M0 - Agent-Based Modelling as a Field

Gidi Brandes

Defining Agent-Based Modelling: The Generative Standard of Explanation

Agent-Based Modelling (ABM) is a computational methods in which a system is represented as a population of autonomous agents, each governed by its own local rules, interacting with each other and with an environment, such that system/macro-level patterns arise from the accumulation of individual/micro-interactions rather then being specified directly. The defining feature of ABM is not the use of a computer as there are plenty of computational methods that exist that are not agent-based. The feature of ABM is the direction of explanation. Classic approaches like regression models or the game-theoretic equilibrium specify a macro-level relationship or a macro-level outcome directly and ask whether the data are consistent with it. An ABM specifies only the micro-level behavioural rules and lets the macro pattern emerge computationally.

In his famous article “Agent-Based Computational Models and Generative Social Science”, Joshua Epstein (1999) gave this move its canonical statement. He argued that a macro-level explanandum is only explained, in the generative sense, once you have demonstrated that a population of heterogeneous, locally interacting agents, following specified rules, is computationally sufficient to grow that pattern from the bottom up. This claim is stronger and different than statistical correlation. Statistical correlation shows association but not the sufficiency of a mechanism. The ABM-claim is also different than an analytical equilibrium proof which shows that a pattern is a fixed point of some idealized rational process. It doesn’t argue that a population of realistically limited actors would actually arrive there. Epstein calls this the “generativist” or “constructive” standard. It is worth to think about the fact that this a a genuinely falsifiable standard: if you cannot grow the pattern from plausible micro-rules, the generative explanation fails, regardless of how intuitively compelling the micro-rules sounded on paper.

This generative logic can be illustrated with the “Coleman boat”. This is a diagram used to explain Coleman’s argument in Foundations of Social Theory (1990). Coleman argued that a macro-level condition produces a micro-level effect on individual, which in turn produces, through the aggregation of many individuals’ actions, a new macro-level outcome. Coleman’s methodological-individualist argument was that most macro-to-macro social explanations smuggle in an unexamined micro-level mechanism. Rigorous explanation requires making that mechanism explicit and showing it actually connects the two macro levels. ABM is, in a very direct sense, Coleman's boat made computationally executable: you write the micro-level rule as code, and the "aggregation", in this case the simulation, is no longer asserted, but run.

Deep History

Cellular Automata and the Origins of Computational Emergence

The mathematical ancestor of ABM is the cellular automaton and its origin is worth knowing. Decades before any social scientist used the technique, it establishes that startling complexity can arise from trivial local rules, which is the core intuition the entire field rests on. John von Neumann's Theory of Self-Reproducing Automata constructed a cellular automaton (1966, edited by Arthur Burks). This automaton is a grid of cells, each updating its state based only on its neighbours' states. The cells were capable of self-reproduction, thereby answering a question motivated by biology (can a machine build a copy od itself?) with a strictly computational, rule-based construction.

Two decades later, Conway's game of Life, popularized in Martin Gardner's Scientific American column (1970), showed the same principle with the simplest possible rule set: a ceil lives, dies, or is born based purely on how many of his eight neighbours are alive. From this trivial rule, the Game of Life produces gliders that travel indefinitely, stable "still life" structures, and the capacity for universal computation, which was proven much later.

Stephen Wolfram's classification of cellular automata (1984; 2002) formalized an intuition in every subsequent modeller uses, often without naming it: automata fall into four qualitative classes: rule that settle into a fixed state, rules that cycle periodically, rules that produce chaos, and a fourth class producing complex, structured, neither-ordered-nor-chaotic patterns. That fourth class (poised in the language later popularized by complexity scientists) is where genuinely interesting emergent structure lives, and it is an intuition worth carrying into every model you design: rules that are too rigid produce boring, predictable output: rules that are too noisy produce structureless output. The interesting research questions live in between.

Schelling and the Discovery of Aggregation Dynamics in Social Science

Schelling's paper *Dynamic Models of Segregation* (1971) is retrospectively the founding demonstration of ABM in the social sciences. Despite predating the term itself, Schelling built his original model by hand, moving coins and buttons on a checkerboard. The mechanism is simple: agents of two types occupy cells on a grid. Each agent has a mild preference for having at least some fraction of same-type neighbours. Any agent whose neighbourhood falls below the threshold relocated to a satisfactory empty cell. The finding that made this paper foundational is that even very mild individual preferences reliably produce nearly complete aggregate segregation.

