ANNUAL REVIEWS

Annual Review of Criminology

Agent-Based Modeling in Criminology

Daniel Birks,[1] Elizabeth R. Groff,[2] and Nick Malleson[3]

[1] School of Law, University of Leeds, Leeds, United Kingdom; email: d.birks@leeds.ac.uk

[2] College of Liberal Arts, Temple University, Philadelphia, Pennsylvania, USA

[3] School of Geography, University of Leeds, Leeds, United Kingdom

Annu. Rev. Criminol. 2025. 8:75–95

The *Annual Review of Criminology* is online at criminol.annualreviews.org

https://doi.org/10.1146/annurev-criminol-022222-033905

Keywords

agent-based models, individual-based models, simulation, computational criminology, complex systems, research methodology

Abstract

An agent-based model is a form of complex systems model that is capable of simulating how the micro-level behavior of individual system entities contributes to macro-level system outcomes. Researchers draw on theory and evidence to identify the key elements of a given system and specify behaviors of agents that simulate the individual entities of that system—be they cells, animals, or people. The model is then used to run simulations in which agents interact with one another and the resulting outcomes are observed. These models enable researchers to explore proposed causal explanations of real-world outcomes, experiment with the impacts that potential interventions might have on system behavior, or generate counterfactual scenarios against which real-world events can be compared. In this review, we discuss the application of agent-based modeling within the field of criminology as well as key challenges and future directions for research.

WHAT IS AN AGENT-BASED MODEL?

Crime is the output of a complex dynamic system. Criminal incidents emerge from a myriad of personal behaviors and interactions involving potential victims, offenders, and crime preventers, all set within their own interdependent social and physical environments. This complexity shapes the specific locations and times where crimes take place and determines who is involved and who may be affected. As criminologists, we rarely observe these individual-level processes and instead are left to analyze their outcomes as expressed through crime trends, hotspots, patterns of repeat victimization, and recidivism rates. Of course, this problem is not unique to criminology—core to the social science endeavor is a desire to connect observed societal outcomes to the actions and interactions of individuals.

Agent-based modeling is a relatively new analytical approach that complements traditional theoretical and empirical research methods by providing a means to explore how individual actions at the micro level are linked to broader outcomes at the macro level. Gaining momentum over the past three decades and informed by a range of fields, including game theory, computer science, artificial intelligence, and complexity theory, agent-based models (ABMs) enable researchers to create artificial societies (Epstein & Axtell 1996) in which simulated agents (most often individuals) act and interact according to hypotheses specified by a researcher. In contrast to traditional analytical models, ABMs are often described as bottom-up or generative in nature, capable of explicitly modeling the dynamic actions of and interactions among individuals and, in turn, examining how these generate higher-order properties such as wealth distributions, segregation patterns, or crime hotspots. Some researchers regard this ability to simulate causal processes and "grow" a societal pattern as more likely to foster a true understanding of the system of interest by revealing dynamics that may not be immediately apparent through traditional analytical methods. As Epstein (2006) posits, "If you didn't grow it, you didn't explain it."

In practice, an ABM is a type of computer simulation. In the model, agents represent individual system entities, which are typically people but could also be households, organizations, or political parties (Bonabeau 2002). Drawing on both theory and empirical evidence, researchers program agents with their own characteristics and behaviors that dictate how they will act in varying circumstances. When a model is executed, each agent operates autonomously, making decisions on the basis of these behaviors and its own circumstances, which might encompass the agent's current and previous states, those of nearby agents, and the simulated social or physical environment in which agents are typically situated. Within the field of criminology, hot-spot generation offers an illustrative example. Researchers who are interested in exploring the potential causal pathways that generate commonly observed concentrations of crime might build an ABM that includes offender agents navigating a street network, perceiving and reasoning about potential victim agents they encounter, and ultimately committing simulated crimes. Over the course of a simulation, crime locations as well as the individual agent and place characteristics are tracked and analyzed.

Key to the agent-based modeling approach is the concept of emergence—complex system outcomes that arise from interactions among individual agents that cannot be reliably predicted by examining individual system entities in isolation. As ABMs run, agents repeatedly interact and researchers observe how different characteristics, behaviors, and initial conditions influence emergent system properties and, in turn, derive insights into the system they seek to simulate (Epstein & Axtell 1996). In the above hot-spot generation illustration, simulated patterns of crime generated within the model are an emergent property of the spatiotemporal behaviors of all victim, offender, and crime preventer agents. These behaviors are influenced by both the morphology of the street network and the nature of the places that can be found within it. Accordingly, the interplay among agent characteristics and preferences that play out within this environmental backcloth influence where, when, and in what contexts crime is perpetrated and prevented.

Proponents of agent-based modeling argue that one of its main strengths is the natural approach that it offers to the study of complex human systems (Bonabeau 2002, Gilbert 2008). Agents are autonomous in that they perceive, reason, act, and interact on the basis of behaviors specified by the researcher. Agents can also be heterogeneous, differing from one another in terms of characteristics and behaviors. ABMs can simulate localized interactions, such that agents interact with one another as a result of social or physical proximity within some simulated environment. Agents can also change over time as a simulation runs, allowing previous states to influence future behavior and, in turn, the development of endogenous effects and feedback loops. Agents need not exhibit absolute rationality or knowledge of their world, with behaviors instead designed to reflect localized perception or bounded decision-making. ABMs typically incorporate stochasticity, such that each time a simulation is run different outcomes may emerge. Consequently, researchers enact multiple simulations, exploring the range of possible futures that may result from the same initial conditions. Returning to our example of modeling potential explanations of crime hotspots, these characteristics of the agent-based modeling approach can play out in myriad ways. We can explore different theoretically or empirically inspired behaviors for navigation, target selection, or crime prevention; potential victims can be perceived as more or less attractive to offenders; such attractiveness can be the product of agent characteristics, offender preferences, and the localized environment in which encounters occur; offender agent motivation can change over time as a function of previous successful or unsuccessful attempts to commit crime or the time since the last offense was committed. As we run the model, we may devise varying configurations that contrast with these assumptions and examine, over hundreds or thousands of distinct simulation contexts, how they influence where, when, and to what degree crime hotspots emerge.

Over the past three decades, ABMs have been applied in a wide variety of fields, including epidemiology and public health (Tracy et al. 2018), economics (Bookstaber 2017), the social sciences (Hedström & Ylikoski 2010, Luke & Stamatakis 2012, Macy & Willer 2002), political science (Marchi & Page 2014), and the ecological and environmental sciences (Grimm & Railsback 2005). Unsurprisingly, then, the approach has also garnered interest within the criminological community, with models developed to explore a range of distinct research questions of interest to criminologists (Gerritsen 2015, Groff et al. 2019, Liu & Eck 2008a). The following sections review some of the motivations for building ABMs specifically within criminology and discuss some relevant literature, before outlining the process of building an ABM and the main challenges and limitations associated with use of the methodology. We conclude with a brief discussion of several key opportunities for future agent-based modeling research in criminology.

WHY BUILD AGENT-BASED MODELS?

The nature of ABMs makes them well suited for constructing a "computational laboratory" (Groff et al. 2019): an ABM of a given system that can then be used to conduct in silico experiments. Although the results of these experiments must necessarily be interpreted differently from real-world experiments (Liu & Eck 2008b), ABMs exhibit several distinct properties that complement traditional theoretical and empirical efforts within criminology.

First, as discussed above, ABMs provide an intuitive means of modeling the emergence of criminological outcomes from the situated actions of and interactions among individuals. Building an ABM necessitates thinking about how individual-level mechanisms generate criminological phenomena and provides a formalized platform with which to consider these causal pathways. This ability to explicitly link micro and macro accounts of crime and criminality remains a fundamental challenge within empirical criminology. For example, crime prevention interventions that aim to affect individual decisions (and thus rely on an effective understanding of those micro-level mechanisms) are often evaluated using aggregate or small-scale sample data. Conversely, studies of habitual offenders or criminal careers remain at the individual level and are rarely linked to area-level crime patterns. At the same time, area-based evaluations common in enforcement-led analyses typically concentrate on retrospective descriptions of localized crime patterns, utilizing aggregate geo-demographic data sets, and consequently may be vulnerable to the ecological fallacy. Collectively, these challenges mean that we are often required to make a leap of faith in ascribing observed outcomes to proposed individual-level behaviors and vice versa, hampering both theoretical and intervention development. Consequently, any approach that may provide insights into this micro/macro divide is a welcome addition to existing methodological tools.

