://www.proquest.com/docview/305492615, accessed 12th October 2022.

Kononowicz, A. A., L. A. Woodham, S. Edelbring, N. Stathakarou, D. Davies, N. Saxena, L. T. Car, J. C.- Duke, J. Car, and N. Zary. 2019. "Virtual Patient Simulations in Health Professions Education: Systematic Review and Meta-Analysis by the Digital Health Education Collaboration." *Journal of Medical Internet Research 21*(7):e14676.

Koob, G. F., and M. Le Moal. 2006. Neurobiological Theories of Addiction. In *Neurobiology of Addiction*, edited by G. F. Koob and M. Le Moal, 377–428. London, Burlington, San Diego: Academic Press imprint of Elsevier.

Koob, G. F., and N. D. Volkow. 2010. "Neurocircuitry of Addiction." *Neuropsychopharmacology: Official Publication of the American College of Neuropsychopharmacology 35*(1):217–238.

Lötsch, J. 2005. "Pharmacokinetic–Pharmacodynamic Modeling of Opioids." *Journal of Pain and Symptom Management 29*(5):90–103.

Murray, J. D. 2002. *Mathematical Biology: I. An Introduction*. 2nd ed. Berlin Heidelberg: Springer.

Newlin, D. B., P. A. Regalia, T. I. Seidman, and G. Bobashev. 2012. "Control Theory and Addictive Behavior." In *Computational Neuroscience of Drug Addiction*, edited by B. Gutkin and S. H. Ahmed, 57–108. New York: Springer.

Preiss, A. J. n.d. *Virtual Opioid User*. https://share.streamlit.io/rtiinternational/virtual-opioid-user/main, accessed 5th, 2022.

Preiss, A. J., M. Duprey, and A. Berghammer. n.d. *Virtual Opioid User: Continuous Model of an Individual's Opioid User Over Time*. https://github.com/RTIInternational/virtual-opioid-user, accessed 5th July, 2022.

Ranapurwala, S. I., M. C. Figgatt, M. Remch, C. Brown, L. Brinkley-Rubinstein, D. L. Rosen, M. E. Cox, and S. K. Proescholdbell. 2022. "Opioid Overdose Deaths Among Formerly Incarcerated Persons and the General Population: North Carolina, 2000–2018." *American Journal of Public Health 112*(2):300–303.

Solomon, R. L., and J. D. Corbit. 1973. "An Opponent-Process Theory of Motivation: II. Cigarette Addiction." *Journal of Abnormal Psychology 81*(2):158–171.

Victor, G., C. Zettner, P. Huynh, B. Ray, and E. Sightes. 2022. "Jail and Overdose: Assessing the Community Impact of Incarceration on Overdose." *Addiction 117*(2):433–441.

AUTHOR BIOGRAPHIES

ALEXANDER PREISS is a Research Data Scientist in the Center for Data Science at RTI International. He received his Master's degrees in International Studies and Analytics from North Carolina State University. He combines expertise in machine learning, natural language processing, and simulation modeling with domain knowledge in public health and international development. His current research spans a variety of methods and domains, including simulation modeling, forecasting, network analysis, causal inference, COVID-19, colorectal cancer screening, and opioid use. His e-mail address is apreiss@rti.org. His website is https://www.rti.org/expert/sandy-preiss.

ANTHONY BERGHAMMER is a Research Data Scientist in the Center for Data Science at RTI International. He has experience working across different industries including manufacturing, technology and government in data analysis, visualization, machine learning and natural language processing capacities. During his time working with large manufacturing clients, he helped organizations plan and forecast demand to drive inventory optimization. His e-mail address is aberghammer@rti.org. His website is https://www.rti.org/expert/anthony-berghammer.

GEORGIY BOBASHEV is an RTI Senior Fellow in the Center for Data Science at RTI International with over 30 years of experience in health research. His current research interest is focused on the applications of predictive modeling to substance use research. Predictive methods often combine machine learning with mechanistic (e.g., agent-based and system dynamics) modeling and require linking multiple datasets of various nature. He applies modeling, statistical analysis and machine learning to health-related areas, including substance use, HIV, child/maternal health, influenza, measles, cancer, diabetes, and violent behavior. His e-mail address is bobashev@rti.org. His website is https://www.rti.org/expert/georgiy-bobashev.