It demonstrates that an aggregate pattern (residential segregation, and by direct analogy, alliance bloc formation or ideological clustering among states) cannot be read backward to infer the intensity or even the sign of individual-level attitudes. A moderate, even benign, individual-level preference for similarity is sufficient to generate an aggregate outcome that looks, from the outside, like strong collective sorting. Anyone using an observed international alignment as direct evidence of strong underlying animosity or ideological commitment at the state level should have Schelling's result in the back of their mind.

Axelrod and the Birth of Computational Political Science

Robert Axelrod's *The Evolution of Cooperation* (1984) is the field's direct bridge into political science and IR. Axelrod ran a computer tournament in which submitted strategies played the iterated prisoner's dilemma against one another. The simple strategy Tit-for-Tat (cooperate first, then do whatever the opponent did last round) won repeatedly against far more complex submissions. Axelrod's analysis identified four properties behind Tit-for-Tat's robustness: it is

- Nice (never defects first)
- Retaliatory (punishes defection immediately)
- Forgiving (returns to cooperation as soon as the opponent does)
- Clear (simple enough for an opponent to recognize and adapt to)

It demonstrates that cooperation among purely self-interested actors can emerge and stabilize without central enforcement, provided interactions are repeated and the future is weighted heavily enough. The substantive claim this supported was directly aimed at, and reframed, a core question in IR theory about cooperation under anarchy.

His follow-up book, *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration* (1997), is important because it explicitly names and formalizes agent-based simulation as, in Axelrod's own framing, a third way of doing science alongside induction (finding patterns in empirical data) and deduction (proving consequences of assumptions). Simulation, Axelrod argues, aids intuition in a different way, by allowing a research to explore consequences of a set of assumptions that are too complex to solve analytically.

The same book introduces landscape theory, which is a formal model for predicting how a set of actors will divide into alignment blocs given a matrix of pairwise relationships (friendly or hostile). The theory works by minimizing "frustration", which is the total number of relationships in the resulting configuration that are unsatisfied (an ally treated as an enemy or vice versa). This is a direct computational formalization of Heider's structural balance theory, who argued "the friend of my friend is my friend; the enemy of my enemy is my friend". Axelrod & Bennett (1993) applied this directly to IR in *A Landscape Theory of Aggregation*. It used the frustration-minimization algorithm to predict the alignment of European powers in the run-up to WWII from a matrix of historical bilateral relationships, and found the model's predicted alignments matched the historical record closely. This paper is, in effect, an analytical predecessor to the alliance formation model that will be build later in this piece (Track C).

The Santa Fe Institute and Epstein and Axtell's Sugarscape

The Santa Fe Institute, founded in 1984, established complexity science as a self-conscious interdisciplinary field studying emergence, adaptation, and self-organization across physical, biological, economic, and social systems. In this regard, agent-based simulation is a central shared method across otherwise disconnected disciplines. Within this environment, Epstein & Axtell (1996) built *Sugarscape*, documented in their article *Growing Artificial Societies: Social Science from the Bottom Up.* Sugarscape place heterogenous agents, differing in metabolism and vision, on a landscape of a regenerating resource ("sugar"), and shows that from nothing but local foraging, reproduction, and movement rules, an entire cascade of macro-institutional looking phenomena spontaneously appears. Examples include skewed wealth distributions resembling real inequality, migration waves, disease transmission, the emergence of trade and endogenous prices that approach the theoretical equilibrium prices of formal trade theory, cultural transmission, and even proto-warfare over resource-rich territory. No institution was programmed into the model directly, but each was grown from agent-level rules. Sugarscape remains the field's manifesto text as it is the most vivid available demonstration that phenomena we ordinarily explain by invoking institutions can, in principle, be explained without assuming those institutions by only assuming the more primitive behaviours that could plausibly give rise to them.

ABM's Epistemological Position: Between Positivism and Interpretivism

IR theory has long organized itself, sometimes too rigid, around a divide between

- Positivist approaches, which seek causal, law-like generalizations testable against observation, in tradition of Hempel's covering law-model of explanation); and
- Interpretivist or constructivist approaches, which treat meaning, discourse, and intersubjective understanding as irreducible and resist flattening politics into law-like regularities.

ABM does not sit comfortably in either camp. On one hand, ABM is formal and mechanistic in a way that satisfies positivist instincts: every assumption is written down as executable code, nothing is left as unspecified black box, and the model construction is fully replicable. On the other hand, the standard by which ABM is judged successful is almost never point prediction against a single observed outcome (the classic positivist test). Rather, ABM is judged on generative sufficiency and plausibility: a model earns scientific value by showing that a particular set of mechanisms is capable, in principle, of producing an observed class of patterns, and not by claiming to have identified the unique true mechanism.