Second, the synthetic nature of ABMs offers researchers levels of experimental control and observation that are largely unattainable in traditional empirical research settings. Any element of an ABM can be systematically manipulated. Researchers can compare how agents behave in different environments by instantiating simulations with different street networks or social structures. Agent characteristics or behaviors can be manipulated to compare agents of a certain type, that behave in a certain way, or do or do not receive a given treatment against agents with other traits. All these manipulations can be undertaken without changing any other element of the model, thereby providing in silico counterfactuals (Nagin & Sampson 2019) for which there are no empirical equivalents. Termed by Holland (1986) the "fundamental problem of causal inference," this issue arises when we cannot directly observe the impact of two different treatments on a single experimental subject.

ABMs also enable absolute observation. Within a model, researchers can observe any system element at any time and any scale, simply by writing appropriate program code to record the variable of interest as the model is run. For example, in our illustrative model of crime hotspots, one might record the locations, times, and types of crimes committed. Such observations are often similar to those made within the real world, allowing simulated data to be compared with real-world data to explore model validity. However, models also enable observations that would be impossible for practical or ethical reasons—like the minute-by-minute movements of an entire population, the internal decision calculi of offenders, or even the locations of near crimes (those where an agent came close to the decision to commit a crime but ultimately decided not to act).

Collectively, these properties mean that criminologists can use agent-based modeling in a variety of contexts and explore a diverse range of scenarios and hypotheses, in turn potentially generating deeper insights into the behavior of complex systems of interest. That said, it is also essential to recognize that ABMs have limitations and assumptions inherent in their design, which may affect the accuracy and generalizability of their findings and can often make alternative approaches preferable. We discuss these limitations in detail in the section titled Key Challenges. In the following section, we discuss various ABMs that demonstrate the scope of their use in criminological research, highlighting several unique types of contributions to the discipline. For a broader overview of relevant research, we refer the readers to Liu & Eck (2008a), Groff & Mazerolle (2008), and Gerritsen & Elffers (2020).

APPLICATIONS OF AGENT-BASED MODELING IN CRIMINOLOGY

Exploring Criminological Theory

The simplest and perhaps the most powerful application of agent-based modeling within criminology is to act as a tool to support thinking about the system of interest (Liu & Eck 2008b). Criminological theories can be thought of as collections of interlinked causal assumptions. The same is true of agent-based modeling. The process of developing an ABM requires researchers to explicitly specify causal mechanisms of interest. To build a model, one must write program code for how action A influences behavior B of agent Y: What makes a peer criminogenic? At what speed do offenders learn about their surroundings? Under what circumstances is an offender dissuaded from offending? This requirement offers epistemic advantages far beyond the model itself by forcing one to formally specify all assumptions used to inform one's own internal model of a given system, which in turn guides any model (be it theoretical, analytical, or computational) that is subsequently built.

Of course, this is often not a simple task and will require the synthesis of diverse theories, data sources, and perspectives. Moreover, accurate data and robust theory are often scant as a direct result of the challenges facing criminology discussed above. Nevertheless, in a range of contexts, an ABM can act as an integrative platform capable of organizing diverse sources of insight, evidence, and hypotheses about a system of interest. In this way, models act to organize proposed explanations. Once a model is constructed, it can be used to systematically investigate the consequences of those explanations and, in turn, their validity. Considering this facet of agent-based modeling, Epstein (2008) argues that the most significant contribution of the modeling enterprise is that it "enforces a scientific habit of mind."

To illustrate, Groff (2007a,b; 2008), Wang et al. (2008), and Zhu & Wang (2021) present various ABMs of street robbery, each of which draws heavily on propositions put forward in Cohen & Felson's (1979) routine activity approach. These models include populations of agents representing potential victims, offenders, and crime preventers that navigate realistic street networks drawn from GIS data and interact with one another. Using these models, the authors conduct various experiments that provide a unique means of testing key predictions made by theory. Groff (2007a) tests a core proposition of the routine activity approach inaccessible to empirical inquiry: that changes in the level of crime can occur as a result of changes in routine activities of people while the size of potential offender and victim populations remains fixed. The model developed by Groff demonstrates that as time spent away from the home increases, so does the likelihood of victimization, even when the proportions of potential offenders and victims remain the same. Similarly, Groff (2007b, 2008), Liu & Eck (2008a), and Zhu & Wang (2021) explore how simulated crime is influenced by varying conceptualizations of theoretical constructs, including activity spaces and activity rhythms. In all cases, the model outcomes exhibit plausible patterns of crime, strengthening confidence in the theoretical constructs that inform them.

Building directly on these efforts, Birks et al. (2012, 2014) describe a model capable of simulating both property and interpersonal victimization. Their approach seeks to test whether mechanisms of movement, target selection, and offender learning derived from key theories of environmental criminology can generate multiple independent regularities of crime, namely spatial concentration, repeat victimization, and the characteristic journey-to-crime curve. In contrast to the models described above, this model uses an abstract street network that is randomly generated for each simulation to discount its potential effects on crime patterns. Results of simulated experiments demonstrate that as agents behave according to mechanisms proposed by theory, patterns of simulated crime become increasingly congruent with those observed in empirical studies.

In related research, Birks & Davies (2017) explore competing theoretical accounts through agent-based modeling. The authors explore the impact of street network structure on crime risk, assessing the veracity of the encounter (Jacobs 1961) and enclosure (Newman 1972) hypotheses often discussed within studies of crime and place and urban design. In their model, victim and offender agents undertake structured routine activities within a street network that is experimentally manipulated, and patterns of crime at the offender, street segment, and aggregate levels are observed. The experimental findings demonstrate a nonintuitive, nonlinear relationship between network permeability and crime risk and provide insights that may inform the interpretation of historically mixed empirical results in the field, in turn providing context-dependent support for both theories.

Collectively, these studies demonstrate that agent-based modeling is uniquely placed to formalize individual-level mechanisms put forward by theory and assess whether they are causally sufficient to explain, and thus generate, observed patterns of crime. The distinction between causal sufficiency and necessity is crucial here, as there may be multiple configurations of mechanisms capable of generating a given outcome. Consequently, agent-based modeling offers a viable means of distinguishing candidate explanations—that is, those that can generate a given outcome—supporting further empirical and theoretical development. By extension, agent-based modeling can also provide evidence toward the falsification of theoretical accounts that cannot generate the outcome they have been proposed to explain. The abovementioned studies also demonstrate that agent-based modeling offers a natural way to model the localized interactions among victims, offenders, and crime preventers that are often of greatest interest to environmental criminologists. Nonetheless, the universality of considering how proposed (but unobserved) individual-level behaviors translate into macro patterns of interest to criminologists means that agent-based modeling has been productively used to explore a diverse array of theories within the field. For example, ABMs have been employed to consider mechanisms of desistance that are of interest to developmental criminology (Cornelius et al. 2017), the dynamics of bystander effects (Gerritsen 2015, Groff & Badham 2020), and configurations of normative behaviors most capable of explaining extortion racket dynamics (Troitzsch 2016), among many others.

Prototyping Intervention Mechanisms

Beyond understanding the system of interest, a core goal of criminological research is to devise strategies that might allow intervention in that system to produce a desired outcome. ABMs are particularly well suited to exploring potential interventions because they enable multiple interventions to be compared with one another and with the outcomes from a baseline counterfactual (i.e., a society without any interventions) (Groff & Birks 2008, Groff et al. 2019). ABMs also offer a valuable tool for testing interventions when the data describing mechanisms and outcomes are poor. Perhaps most importantly, ABMs offer the opportunity to look "under the hood" at the mechanisms that produce outcome patterns so that the modeler can investigate why a simulated intervention does or does not produce expected outcomes. Even when such investigations are possible, simulated experiments can play a core role in the intervention theory development process because they can be conducted much more rapidly than field experiments, at a relatively minimal cost, and without navigating what may be complex logistical and ethical landscapes. In such scenarios, simulations may offer criminologists a method adjacent to in vitro experimentation, allowing researchers to explore numerous potential interventions to identify those that are most promising and for which the causal pathways to a desired outcome seem most robust—subsequently informing the design and evaluation of necessary but resource-intensive field trials. ABMs that seek to develop or evaluate interventions can take various forms, which in turn influence how the insights they generate should be applied.

