



Nordic Centre
for Internet and Society

Designing Tech Regulation for Complex Systems: Insights from Complexity Science


Nordic Conference for Law and IT
Oslo, 6 Nov 2025
Samson Esayas, s.y.esayas@jus.uio.no

What do ant colonies, pandemics, and online social networks have in common?



EWS

Ant colonies



Pandemics



Online social networks

Emergence • Nonlinearity • Feedback Loops

AI as a complex system

Emergent Abilities in Large Language Models: A Survey

Leonardo Berti, Flavio Giorgi, Gjergji Kasneci

Abstract—Large Language Models (LLMs) are leading a new technological revolution as one of the most promising research streams toward artificial general intelligence. The scaling of these models, accomplished by increasing the number of parameters and the magnitude of the training datasets, has been linked to various so-called *emergent abilities* that were previously unobserved. These emergent abilities, ranging from advanced reasoning and in-context learning to coding and problem-solving, have sparked an intense scientific debate: *Are they truly emergent, or do they simply depend on external factors, such as training dynamics, the type of problems, or the chosen metric?* What underlying mechanism causes them? Despite their transformative potential, emergent abilities remain poorly understood, leading to misconceptions about their definition, nature, predictability, and implications. In this work, we shed light on emergent abilities by conducting a comprehensive review of the phenomenon, addressing both its scientific underpinnings and real-world consequences. We first critically analyze existing definitions, exposing inconsistencies in conceptualizing emergent abilities. We then explore the conditions under which these abilities appear, evaluating the role

of numerous simple elements, he observed that collective phenomena, such as stable magnetic orientations or vortex patterns in fluid dynamics, arise from the interactions of these basic elements. This observation prompted Hopfield to investigate whether the computational capabilities of neural networks could be understood as an emergent property resulting from the interactions of many simple neuronal units. Anderson's and Hopfield's insights laid the foundation for understanding how complex behavior can emerge from simple interactions, a principle that continues to influence modern artificial neural networks. This idea has become particularly relevant in deep learning with the advent of large language models (LLMs) [11, 82, 102]. These models have fundamentally revolutionized the field of natural language processing, achieving state-of-the-art performance through novel techniques such as in-context learning and chain-of-thought prompting. By leveraging a few examples within the input prompt, LLMs demonstrate

[Check for updates](#)

SCIENCE ADVANCES | RESEARCH ARTICLE

SOCIAL SCIENCES

Emergent social conventions and collective bias in LLM populations

Ariel Flint Asbery¹, Luca Maria Aiello^{2,3}, Andrea Baronchelli^{1,4*}

Social conventions are the backbone of social coordination, shaping how individuals form a group. As growing populations of artificial intelligence (AI) agents communicate through natural language, a fundamental question is whether they can bootstrap the foundations of a society. Here, we present experimental results that demonstrate the spontaneous emergence of universally adopted social conventions in decentralized populations of large language model (LLM) agents. We then show how strong collective biases can emerge during this process, even when agents exhibit no bias individually. Last, we examine how committed minority groups of adversarial LLM agents can drive social change by imposing alternative social conventions on the larger population. Our results show that AI systems can autonomously develop social conventions without explicit programming and have implications for designing AI systems that align, and remain aligned, with human values and societal goals.


Copyright © 2025 The Authors, some rights reserved; exclusive license American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).

INTRODUCTION

Last, a third question concerns the robustness of social conven-



Nonlinear jumps



Emergent capabilities

BI

Much of legal thinking remains rooted in an industrial-age, reductionist mindset



Atomistic approach

Focus on isolating specific units of analysis—a specific market, data processing, AI system—as a target of regulation



Additive outcomes

A system is equal to the sum of its parts



Causal relationships

Between input and output; a behaviour and its outcomes



- Focusing on a single component as a target of regulation in the face of multiple interacting components
- Studying single ant in ant colonies

- Focusing on a single component as a target of regulation in the face of multiple interacting components
- Studying single ant in ant colonies

- Data protection
- Competition law
- AI regulation (AI Act)

- Data protection
- Competition law
- AI regulation (AI Act)


- Difficult to isolate an atom
- Misses emergent, system-level risks:

- Difficult to isolate an atom
- Misses emergent, system-level risks:


Atomized approach in data protection and competition law



- **Focus on isolating**
 - Single data processing operation (DPL)
 - Relevant product/service market (CL)
- **Regulatory challenges in light of complexity from**
 - Emergence of companies processing personal data for multitude of purposes/services
 - Emergences of multi-product & ecosystem markets
- **Atomized approach**
 - Misses emergent, system-level behaviors:
 - Privacy risks arising from totality of data processing operations
 - Economic power from managing relations and data access