Epstein has been explicit that this licenses a kind of "explanatory pluralism", meaning several distinct generative mechanisms may each be independently sufficient to produce the same macro pattern. Demonstrating this is itself valuable, because it clarifies what is logically necessary for the pattern to arise versus what is merely one possible sufficient pathway among several. This approach is much closer to the "structured environment for exploration" framing than to a claim of discovered law. It thus leaves room for the kind of reflexivity about how models shape perception that interpretivist scholars insist on, while still producing the kind of explicit, falsifiable, mechanism-level claims that positivist scholars value.

It is also worth noting that the field has its own internal positivist-flavoured rigour debate. It is worth being aware of precisely because it shows ABM policing its own generative pluralism rather than treating it's just an explanatory through experiment as an excuse for sloppiness. Axtell et al. (1996) introduced the concept of "docking", which is the practice of verifying that two independently implemented models of the supposedly same mechanism actually produce equivalent results. This is a replication check unique to the method as you cannot "dock" a regression the same way, but you very much can, and should, dock a simulation. This concern reappears later in D2 where calibration, validation, and the ODD protocol are discussed. It is one of the clearest signs that ABM, despite its generative and pluralist epistemology, hold itself to serious standards of reproducibility.

ABM and Game Theory: Complement, not Substitutes

Classical, analytical game theory specifies a game from and solves for an equilibrium concept typically under an assumption of common knowledge of rationality among the players. This buys enormous analytical power: a proof that a particular outcome is an equilibrium is a proof, not a simulation result contingent on parameter choices. What it does not tell you by construction is whether a large, heterogeneous population of only boundedly rational actors would actually discover that equilibrium, how long it would take, what transient dynamics look like along the way, or whether the system might instead settle into a different, locally stable but globally suboptimal configuration.

Evolutionary game theory already loosened the full-rationality assumption in an important way. Smith and Price's (1973) *The Logic of Animal Conflict* introduced the Evolutionary Stable Strategy (ESS) concept. Rather than assuming individual rational choice, strategies compete and propagate through a population according to their relative payoffs. An ESS is a strategy that, once dominant, cannot be successfully invaded by a rare alternative. This is usually formalized through replicator dynamics, which are differential equations describing how strategy frequencies change over time in a large, well-mixed population.

ABM is best understood as the fully general computational extension of this move. Whereas replicator dynamics still assume a well-mixed, typically homogeneous population interacting via aggregate frequencies, ABM allows explicit heterogeneity, explicit spatial or network structure, and observation of the system mid-transition rather than only at its long-run attractor. Axelrod's iterated prisoner's dilemma tournaments were, in this sense, already an agent-based approach to a nominally game-theoretic question. He did not solve for the equilibrium of the repeated game analytically, but he simulated a population of concrete strategies interacting and let the empirically dominant strategy reveal itself.

Game theory and ABM are best used together: game theory characterizes what the equilibria of a given strategic structure are and why they are stable once reached. ABM tells you whether a realistic population, with realistic cognitive and informational limit, is likely to find a particular equilibrium, how fast, and what happens if it does not.

ABM and Network Analysis: Generative Mechanism vs. Statistical Inference

Standard social network analysis, particularly the exponential random graph model (ERGM) (Wasserman & Pattison, 1996; Snijders et al., 2006) is fundamentally inferential. Given an observed network, it specifies a probability distribution over graphs with particular structural features. Then it estimated, typically via Markov Chain Monte Carlo, which structural tendencies are statistically over- or underrepresented relative to chance. It allows formal hypothesis tests about tie-formation processes.

Agent-based network-formation models, by contrast, are generative. You specify a micro-level behavioural rule for how and why an actor forms or breaks a tie and grow the network forward from an initial condition to then compare the resulting structure to what it observed empirically. The two traditions are converging in an important way. Snijders' (2001) stochastic actor-oriented models (SAOM) are essentially a hybrid. They specify an agent-level micro-mechanism for tie change, in the spirit of ABM, but then estimate the parameters of that mechanism from real longitudinal network data using formal statistical inference in the spirit of ERGMs. This bridge is still underexploited.

ABM and Machine Learning/Artificial Intelligence: 4 Distinct Points of Contact

The relationship between ABM and the current wave of machine learning and AI is not a single connection, but at least four analytically distinct ones. Conflating them is a common course of confusion worth avoiding from the outset.