To illustrate, one substantive focus of intervention ABMs is on strategies to disrupt criminal networks (Calderoni et al. 2021, Duxbury & Haynie 2019, Weisburd et al. 2022). These studies use agent-based modeling because data describing individuals and processes involved in organized crime are currently unavailable. In the case of terrorist groups, the small numbers of individuals recruited present challenges to achieving statistical power. In such cases, agent-based modeling offers another avenue to test a variety of interventions and observe dynamic interactions over time.

Duxbury & Haynie (2019) note that empirical research applying social network analysis has identified the security–efficiency trade-off in criminal networks and the tremendous variation among organizations in how they balance those constraints. However, these studies were limited to cross-sectional examinations of networks and, except for drugs, were unable to tailor interventions to the network. In response, Duxbury & Haynie (2019) developed an abstract, network-based ABM capable of network resiliency both to initial disruption and over time. Using this model, they explored whether the outcomes of implementing three potential disruption strategies differ according to the organizational characteristics of a network. The results of these in silico experiments demonstrate that the simulated criminal networks do not recover quickly from a disruption, that there are differences in whether organizations are security or efficiency oriented, and that broker targeting is the most damaging strategy both immediately and over the long term. Such insights can be obtained only in simulation settings, where networks of different forms can be constructed and repeatedly subjected to various interventions with the aim of understanding those that are most likely to be disruptive.

Two closely related ABMs examine different policy prescriptions for reducing recruitment into organized crime (Calderoni et al. 2021, Weisburd et al. 2022). Calderoni et al. (2021) compare a counterfactual of no intervention with four potential interventions. Two interventions focus on disrupting the organization by arresting leaders and individuals with important skills (facilitators). The other two disrupt socialization within the family and the peer network. The ABM is theoretically informed, incorporating both social environments and social influences (differential association and social learning theories) and individual influences (general theory of crime). It is also empirically grounded with data from Sicily and includes both individual and social network characteristics of agents. The baseline model generates realistic numbers of crimes and organized crime members. The results indicate that although all four policies significantly reduce the number of active members, only those interventions targeting leaders or facilitators and secondary socialization reduce the number of newly recruited members. Weisburd et al. (2022) study interventions focused on reducing recruitment to terrorist groups. They test the assumption that reducing radicalization will, by extension, reduce recruitment by implementing three commonly used interventions in an ABM. They find that both community workers at community centers and community-oriented policing reduce radicalization but have no significant effect on recruitment, whereas employment centers significantly reduce recruitment but have no effect on radicalization.

Earlier ABMs also examine how different crime prevention and policing strategies affect where and when offenders choose to commit crime. Malleson et al. (2010, 2012, 2013) develop a model of residential burglary and use it to estimate the impacts of potential built-environment manipulations and target hardening policies on spatial patterns of offending. To increase confidence in these insights, the authors validate the output of the burglary baseline model against empirical burglary data. The model includes agents with cognitive maps, namely an awareness of the local environment that evolves as the agents navigate the virtual city. The findings indicate that certain configurations of the built environment affect agents' cognitive maps such that a small number of houses are at a higher risk of burglary after implementation of an urban regeneration scheme and should be prioritized for target hardening by the police or local government.

Dray et al. (2008) evaluate the effectiveness of differing police tactics that seek to disrupt drug markets. Their ABM examines drug enforcement strategies by comparing standard patrol (random police movement), hot-spot policing (police patrol in specific hot grid cells), and problem-oriented policing (patrol in hot grid cells but also partnership with outreach workers). Their model combines an abstract landscape with a complex set of agents, including police officers, outreach workers, drug dealers, and wholesalers. In their model, problem-oriented policing strategies were most effective at disrupting street drug markets.

The models discussed in this section demonstrate the utility of ABMs for prototyping interventions. In each case, an ABM is used because process and/or outcome data were unavailable or of poor quality; multiple interventions can be tested against a counterfactual; and mechanisms underlying differences in outcome measures for each intervention can be identified and systematically explored.

Scaling Up Field Trials Through Simulation

Beyond exploring potential interventions, a question of particular interest to policymakers and practitioners is whether interventions shown to be effective in randomized experiments are likely to be effective when implemented as jurisdictional policies (Nagin & Sampson 2019). The limitations of randomized experiments for answering this question have received increasing attention in the academic literature (see, in particular, Nagin & Sampson 2019; see also Deaton & Cartwright 2018, Mowat et al. 2018). Because ABMs enable the exploration of multiple counterfactual worlds emerging from treatment regimes, in comparison to the status quo, they offer a clear contribution to the understanding of outcomes from taking a program to scale.

The area of hot-spot policing offers examples of how ABMs might support taking a program to scale. A significant body of literature using randomized controlled studies or related designs has found that sending police to small, high-crime hotspots reduces crime (Braga et al. 2019). Recognizing the limitations of empirical methods to answer whether jurisdiction-wide hot-spot policing would reduce overall crime in the city, three different ABMs examine the effect of police patrol strategies on jurisdiction-level crime. An early ABM created as an illustration of the method compares the effects of random patrol patterns with hot-spot policing and finds fewer overall burglaries in the hot-spot policing condition (Johnson 2009). A better-developed example creates four different police patrol scenarios—no police, random, low-intensity hot-spot policing, and high-intensity hot-spot policing—and examines their relative effects on street robbery (Weisburd et al. 2017). The findings indicate that in silico hot-spot police patrol significantly reduces robbery at the jurisdictional level. Using the same ABM, Wooditch (2021) examines whether using unallocated time in hot-spot patrol reduces robbery in larger areas. She finds that spending up to 30% of unallocated time on patrol in hotspots reduces crime even further both in the hot-spot areas and at the larger jurisdictional level.

The examples of ABMs covered in this section offer a sample of the diverse applications that exist. They also illustrate the continuum of model applications from largely abstract to increasingly complex. Regardless of the level of complexity, common practice involves the creation of a baseline model followed by comparisons between the outputs of the baseline model and stylistic or empirical values (i.e., validation). Once the baseline model has been validated, a small set of model characteristics is systematically manipulated and the changes in model outputs analyzed. Existing ABMs have barely scratched the surface of what is possible with intervention prototyping. ABMs have the potential to disentangle the effects of the intervention from other, related factors. As with empirical research, none of these modeling efforts are without flaws. Some are stronger, others are weaker. Similar to empirical studies, it is the accumulation of evidence across models that produces the most value.

WHAT IS INVOLVED IN BUILDING AN AGENT-BASED MODEL?

Building a robust ABM is a challenging endeavor, partly because the field is still relatively new but also because every ABM is unique. This uniqueness is both a blessing and a curse: It means that bespoke models can be created to explore a diverse range of applications, but it also means that the field lacks standardized techniques for evaluating model reliability or accuracy. That said, although no standard suite of methods for the development of ABMs exists, many modelers

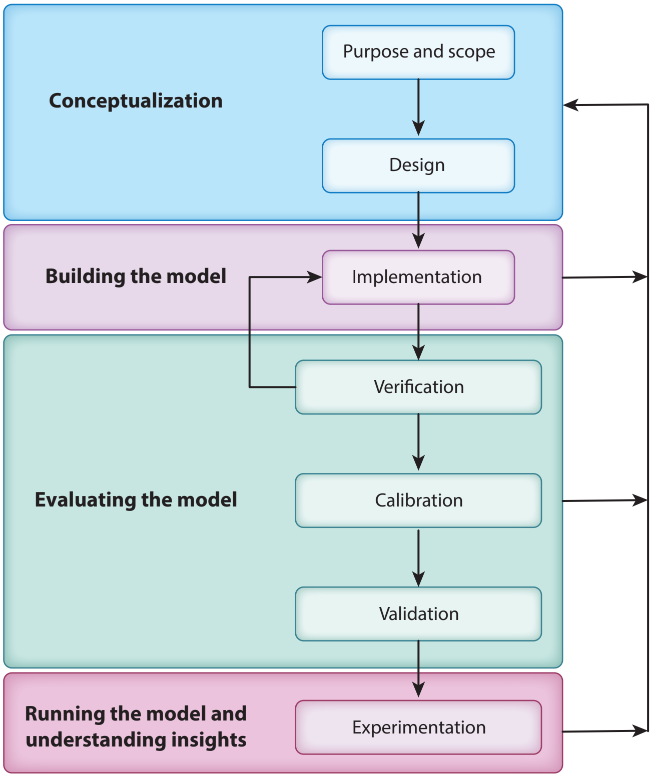

Figure 1

Key stages in agent-based model development.

follow a similar approach to model building, and these shared approaches foster the creation of more robust models. This section outlines key activities in the model development process (**Figure 1**). Note that the workflow is rarely linear, and the process of developing and experimenting with a model often leads to unanticipated outcomes that require reevaluation of the core model conceptualization. As discussed above, this is one of the most powerful aspects of agent-based modeling: The process of building a model can structure thinking about, and potentially reveal insights into, the underlying system one seeks to simulate.

Conceptualization

Before an ABM can be developed, it must be conceptualized. Conceptualization guides creation of the key elements of the model: the definition of agents, their interactions, and the environment. Depending on the type and purpose of the model, this phase can draw on insights from theory, previous empirical studies, and domain experts who are familiar with the system to be simulated. Grounding the conceptualization of the model in these diverse sources of evidence brings us closer to a model that is meaningful and applicable to real-world scenarios. In addition, both theory and domain expertise can play key roles in formulating hypotheses that the model can explore. Practically speaking, conceptualization begins by defining the model's purpose and scope. This process includes identifying the agents (e.g., types, number), their attributes (e.g., health, age, preferences, state), the ways that they interact (e.g., sending messages over a social network, interacting with nearby agents in a physical environment), their behaviors (e.g., compliance with social norms, offender perceptions of risk and reward, responses to changing environmental or social conditions), and the environment in which they will operate. During conceptualization, it is important to define the objectives and boundaries of the model and to formulate the questions that it will be designed to answer. It is easy to overcomplicate models by including a diverse array of behavioral factors that, although present in the real system, are unlikely to have a significant impact on system outcomes. For example, everyone needs to eat, but is eating an important activity that needs to be included in a model of interpersonal violence? The risk of including too many possible behaviors is that the model becomes as complicated, and as difficult to understand, as the real system.

Interestingly, the process of formally conceptualizing a model might result in the model being deemed unnecessary, or it might serve to highlight alternative research that needs to be conducted prior to model development. For example, when identifying important agent attributes—say, the deterrence effect exerted on offenders by a particular crime prevention measure—it might become clear that additional data must be collected first to adequately quantify this attribute. Similarly, the process of converting current understandings of a system into elements of a model might highlight aspects of the system of interest that are not well theorized or understood. This could necessitate further theoretical development, engagement with domain experts, or additional empirical research. Ultimately, a lack of sufficient insight to inform an initially conceptual model may motivate the development of a simpler model that can help explore the ambiguous elements of the system. In these ways, model conceptualization serves to further scientific investigation of the system prior to model development.

Development and Tools

Once a model has been adequately conceptualized, it can be developed. An early decision is the choice of modeling platform. In recent years, many alternatives have emerged, from fully graphical tools that require no previous programming expertise to detailed libraries that require advanced software development skills (for a straightforward, step-by-step example, see Birks & Malleson 2023). The choice of platform will ultimately depend on the preferences of the researcher and on the scope of the model. A platform that is suitable for developing instructive but parsimonious models may not be suitable for more realistic models that encompass fine-grained calibration data and numerous complex agents.

Regardless of the tools used, model testing is essential. The nature of agent-based modeling means that as small pieces of code are replicated to control the behavior of hundreds or even millions of agents, small errors can quickly lead to model outcomes that are effectively meaningless (but might not seem so). In addition, thorough sensitivity analyses where key model parameters are systematically manipulated and outcomes observed can be valuable for assessing whether such parameters are having their intended impacts (see the "cookbook" developed in Thiele et al. 2014). The process of checking that the ABM has been implemented correctly is often termed verification in the agent-based modeling literature.

Finally, we note that, unlike in other modeling techniques where the process of developing the model might mean simply converting a conceptual plan into computer code, the implementation of an ABM is often an iterative process (**Figure 1**). Early versions of the model may produce unexpected behavior, or the implementation of a particular feature might prove problematic for reasons that become clear only once implementation has begun. Here a more agile process that draws reflectively on insights gained through development of the model to contextualize theory, expertise, and empirical evidence previously deployed in model conceptualization often produces the most useful models. In addition, the process of experimentation and discussion over early model versions has value as a scientific exercise in its own right; finding that our initial assumptions lead to unexpected macro effects might generate interesting insights into the dynamics of the system itself or force us to reconsider those assumptions.

Evaluation

Typically, the utility of a model depends on the degree to which it is a sufficient representation of the system it seeks to emulate. A common pipeline that is used to assess this degree of representation consists of (*a*) verification, or ensuring the model implementation matches its design and has been programmed correctly; (*b*) calibration, or adjusting a model's parameters to replicate known data; and (*c*) validation, or the application of the model to a new context to evaluate its accuracy. Note, however, that validation is one of the most contentious aspects of the ABM process and is still hotly debated (An et al. 2021, Malleson et al. 2022, Manson et al. 2020). Although some authors argue for validation as a means of proving that the model can adequately reflect some real-world patterns (Berk 2008), others argue that validation of empirical models is impossible, largely application dependent, or even wholly unnecessary (Augusiak et al. 2014). In this review, rather than outlining the typical approach to calibration and validation—given the numerous tutorials that do so (e.g., Crooks et al. 2019, Fagiolo et al. 2007)—we briefly outline the main challenges to evaluation.

First, modelers must decide which elements of the system they seek to validate against. As discussed above, ABMs can produce many outcomes at varying scales, all of which may be compared with their real-world equivalents for the purposes of validation. These outcomes range from the specific behaviors or actions that individual agents take to the aggregate patterns of some outcome such as crime trends or patterns. Models in criminology often use aggregate characteristics of distributions instead of actual numbers for comparison (Groff et al. 2019). Two well-known characteristics of crime-related data are that crimes are usually spatially and temporally concentrated and that small numbers of repeat offenders and victims are responsible for and experience disproportionate levels of victimizations, respectively. Thus, any model, and by extension the theories upon which it is based, that produces a random or dispersed pattern of crime is unlikely to be considered valid. Comparing the empirical regularities or patterns found within data offers a defensible alternative to using count data (Birks et al. 2012, Grimm et al. 2005). A comprehensive literature from ecology on these ideas, termed pattern-oriented modeling (Grimm & Railsback 2005), continues to develop. Through this approach, confidence in the credibility of the model is increased by confirming that its behavior matches that of the underlying system across multiple dimensions and at multiple scales.

In addition to identifying the specific patterns to be compared, it is necessary to decide how to assess the similarity between the model outputs and the associated real-world observations. This similarity can be assessed qualitatively through inspection of so-called distributional equivalences or through formal quantitative measures, but again, these can vary significantly, from standard statistical tests to complex spatiotemporal algorithms. Even with quantitative measures, there is no set threshold error level that a model must meet; consequently, it is up to the modelers to assess whether the simulated data are close enough to the real-world data. An alternative to quantitative approaches, or one that can be used in conjunction with them, is termed face validation. Face validation is expressed as the informal process of evaluating a model by having relevant experts review its outputs to determine whether the model is operating in a manner that is consistent with their understanding of the underlying system (for example, see Gong et al. 2023 for recent research on effective face validation with police investigators and crime analysts).

A complicating factor in comparisons of simulated and real-world data arises as a result of the probabilistic nature of ABMs. As agents' decisions usually involve an element of randomness (e.g., agent A has a 40% chance of taking action X and a 60% chance of taking action Y), models must be executed multiple times so that we can be confident that an outcome is not a rogue result of a particularly unusual model evolution. Consequently, some averaging of model outputs often takes place before they can be compared with real data, but this averaging should not neglect the uncertainty in the model outputs. If the simulated data vary wildly from one model run to another, but the average happens to be similar to that of real-world data, can we still be confident that the model is reliably replicating the behavior of the underlying system? Maybe we can, but a thorough assessment of the uncertainty in the model (e.g., McCulloch et al. 2022) might need to be undertaken.