First, machine learning as a calibration and estimation tool for ABM. Because agant-based models typically lack a tractable lilelihood function, fitting them to real data has historically been difficult. Beaumont's survey (2010) on Approximate Bayesian Computation (ABC) addresses this by simulating the model many times across a range of parameter values and retaining only those parameter draws that produce simulated data sufficiently close to the observed data. This is a likelihood-free inference framework built almost entirely around the practice of running the simulation repeatedly, which is precisely what the NetLogo function BehaviourSpace is for. A related and increasingly common practice trains a machine-learning surrogate on a sample of expensive simulation runs, to approximate the model's input-output relationship cheaply enough search for well-fitting parameter regions without re-running the full simulation thousands of times.

Second, ABM as a synthetic data generator for machine learning. When real data on a phenomenon is sparse, sensitive, or simply unobserved, calibrated agent-based model can generate large synthetic datasets on which other statistical or machine-learning models can be trained or stress-tested.

Third, multi-agent reinforcement learning (MARL) (Busoniu et al., 2008). Rather then hand-coding a fixed behavioural rule for an agent, MARL has agents learn a policy through repeated interaction and reward feedback, converging on strategies the researcher did not specify in advance. This is, in essence, the AI community studying the identical question (emergent behaviour from interacting adaptive agents) from an optimization and control angle rather then a social-scientific explanatory angle. The overlap with ABM is close enough that the same NetLogo model can, in principle, be extended so its agents learn via reinforcement rather then following fixed rules.

Fourth, and newest, LLMs embedded directly as an agent's decision function. Park et al. (2023) placed LLM-driven characters in a simulated town, where each agent formed memories, made plans, and reacted to other agents, producing the emergent coordinated behaviour purely from language-model-driven reasoning rather than hand-coded rules. It is a working demonstration about human-machine collaboration, delegation of analytical tasks, and trust in algorithmic outputs.

ABM and Complex Systems Theory: Emergence Made Operational

Several core ideas from complexity science are not just closely linked or related to ABM. Rather, ABM puts these ideas into practice. Naming them clearly is useful, as a lot of ABM-research relies on their terms. Self-organized criticality describes how systems naturally move toward a critical point. At this point, minor triggers can cause cascading chain reactions of any size following a power law, without needing outside adjustments (Bak et al., 1987). This directly addresses the question: why do certain crises spread catastrophically while similar ones disappear, even without major external shock? Similarly, path dependence and lock-in explain how small early events can steer a system toward a specific long-term state under increasing-returns dynamics (Arthur, 1989). This applies straight to market competition and technical standards, where early adoption can lock in a dominant technology regardless of superior alternatives. Both concepts highlight a mainly methodological lesson: linear statistical models fail to capture systems that stay stable for long periods before suddenly crossing a threshold into new behaviour. Agent-based models are naturally suited to modelling exact these threshold dynamics by building macro outcomes from the bottom up.

Where ABM Sits in International Relations and Security Studies Specifically

Within IR and security studies specifically, ABM remains a minority tradition. However, it is a well-established and growing one, and has a recognizable set of anchor works and scholar. Cederman's *Emergent Actors in World Politics: How States and Nations Develop and Dissolve* (1997) applies the generative ABM-logic to state formation and dissolution. He treats the modern state system itself as an emergent, historically contingent outcome of micro-level competitive and alliance dynamics rather than a natural or inevitable unit of analysis. His later article *Modelling the Size of Wars: From Billard Balls to Sandpiles* (2003) is particularly worth reading because it directly relates to Lewis Fry Richardson's work. He observed that war sizes (battle deaths) follow something close to a power-law distribution to a self-organized-criticality-style agent-based mechanism. It ties together literature from historical arms-race, complex systems, and formal international relations theory in a single argument.

Other relevant ABM-IR work includes Bhavani & Backer's (2000) computational work on ethnic conflict applies agent-based reasoning to genocide dynamics at the sub-state level and Weidman's (2010) work combining geographic information systems with conflict modelling.