Where confidence in a model's validity increases, the implicit assumption may be that the mechanisms enacted within the model are an appropriate abstraction of those that operate in reality and, thus, that the causal pathways that generate simulated outcomes are analogous to those that operate in the real world. Here, as discussed above, agent-based modeling faces a significant challenge: equifinality. Equifinality refers to the observation that in open systems there may be any number of initial conditions and causal pathways capable of generating the same system outcome (von Bertalanffy 1950). Practically speaking, this means that as models generate increasing numbers of independent real-world patterns, confidence in their validity grows; however, authors should always be careful with the inferences they draw from model outcomes alone.

Experimentation

The final stage involves using the model to conduct experiments that aim to answer research questions of interest. Although model uses might vary dramatically from one study to another, experiments within criminology typically mirror the types discussed in the previous sections. Broadly speaking, these can be divided into two categories. The first relates to models that seek to explore, understand, and test propositions (and related theories) regarding complex criminological systems. As agent-based modeling offers a natural means of exploring causal processes proposed by criminological theory, it follows that many research studies conduct experiments comparing the outcomes generated by various agents and behavioral rules derived from theory. The second use involves scenario analyses. Here, ABMs are used to simulate "what-if" scenarios, which estimate the possible and plausible impacts of changes to the system of interest—be they population changes, the implementation of specific crime reduction tactics, or the introduction of new policies or regulations. By simulating these scenarios, researchers hope to gain insight into how the system might respond under different conditions and make more informed decisions about potential courses of action.

In both cases, models can span varying levels of complexity—from the very abstract to the highly realistic—in turn dictating the criteria against which models should be validated and the types of insights that can be drawn from their outputs. Ultimately, the success of an experiment depends on the extent to which it can answer the research question. In some cases, a thorough quantitative evaluation will be the only means by which researchers can be confident that their results have value, but in other cases, it may be sufficient to simply highlight outcomes that are unexpected or counterintuitive.

KEY CHALLENGES

Although agent-based modeling may provide unique complementary insights into criminological phenomena, much of its potential to make significant contributions to the field remains unrealized, for various reasons. Many of the challenges facing agent-based modeling stem from its use in examining complex systems, which is a fundamentally difficult endeavor for any method to tackle

(Batty & Torrens 2005). Some challenges are intrinsic to agent-based modeling, such as defining appropriate model complexity or communicating model insights and uncertainty. Others are faced by other disciplines but are especially challenging in criminology, such as the relative paucity of empirical data to calibrate and validate models of crime and criminality or the opacity of the human decision-making processes that are at the heart of observed criminal behavior. This section elaborates on several of these key challenges and discusses how they have prevented agent-based modeling from becoming as widely adopted in criminology as it might be.

Perhaps the most oft-discussed challenge associated with agent-based modeling relates to the observation that models are simply collections of assumptions. Of course, this observation applies not only to ABMs but also to many other forms of model-based abstraction used within the sciences. Importantly, then, the utility of a given ABM is a function of both the context in which the model is used and the collective validity of the assumptions made in developing it. For instance, if a model is used to anticipate possible outcomes that result from competing crime reduction interventions, the utility of simulated scenarios in informing real-world action is directly tied to the collective validity of the assumptions that underpin the model. In this case, the model is used to anticipate the downstream consequences of those assumptions; thus, the plausibility of the possible futures it generates is directly related to their validity. Conversely, as discussed above, models are often used to test the validity of existing assumptions—be they theoretically or empirically informed. In this case, the notion that the model is a collection of assumptions is its primary purpose, such that it provides a framework to organize these assumptions and explore their collective validity in generating known outcomes.

Nevertheless, both of these cases are made more challenging within criminology by the relative scarcity of reliable data or robust theories through which to specify model assumptions or validate model outcomes. Where models are used for scenario analysis, there is often insufficient empirical evidence to validate initial assumptions to the degree that would be desirable. When models are used to test a theory, criminological theories are often insufficiently specified to enable direct implementation within the model. That said, we note that these challenges are not unique to criminology but rather are common in agent-based modeling within the social sciences generally (Antosz et al. 2023). Moreover, such challenges are a primary rationale for building models in scenarios where empirical and theoretical efforts alone are insufficient to unravel complex social systems. Consequently, a core responsibility of the modeler is to think critically about the goal of constructing a model, the veracity of evidence that underpins it, and the types of inferences that can and cannot be reliably made from its output.

An important related topic is model complexity—that is, what should be included within a model and at what level of realism. Just as the map is a simplified representation of a complex place, an ABM is a simplified representation of a complex social or physical phenomenon. But finding the balance between realistic empirically detailed models and more abstract theoretically grounded models can be challenging. Current discourse in agent-based modeling contrasts two approaches. The first favors parsimony, proposing that models should be as simple as possible and become more complex only when necessary (O'Sullivan et al. 2012). Parsimonious models offer several advantages. They are easier to understand; they have fewer parameters to operationalize; they are less susceptible to the reproduction of errors that can occur when several parameters are interacting (O'Sullivan et al. 2012); and the simpler the model is, the easier it will be to interpret the interactions and how those interactions affect the results. The second approach argues that although this desire for parsimony undoubtedly offers practical advantages in model development and interpretation, it does not necessarily confer advantages in terms of model validity and utility. Put simply, simpler models are not inherently better representations of the system of interest (Edmonds 2007). Proponents of this approach suggest that models should seek to describe the target system in the most direct way possible while incorporating all available evidence and should defer to simplification only if and when the model and evidence provide sufficient justification (Edmonds & Moss 2004).

Although these challenges have no simple solution, a key factor in increasing the utility derived from agent-based modeling is to insist that researchers be transparent in documenting the decisions made in developing a model. Within (and beyond) criminology, too many papers do not adequately describe these assumptions or, more generally, the models they utilize (Groff et al. 2019). Simply cataloging agents, behaviors, and relationships is a first step, but the assumptions that underlie the model and their rationale—be they empirically, theoretically, or intuitively informed—must be made explicit (Groff et al. 2019, Townsley & Birks 2008). Additionally, the development of standardized practices for model building would help guide modelers through these decisions (Groff et al. 2019, O'Sullivan 2004).

In the past 20 years, these issues have become widely recognized in all fields applying agent-based modeling. One proposed solution is to develop protocols for describing, analyzing, and reporting results from ABMs. Such standards have the dual benefit of clearly communicating what the modeler did (including the design, sensitivity testing, and analysis of the model) and making it easier for subsequent scholars to accurately replicate the model. Protocols such as ODD (overview, design concepts, and details) and ODD + D (ODD plus human decision-making) represent early options (Müller et al. 2013). That said, the comprehensiveness of such frameworks means that they are time consuming to complete, and models often do not fit all framework categories well. They can also increase article length considerably (although this can be avoided by taking advantage of the supplemental materials that many journals permit). Despite these difficulties, simply attempting to use standardized protocols improves the documentation of models.

A related challenge concerns the consistent lack of model replication, which has hampered the evolution of trust in agent-based modeling. Replication is a type of validation that tests the claims of prior research. It involves using the same design as a prior study and determining whether similar results can be obtained. Consistent outcomes suggest that an original study's findings are likely to hold for other populations and other places, allowing research to move from the specific to the generalizable. Replication within agent-based modeling can take several forms. One is the reimplementation of the same model through the use of a different software platform. A second form uses the same model but different empirical calibration data. A third consists of theoretical replications that involve recreating the model using the same theory or set of theories to model the same outcome (Liu & Eck 2008a). Similar results across different operationalizations support the soundness of assumptions and, in turn, the evidence from which they are derived. Differences in results can indicate places of sensitivity with regard to the operationalization of assumptions or unintended disparities in implementation. Either way, model insights that cannot be reproduced through replication should be subject to the same skepticism as empirical studies that cannot be reproduced. Traditionally, obtaining funding for and publishing empirical replication studies have been more difficult than testing new ideas and, thus, are less attractive to researchers. To a lesser extent, this bias also exists in agent-based modeling since academic career progression, irrespective of method, often favors the development of novel ideas. As a result, most agent-based modeling studies present new models, rather than incrementally increasing confidence in existing models and applying them to answer multiple research questions, a situation that has been lamented both within criminology (Birks & Townsley 2013) and further afield (O'Sullivan et al. 2016).

Furthermore, there is often a lack of detailed and accurate data to inform criminological ABMs to the degree that would be desirable. In many cases, the individual-level data necessary for ABMs have simply not been produced by extant empirical research. These challenges relate to the general assumptions that underlie models and the specific requirements for initial simulation conditions.

For example, simulations of human movement, and the criminological theories from which they are informed, contend that anchor nodes play a key role in the emergence of criminal opportunities (Brantingham & Brantingham 1993). However, until recently there was little empirical research to guide the number of anchor nodes assigned to agents in models. Similarly, only a small (nonrandom) selection of offenders is known; thus, calibration of the degree to which characteristics such as motivation vary between offenders or, more specifically, the spatial distribution of offender residences within a particular locality is largely unknown. Although this lack of data does not preclude the use of agent-based modeling in these contexts (indeed, it may offer a unique means of exploring the plausibility of various hypothetical configurations), it does constrain what model insights can be effectively used for.

Relatedly, official data describing crime incidents have well-known and significant deficiencies that result in their capturing only a portion of the actual crime that occurs (Maguire 2007). Although this is a known problem that similarly afflicts empirical research, it has clear ramifications for the validation of ABMs and is especially problematic for models in which crime is the outcome variable. Recorded crime data contain only incidents that were perceived by the victim as a crime, reported to the police, and recognized by the police officer as a crime worthy of a crime report. This problem is especially acute for agent-based modeling because the mechanisms represented in a model may produce crime patterns that reflect actual crime patterns but be inconsistent with official crime data (Eck 2007). For agent-based modeling to be validated with recorded crime data, one would also need to model the abovementioned selection and attrition processes (Eck 2007). A direct corollary of this challenge is that as the criminological evidence base becomes both more diverse and robust, the validity of models and the utility that can be derived from them will increase.

The final challenge facing agent-based modeling relates to the nature of model insights, particularly when compared with more traditional modeling approaches. Although ABMs may be considered a natural approach for modeling human systems, this clarity is offset by the fact that ABMs do not always provide simple answers to a given question. Even the simplest ABM will produce vast quantities of outcome data at various scales and aggregations. Revisiting our crime hot-spots example from the introduction, our ABM of crime pattern formation would enable examination of diverse constructs of interest, including the internal decision-making of agents, target evaluations, and the influence of route choices on the targets encountered, among others. The stochastic nature of ABMs also means that these measures may require analysis over hundreds or thousands of simulation runs associated with various sets of initial conditions. The very nature of complex systems means that these arrays of outcomes may often reveal, for example, nonlinear outcomes, tipping points, and complex context-dependent relationships. The challenge is in communicating these system properties in contexts where simple outcome measures may be most desirable.

FUTURE DIRECTIONS

In this final section, we highlight three core areas where we believe future agent-based modeling research will support criminology.

The Rise of Novel Data to Inform Agent-Based Models

Many of the challenges associated with the use of agent-based modeling in criminology stem from a lack of appropriate evidence for how the underlying criminological systems behave. This is particularly true when it comes to encoding human behavioral rules that drive agents; there are many theories in criminology, but even when theories have the support of a strong evidence base, they may be too weakly specified to form concrete behavioral rules in a model. The emergence of big data (Mayer-Schönberger & Cukier 2013) and the rise in related data science techniques may provide new opportunities for observing and interpreting actions and behaviors, often at the level of the individual. For example, recent advances in mobility data have produced empirical descriptions of activity spaces that can be used to calibrate the characteristics of activity spaces in agent-based modeling (Candia et al. 2008, Song et al. 2010). The use of social media data might also provide a supplementary source of information to complement police-recorded crime data as a means of better understanding crime patterns and enable new tests of related theories (Williams et al. 2017). More broadly, this combination of novel data and ABMs might foster a new generation of societal digital twins (Birks et al. 2020) that could provide a new avenue for criminologists and other social scientists to simulate and test interventions in a realistic, yet controlled manner.

At the same time, we do not believe that this new "age of data" will dictate that theory has a less important role in the models criminologists build. Indeed, we believe that if ABMs are to be of use, they must combine solid theoretical underpinnings with diverse sources of data. Such a combination would allow their use to increase our understanding of how potential futures might develop (Elsenbroich & Badham 2023). Thus, in contrast to other methods, such as increasingly popular forms of machine learning, ABMs provide new opportunities to utilize novel detailed data sets to build, calibrate, and validate data-informed but theoretically grounded and causally explicit models of the complex systems of interest to criminologists.

Better Linking of Empirical and Modeling Research

A tighter coupling between empirical and modeling research is one of the keys to developing policies that can address the complex problems society faces. We envisage two main ways that this integration can be achieved. The first is through the addition of more empirically grounded agent behavior to models. Both abstract and more realistic models benefit from stronger empirical evidence to support agent behavior and decision-making. As discussed above, ever-increasing and novel data sources are making it possible to empirically measure characteristics of human behavior in time and space that were previously inaccessible. These data have the potential to inform a host of new ABMs of interest to criminologists.

Second, there is considerable potential for the addition of agent-based modeling into the criminological methodological cycle, which will be most powerful when in silico and empirical studies are part of a coordinated program. Groff and colleagues (Groff & Birks 2008, Groff et al. 2019) propose a methodological cycle that uses in silico experiments to interrogate the processes and mechanisms underlying a potential treatment. The knowledge gained from model building and experimentation within an ABM is employed to strengthen empirical studies using secondary data or quasi-experimental designs. Results from the empirical phase are then used to inform the ABM. Further exploration of the enhanced ABM, in turn, provides a better understanding of how the treatment might work when deployed in a field experiment. These additional in silico tests can be undertaken quickly and for a fraction of the cost of additional field experiments. In this way, an ABM-empirical research program can maximize our understanding of why a treatment works, and under what conditions, for less time and money than a traditional program.

Although implementation of such a program remains aspirational, a future where ABMs are used to instantiate theoretically grounded, data-informed simulations; test interventions; and understand mechanisms as part of a methodological cycle that includes both simulated and empirical investigations could rapidly increase our ability to develop complex solutions to complex societal problems. The transformation needed to achieve this reorientation in research paradigms is analogous to the recent methodological evolution toward mixed methods to develop a more nuanced understanding of complex issues.

Models to Bridge Silos

One thing that is clear from the application of ABMs in various fields, including criminology, is that the complexity of the modeling enterprise and the problems that agent-based modeling seeks to address often benefit from coproductive, multidisciplinary, and multisectoral ways of working. This form of cross-silo cooperation can manifest in multiple ways. Scholars of criminology are rarely well versed in software development. Similarly, those with expertise in software development often lack the domain knowledge necessary to understand the gaps in knowledge that could be addressed through the application of agent-based modeling. Teams enable criminologists to contribute their knowledge of theory and extant empirical literature. Programmers and/or modelers typically found in other disciplines can contribute their knowledge of software development and testing as well as simulation modeling. Although individuals with both sets of skills exist (several of whom are discussed above), there is still much to gain from productive collaboration between criminologists, data scientists, and, in many cases, experts from other disciplines. Notwithstanding the challenges associated with shared language, such partnerships leverage the strengths of team members and have the advantage of bringing together multiple perspectives (Birks & Townsley 2013, Groff et al. 2019). Relatedly, despite their computational nature, the development of robust ABM necessitates methodological integration. ABMs require insights and evidence that are both qualitative and quantitative in nature. Ethnographic studies, accounts of those with lived experience, and researcher field expertise can all be used to inform substantive elements of ABMs. Such insights can be effectively integrated with individual- and census-level statistics, odds ratios derived from previous empirical studies, and spatial data regarding real-world environments to create rich multimodal models of complex systems.

Similarly, the nature of some ABMs that are in use in the field of criminology, particularly those considering intervention development or prototyping, means that expert knowledge of the system of interest is often essential for building models that will be of practical utility outside of academia. For example, one of the authors of this review has expended considerable effort working closely with police and related practitioners to build ABMs of police supply-and-demand dynamics (Laufs et al. 2021) to produce models capable of exploring possible futures related to changes in police demand and how resources might be organized to respond. Models of this nature must be coproduced with practitioners. Doing so ensures that they best represent the system of interest, draw on the powerful and diverse insights that exist within organizations that operate within that system, can explore questions of interest to those organizations, and are cognizant of how insights derived through the model will be used to inform decision-making. Although such engagement with agent-based modeling within criminology is still in its infancy, powerful recent examples from public health have informed responses to the global COVID-19 pandemic (Badham et al. 2021, Blakely et al. 2021, Thompson et al. 2021).