It is equally important to state the field's genuine limitations rather than only its promise. Three criticisms recur consistently in methodological critiques of ABM and IR in the social sciences more broadly.
First, the "garbage in, garbage out" concern. Because agent-level rules and parameters are often chose for plausibility rather then estimated from data, a model can be tuned, consciously or not, to produce whatever macro pattern the researcher expected. It undermines the generative claim's evidentiary force unless calibration and validation are taken seriously.
Second, the small-N problem is, if anything, worse at the level of great-power iIR than in most domains where ABM as flourished. Because there are very few historical instances of, say, systemic great-power war, it is difficult to make pattern-oriented validations against real macro-historical outcomes.
Third, reproducibility and the docking problem discussed above are not fully solved by having code available. Subtle implementation choices (e.g. the order in which agents are asked to act, the specific pseudo-random number generator used, floating-point rounding) can produce materially different results across nominally identical model specifications. It is precisely why the ODD protocol and disciplined use of BehaviorSpace across many random seeds are not optional polish but load-bearing parts of a credible research design.

Discussion Questions

1. Epstein's generative standard requires only that a mechanism be sufficient to produce a pattern, not that it be the unique true mechanism. Is this a strength (explanatory pluralism) or a weakness (unfalsifiability) for a field aiming at policy relevance? Defend your answer against the opposite reading.
2. Schelling's segregation result shows that mild individual preferences can produce strong aggregate sorting. Pick one real international alignment pattern (a specific alliance bloc, a specific instance of great-power alignment) and argue both for and against the claim

that it reveals strong underlying state-level animosity or ideological commitment, rather than a Schelling-style aggregation of milder individual tendencies.

3. Axelrod's landscape theory formalizes Heider's structural balance theory as frustration-minimization. What does this formalization assume about state preferences that a constructivist IR scholar would likely object to? Is the objection fatal to the model's usefulness, or does it simply define the model's scope condition?

4. One of the above sections argues game theory and ABM are complements. Construct a concrete counterexample or edge case where you think a policymaker would be positively misled by relying on an ABM's transient dynamics instead of a game-theoretic equilibrium analysis, or vice versa.

5. The above section distinguishes four separate ABM-AI relationships. Which of the four is most directly relevant to a specific project you might pursue in the next two years, and why do the other three matter less for that specific project (even though all four are legitimate research directions in general)?

6. The chapter connects self-organized criticality to the question of why some crises cascade and others don't. Name one historical international crisis you think plausibly exhibits self-organized criticality-like dynamics (a long quiet build-up, then a disproportionate cascade from a small trigger), and one you think does not, and explain the difference in your reasoning, not just the conclusion.

7. The chapter raises the small-N problem as a distinctive weakness of ABM applied to great-power politics specifically. Propose one concrete strategy (empirical, methodological, or both) you could use in your own future work to partially mitigate this weakness, beyond simply acknowledging it exists.

8. Having read the full synthesis at the end of this chapter, write, in no more than 150 words, the single sentence you would put in a grant application or a research proposal explaining why agent-based modelling, specifically, rather than any other computational method available to you, is the right tool for your research agenda. Be prepared to defend every clause of that sentence.

Independent Exercise

1. Locate and skim Epstein's 1999 Complexity article "Agent-Based Computational Models and Generative Social Science" directly (not a secondary summary). Write a 300–500 word précis in your own words of his generative-sufficiency argument, and note one point at which you think his argument is under-defended.
2. Read Schelling's original 1971 paper or a detailed secondary account of it, and attempt to state, in one paragraph, the minimum similarity threshold below which no segregation emerges in his model, based on his own reported results.

3. Find Axelrod and Bennett's 1993 "A Landscape Theory of Aggregation" (British Journal of Political Science). Sketch, in pseudocode or plain English, the frustration-minimization procedure it describes. You will use a working version of this logic in Module C2.
4. Pick any one JASSS article (browse jasss.org) from the last five years that touches on conflict, security, or international politics. Identify which of the four ABM–AI relationships from section 2.9, if any, it uses, and whether it engages with calibration/validation (section 2.6/module D2 preview) at all.
5. Draft a one-paragraph generative-sufficiency argument (in the Epstein sense) for one specific IR phenomenon central to your own research agenda. For example, the persistence of a particular alliance structure, or the escalation trajectory of a particular historical crisis. State explicitly what micro-level mechanism you are proposing and what macro-level pattern it would need to grow in order to count as a successful generative explanation.
6. Write a half-page comparison of how a game theorist, a network analyst, and an agent-based modeller would each approach the same question: "why did a particular wartime alliance hold together while another one collapsed?" Be specific about what each method would treat as its primary explanandum and what each would leave unexplained.