One of the oft-cited strengths of ABMs is their relative accessibility versus that of other complex system modeling approaches (Gilbert 2008). That is, specifying a complex system in terms of individuals and the decisions they make is considerably less opaque than abstracting such relationships to a set of differential equations. At the same time, Elsenbroich & Badham (2023) rightly point out that models can be overly alluring to nonmodelers, such that they are interpreted as facsimiles of reality when they are not. We agree that appropriate caution is warranted in the interpretation of ABMs. Yet, we also believe that models provide something concrete around which individuals from diverse perspectives can converge. Models not only organize proposed explanations but also structure and productively constrain discussion of the real-world system, bringing together researchers from different backgrounds to ask key questions: What assumptions underlie our current thinking about the system? What elements of the system are important? What elements of the system do we know a lot about, and what do we know little about? What evidence is available to provide these insights? What are its strengths and weaknesses? What new evidence is necessary? These discussions are essential if we are to increase our understanding of the complex systems that make up the world around us.

DISCLOSURE STATEMENT

The authors are not aware of any affiliations, memberships, funding, or financial holdings that might be perceived as affecting the objectivity of this review.

ACKNOWLEDGMENTS

Daniel Birks is partly supported by the Economic and Social Research Council (grant number: ES/W002248/1).

LITERATURE CITED

An L, Grimm V, Sullivan A, Turner BL II, Malleson N, et al. 2021. Challenges, tasks, and opportunities in modeling agent-based complex systems. *Ecol. Model.* 457:109685

Antosz P, Birks DJ, Edmonds B, Heppenstall A, Meyer R, et al. 2023. What do you want theory for? A pragmatic analysis of the roles of "theory" in agent-based modelling. *Environ. Model. Softw.* 168:105802

Augusiak J, van den Brink PJ, Grimm V. 2014. Merging validation and evaluation of ecological models to 'evaludation': a review of terminology and a practical approach. *Ecol. Model.* 280:117–28

Badham J, Barbrook-Johnson P, Caiado C, Castellani B. 2021. Justified stories with agent-based modelling for local COVID-19 planning. *J. Artif. Soc. Soc. Simul.* 24(1):8

Batty M, Torrens PM. 2005. Modelling and prediction in a complex world. *Futures* 37:745–66

Berk R. 2008. How you can tell if the simulations in computational criminology are any good. *J. Exp. Criminol.* 4(3):289–308

Birks DJ, Davies T. 2017. Street network structure and crime risk: an agent-based investigation of the encounter and enclosure hypotheses. *Criminology* 55(4):900–37

Birks DJ, Heppenstall A, Malleson N. 2020. Towards the development of societal twins. In *Proceedings of the 24th European Conference on Artificial Intelligence* (*ECAI 2020*), pp. 2883–84. Amsterdam: IOS

Birks DJ, Malleson N. 2023. Agent-based simulation models. In *Understanding Crime and Place: A Methods Handbook*, ed. ER Groff, CP Haberman, p. 304. Philadelphia: Temple Univ. Press

Birks DJ, Townsley M. 2013. Making crime simulation richer: 8 steps on the path to criminology's computational enlightenment. In *Wis en Waarachtig: Liber Amicorum voor Henk Elffers*, ed. S Ruiter, W Bernasco, W Huisman, G Bruinsma, pp. 477–90. Amsterdam: NSCR/Vrije Univ. Amsterdam

Birks DJ, Townsley M, Stewart A. 2012. Generative explanations of crime: using simulation to test criminological theory. *Criminology* 50(1):221–54

Birks DJ, Townsley M, Stewart A. 2014. Emergent regularities of interpersonal victimization: an agent-based investigation. *J. Res. Crime Delinquency* 51(1):119–40

Blakely T, Thompson J, Bablani L, Andersen P, Ouakrim DA, et al. 2021. Determining the optimal COVID-19 policy response using agent-based modelling linked to health and cost modelling: case study for Victoria, Australia. medRxiv 2021.01.11.21249630. **https://doi.org/10.1101/2021.01.11.21249630**

Braga AA, Turchan BS, Papachristos AV, Hureau DM. 2019. Hot spots policing and crime reduction: an update of an ongoing systematic review and meta-analysis. *J. Exp. Criminol.* 15:289–311

Brantingham PL, Brantingham PJ. 1993. Nodes, paths and edges: considerations on the complexity of crime and the physical environment. *J. Environ. Psychol.* 13:3–28

Bonabeau E. 2002. Agent-based modeling: methods and techniques for simulating human systems. *PNAS* 99(3):7280–87

Bookstaber R. 2017. Agent-based models for financial crises. *Annu. Rev. Financ. Econ.* 9:85–100

Calderoni F, Campedelli GM, Szekely A, Paolucci M, Andrighetto G. 2021. Recruitment into organized crime: an agent-based approach testing the impact of different policies. *J. Quant. Criminol.* 38:197–237

Candia J, Gonzalez MC, Wang P, Schoenharl T, Madey G, Barabasi A-L. 2008. Uncovering individual and collective human dynamics from mobile phone records. *J. Phys. A* 41:224015

Cohen LE, Felson M. 1979. Social change and crime rate trends: a routine activity approach. *Am. Sociol. Rev.* 44:588–608

Cornelius CV, Lynch CJ, Gore R. 2017. Aging out of crime: exploring the relationship between age and crime with agent based modeling. In *Proceedings of the Agent-Directed Simulation Symposium* (*ADS '17*), Art. 3. New York: ACM

Crooks A, Malleson N, Manley E, Heppenstall A. 2019. *Agent-Based Modelling and Geographical Information Systems: A Practical Primer*. Thousand Oaks, CA: SAGE

Deaton A, Cartwright N. 2018. Understanding and misunderstanding randomized controlled trials. *Soc. Sci. Med.* 210:2–21

Dray A, Mazerolle L, Perez P, Ritter A. 2008. Drug law enforcement in an agent-based model: simulating the disruption to street-level drug markets. In *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*, ed. JE Eck, L Liu, pp. 352–71. Hershey, PA: IGI Global

Duxbury SW, Haynie DL. 2019. Criminal network security: an agent-based approach to evaluating network resilience. *Criminology* 57(2):314–42

Eck J. 2007. *The devil's computer: a theory of crime patterns*. Paper presented at the Environmental Criminology and Crime Analysis Annual Meeting, Chilliwack, BC, July 26–28

Edmonds B. 2007. Simplicity is *not* truth-indicative. In *Worldviews, Science and Us: Philosophy and Complexity*, ed. C Gershenson, D Aerts, B Edmonds, pp. 65–80. Liverpool, UK: Univ. Liverpool Press

Edmonds B, Moss S. 2004. From KISS to KIDS—an 'anti-simplistic' modelling approach. In *International Workshop on Multi-Agent Systems and Agent-Based Simulation*, pp. 130–44. Berlin: Springer

Elsenbroich C, Badham J. 2023. Negotiating a future that is not like the past. *Int. J. Soc. Res. Methodol.* 26(2):207–13

Epstein J. 2006. *Generative Social Science: Studies in Agent-Based Computational Modeling*. Princeton, NJ: Princeton Univ. Press

Epstein JM. 2008. Why model? *J. Artif. Soc. Soc. Simul.* 11(4):12–17

Epstein J, Axtell R. 1996. *Growing Artificial Societies: Social Science from the Bottom Up*. Cambridge, MA: MIT Press

Fagiolo G, Moneta A, Windrum P. 2007. A critical guide to empirical validation of agent-based models in economics: methodologies, procedures, and open problems. *Comput. Econ.* 30:195–226

Gerritsen C. 2015. Agent-based modelling as a research tool for criminological research. *Crime Sci*. 4(1):2

Gerritsen C, Elffers H, eds. 2020. *Agent-Based Modelling for Criminological Theory Testing and Development*. London: Routledge

Gilbert N. 2008. *Agent-Based Models*. London: Sage

Gong Y, Gu F, Dai M. 2023. A methodology of face validation with domain experts for agent-based crime risk prediction. In *IEEE 3rd International Conference on Software Engineering and Artificial Intelligence* (*SEAI*), pp. 96–103. Piscataway, NJ: IEEE

Grimm V, Railsback SF. 2005. *Individual-Based Modeling and Ecology*. Princeton, NJ: Princeton Univ. Press

Grimm V, Revilla E, Berger U, Jeltsch F, Mooij WM, et al. 2005. Pattern-oriented modeling of agent-based complex systems: lessons from ecology. *Science* 310(5750):987–91

Groff ER. 2007a. Simulation for theory testing and experimentation: an example using routine activity theory and street robbery. *J. Quant. Criminol.* 23:75–103

Groff ER. 2007b. 'Situating' simulation to model human spatio-temporal interactions: an example using crime events. *Trans. GIS* 11(4):507–30

Groff ER. 2008. Adding the temporal and spatial aspects of routine activities: a further test of routine activity theory. *Secur. J.* 21:95–116

Groff ER, Badham J. 2020. Examining guardianship against theft. See Gerritsen & Elffers 2020, pp. 71–103

Groff ER, Birks DJ. 2008. Simulating crime prevention strategies: a look at the possibilities. *Policing* 2(2):175–84

Groff ER, Johnson SD, Thornton A. 2019. State of the art in agent-based modeling of urban crime: an overview. *J. Quant. Criminol.* 35:155–93

Groff ER, Mazerolle L. 2008. Simulated experiments and their potential role in criminology and criminal justice. *J. Exp. Criminol.* 4:187–93
Hedström P, Ylikoski P. 2010. Causal mechanisms in the social sciences. *Annu. Rev. Sociol.* 36:49–67
Holland P. 1986. Statistics and causal inference. *J. Am. Stat. Assoc.* 81(396):945–60
Jacobs J. 1961. *The Death and Life of Great American Cities*. New York: Vintage Books
Johnson SD. 2009. Potential uses of computational methods in the evaluation of crime reduction activity. In *Evaluating Crime Prevention: Crime Prevention Studies*, ed. J Knuttson, N Tilley, pp. 175–217. Monsey, NY: Crim. Just.
Laufs J, Bowers K, Birks D, Johnson SD. 2021. Understanding the concept of 'demand' in policing: a scoping review and resulting implications for demand management. *Polic. Soc.* 31(8):895–918
Liu L, Eck J. 2008a. Varieties of artificial crime analysis: purpose, structure, and evidence in crime simulations. In *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*, ed. L Liu, J Eck, pp. 413–32. Hershey, PA: IGI Glob.
Liu L, Eck J. 2008b. Contrasting simulated and empirical experiments in crime prevention. *J. Exp. Criminol.* 4:195–213
Luke DA, Stamatakis KA. 2012. Systems science methods in public health: dynamics, networks, and agents. *Annu. Rev. Public Health* 33:357–76
Macy WM, Willer R. 2002. From factors to actors: computational sociology and agent-based modeling. *Annu. Rev. Sociol.* 28:143–66
Maguire M. 2007. Crime data and statistics. In *Oxford Handbook of Criminology*, ed. M Maguire, R Morgan, R Reiner, pp. 241–301. Oxford, UK: Oxford Univ. Press
Malleson N, Birkin B, Birks DJ, Ge J, Heppenstall A, et al. 2022. *Agent-Based Modelling for Urban Analytics: State of the Art and Challenges*, ed. SV Albrecht, M Woolridge. *AI Commun*. 35(4):393–406
Malleson N, Heppenstall A, See L. 2010. Crime reduction through simulation: an agent-based model of burglary. *Comput. Environ. Urban Syst.* 34(3):236–50
Malleson N, Heppenstall A, See L, Evans A. 2013. Using an agent-based crime simulation to predict the effects of urban regeneration on individual household burglary risk. *Environ. Plan. B* 40:405–26
Malleson N, See L, Evans A, Heppenstall A. 2012. Implementing comprehensive offender behaviour in a realistic agent-based model of burglary. *Simulation* 88(1):50–71
Manson S, An L, Clarke KC, Heppenstall A, Koch J, et al. 2020. Methodological issues of spatial agent-based models. *J. Artif. Soc. Soc. Simul.* 23(1):3
Marchi SD, Page SE. 2014. Agent-based models. *Annu. Rev. Political Sci.* 17:1–20
Mayer-Schönberger V, Cukier K. 2013. *Big Data: A Revolution That Will Transform How We Live, Work and Think*. London: Murray
McCulloch J, Ge J, Ward JA, Heppenstall A, Polhill JG, Malleson N. 2022. Calibrating agent-based models using uncertainty quantification methods. *J. Artif. Soc. Soc. Simul.* 25(2). **https://doi.org/10.18564/jasss.4791**
Mowat R, Subramanian SV, Kawachi I. 2018. Randomized controlled trials and evidence-based policy: a multidisciplinary dialogue. *Soc. Sci. Med.* 210:1–90
Müller B, Bohn F, Dreßler G, Groeneveld J, Klassert C, et al. 2013. Describing human decisions in agent-based models—ODD? D, an extension of the ODD protocol. *Environ. Model. Softw.* 48:37–48
Nagin DS, Sampson RJ. 2019. The real gold standard: measuring counterfactual worlds that matter most to social science and policy. *Annu. Rev. Criminol.* 2:123–45
Newman O. 1972. *Defensible Space: Crime Prevention Through Environmental Design*. New York: Macmillan
O'Sullivan D. 2004. Complexity science and human geography. *Trans. Inst. Br. Geogr.* 29(3):282–95
O'Sullivan D, Evans T, Manson S, Metcalf S, Ligmann-Zielinska A, Bone C. 2016. Strategic directions for agent-based modeling: avoiding the YAAWN syndrome. *J. Land Use Sci.* 11(2):177–87
O'Sullivan D, Millington J, Perry G, Wainwright J. 2012. Agent-based models–because they're worth it? In *Agent-Based Models of Geographical Systems*, ed. AJ Heppenstall, AT Crooks, LM See, M Batty, pp. 109–23. Dordrecht, Neth.: Springer
Song C, Konen T, Wang P, Barabási A-L. 2010. Modelling the scaling properties of human mobility. *Nat. Phys.* 6(10):818–23

Thiele JC, Kurth W, Grimm V. 2014. Facilitating parameter estimation and sensitivity analysis of agent-based models: a cookbook using NetLogo and "R." *J. Artif. Soc. Soc. Simul.* 17(3):11

Thompson J, McClure R, Scott N, Hellard M, Abeysuriya R, et al. 2021. A framework for communicating the utility of models when facing tough decisions in public health. *Health Res. Policy Syst.* 20:107

Townsley M, Birks DJ. 2008. Building better crime simulations: systematic replication and the introduction of incremental complexity. *J. Exp. Criminol.* 4:309–33

Tracy M, Cerdá M, Keyes KM. 2018. Agent-based modeling in public health: current applications and future directions. *Annu. Rev. Public Health* 39:77–94

Troitzsch KG. 2016. Extortion rackets: an event-oriented model of interventions. In *Social Dimensions of Organised Crime: Modelling the Dynamics of Extortion Rackets*, ed. C Elsenbroich, D Anzola, N Gilbert, pp. 117–31. Cham, Switz.: Springer

von Bertalanffy L. 1950. The theory of open systems in physics and biology. *Science* 111(2872):23–29

Wang X, Liu L, Eck JE. 2008. Crime simulation using GIS and artificial intelligent agents. In *Artificial Crime Analysis Systems: Using Computer Simulations and Geographic Information Systems*, ed. JE Eck, L Liu, pp. 209–24. Hershey, PA: IGI Global

Weisburd D, Braga AA, Groff ER, Wooditch A. 2017. Can hot spots policing reduce crime in urban areas? An agent-based simulation. *Criminology* 55(1):137–73

Weisburd D, Wolfowicz M, Hassisi B, Paolucci M, Andrighetto G. 2022. What is the best approach for preventing recruitment to terrorism? Findings from ABM experiments in social and situational prevention. *Criminol. Public Policy* 21(2):461–85

Williams ML, Burnap P, Sloan L. 2017. Crime sensing with big data: the affordances and limitations of using open-source communications to estimate crime patterns. *Br. J. Criminol.* 57(2):320–40

Wooditch A. 2021. The benefits of patrol officers using unallocated time for everyday crime prevention. *J. Quant. Criminol.* 39:161–85

Zhu H, Wang F. 2021. An agent-based model for simulating urban crime with improved daily routines. *Comput. Environ. Urban Syst.* 89:101680