Further Reading

Foundational method and epistemology

- Epstein, J. M. (1999). "Agent-Based Computational Models and Generative Social Science." Complexity, 4(5). The field's clearest statement of the generative standard of explanation; read this first.
- Epstein, J. M., & Axtell, R. (1996). Growing Artificial Societies: Social Science from the Bottom Up. Brookings Institution Press / MIT Press. The Sugarscape manifesto; dense but foundational, worth reading at least the first three chapters in full.
- Coleman, J. S. (1990). Foundations of Social Theory. Harvard University Press. The methodological-individualist backdrop behind the "Coleman boat" framing of micro-macro explanation.
- Axtell, R., Axelrod, R., Epstein, J. M., & Cohen, M. D. (1996). "Aligning Simulation Models: A Case Study and Results." Computational and Mathematical Organization Theory, 1(2). The original "docking" paper; short and directly useful once you reach Module D2.

Historical and computational lineage

- von Neumann, J. (1966). Theory of Self-Reproducing Automata (ed. A. W. Burks). University of Illinois Press. Read a secondary summary unless you have substantial patience for dense formalism.

- Wolfram, S. (1984). "Universality and Complexity in Cellular Automata." Physica D, 10(1–2). The original four-class cellular automaton taxonomy.
- Schelling, T. C. (1971). "Dynamic Models of Segregation." Journal of Mathematical Sociology, 1(2). Short, remarkably readable, and worth reading in full rather than only in summary.
- Axelrod, R. (1984). The Evolution of Cooperation. Basic Books. Highly readable; the field's most successful crossover book.
- Axelrod, R. (1997). The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration. Princeton University Press. The explicit ABM methodology chapter (chapter 1) is essential; the landscape theory chapter is directly relevant to Module C2.
- Axelrod, R., & Bennett, D. S. (1993). "A Landscape Theory of Aggregation." British Journal of Political Science, 23(2). Read before Module C2.

Game theory, networks, and complexity

- Maynard Smith, J., & Price, G. R. (1973). "The Logic of Animal Conflict." Nature, 246. The original Evolutionarily Stable Strategy paper.
- Wasserman, S., & Pattison, P. (1996). "Logit Models and Logistic Regressions for Social Networks." Psychometrika, 61(3). Foundational ERGM paper; technical.
- Snijders, T. A. B. (2001). "The Statistical Evaluation of Social Network Dynamics." Sociological Methodology, 31(1). The origin of the stochastic actor-oriented model / SIENA framework.
- Bak, P., Tang, C., & Wiesenfeld, K. (1987). "Self-Organized Criticality: An Explanation of 1/f Noise." Physical Review Letters, 59(4). The sandpile model paper.
- Arthur, W. B. (1989). "Competing Technologies, Increasing Returns, and Lock-In by Historical Events." Economic Journal, 99(394). Path dependence and technology lock-in.

Machine learning and AI

- Beaumont, M. A. (2010). "Approximate Bayesian Computation in Evolution and Ecology." Annual Review of Ecology, Evolution, and Systematics, 41. The standard review of ABC, the main likelihood free calibration framework relevant to ABM.
- Buşoniu, L., Babuška, R., & De Schutter, B. (2008). "A Comprehensive Survey of Multiagent Reinforcement Learning." IEEE Transactions on Systems, Man, and Cybernetics, Part C, 38(2).

- Park, J. S., O'Brien, J., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. (2023). "Generative Agents: Interactive Simulacra of Human Behavior." Proceedings of UIST 2023. Directly relevant to Module D5. Computational international relations and security studies
- Cederman, L.-E. (1997). Emergent Actors in World Politics: How States and Nations Develop and Dissolve. Princeton University Press.
- Cederman, L.-E. (2003). "Modeling the Size of Wars: From Billiard Balls to Sandpiles." American Political Science Review, 97(1). Read before or alongside Module D2.
- Bhavnani, R., & Backer, D. (2000). "Localized Ethnic Conflict and Genocide: Accounting for Differences in Rwanda and Burundi." Journal of Conflict Resolution, 44(3).
- Weidmann, N. B., & Ward, M. D. (2010). "Predicting Conflict in Space and Time." Journal of Conflict Resolution, 54(6). Relevant to Module B2 and D3 (GIS extension).
- Richardson, L. F. (1960). Statistics of Deadly Quarrels. Boxwood Press. The empirical backdrop to both the arms race model (Module C1) and Cederman's sandpile argument above.

Journals to browse regularly

- Journal of Artificial Societies and Social Simulation (JASSS) — jasss.org, fully open access, the field's dedicated venue.
- International Studies Quarterly, Journal of Conflict Resolution, and Complexity — occasional computational/ABM pieces in a mainstream IR and security studies context