One Google Account for everything Google



Atomized approach in AI regulation

- **AI Act - single AI system, single risk logic**
- **Challenges from emergent behavior from multi-agent systems**
 - Collective bias or manipulation
- **Recognizes systemic risk but**
 - Parameter still focuses on capabilities of a single model
- **Emergence from MAS → implications for many other domains**
 - Liability



SOCIAL SCIENCES

Emergent social conventions and collective bias in LLM populations

Ariel Flint Ashery¹, Luca Maria Aiello^{2,3}, Andrea Baronchelli^{1,4*}

Social conventions are the backbone of social coordination, shaping how individuals form a group. As growing populations of artificial intelligence (AI) agents communicate through natural language, a fundamental question is whether they can bootstrap the foundations of a society. Here, we present experimental results that demonstrate the spontaneous emergence of universally adopted social conventions in decentralized populations of large language model (LLM) agents. We then show how strong collective biases can emerge during this process, even when agents exhibit no bias individually. Last, we examine how committed minority groups of adversarial LLM agents can drive social change by imposing alternative social conventions on the larger population. Our results show that AI systems can autonomously develop social conventions without explicit programming and have implications for designing AI systems that align, and remain aligned, with human values and societal goals.

INTRODUCTION

Social conventions shape social and economic life, determining how individuals behave and their expectations (1–4). They can be defined as unwritten, arbitrary patterns of behavior that are collectively shared by a group. Examples range from conventional greetings like handshakes or bows to language and moral judgments (5, 6). Recent numerical (7, 8) and experimental (9) results have confirmed the hypothesis that conventions can arise spontaneously, without the intervention of any centralized institution (3, 5, 10, 11). Individual efforts to coordinate locally with one another can generate universally accepted conventions.

Do universal conventions also spontaneously emerge in populations of large language models (LLMs), i.e., in groups of N simulated agents instantiated from an LLM? This question is critical for predicting and managing artificial intelligence (AI) behavior in real-world applications, given the proliferation of LLMs using natural language to interact with one another and with humans (12–14). Answering it is also a prerequisite to ensure that AI systems behave in ways aligned with human values and societal goals (15).

A second key question concerns how the biases of individual

Last, a third question concerns the robustness of social conventions. Recent theoretical (22) and empirical (23) results have shown how a minority of adversarial agents can exert an outsized influence on the group, provided that they reach a threshold or “critical mass” (24–26). Investigating how conventions change through critical mass dynamics in a population of LLMs will help anticipate and potentially steer the development of beneficial norms in AI systems, while mitigating risks of harmful norms (27). It will also provide valuable models for how AI systems might play a role in shaping new societal norms to address global challenges such as antibiotic resistance (28) and the post-carbon transition (29).

Here, we address these three key questions—on the spontaneous emergence of conventions, the role of individual biases, and critical mass dynamics—in populations of LLM agents. Drawing from recent laboratory experiments with human subjects (9, 23, 30), we follow the well-established practice of using coordination on a naming convention as a general model for conventional behavior (5, 7, 30–33). In this setting, agents are endowed with purely local incentives and conventions may (or may not) emerge as an unintended consequence of individual attempts to coordinate locally with one another. This

Copyright © 2025 The Authors. Some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).

Downloaded from https://www.science.org on June 15, 2025

EU Commission’s preliminary views on AI agents raise legal applicability question



By Luca Bertuzzi

November 5, 2025, 11:07 GMT | Insight


EU digital chief Henna Virkkunen shared the European Commission’s initial views, prompting questions on whether AI agents always qualify as AI systems under the AI Act. Lawmaker Sergey Lagodinsky asked the commission to clarify its position, warning that the law’s current definition may not capture agents’ autonomous actions. The commission plans to monitor developments and may take further action.

RELATED SECTIONS

Artificial Intelligence

Additive logic & the fallacy of composition

- **Additive logic: a system = sum of its parts**
 - Law: if each component complies, the system is deemed compliant
 - Risk of ignoring emergent risks
- **Illustration - US constitutional privacy**
 - *US v. Jones* (2010)
 - “[t]he sum of an infinite number of zero-value parts is also zero.”
 - Classic fallacy of composition
 - Reversed in *Carpenter v. U.S.* (2018)
 - Aggregate data can violate privacy even if individual datasets don’t
 - Recognizing emergent behavior from scale – ‘more is different’



This house is made of bricks. A brick is light in weight. Therefore, this house is also light in weight.

Similar tendencies in copyright law, LLMs & training data

- **Key challenge: reproduction vs. transformation**
 - Difficult to prove substantial similarity between training data and outputs
- **Claimants' systemic argument**
 - LLMs are “compressed copies” of training data.
 - Every output a “recasting or adaptation” of protected works
 - Infringement occurs at the system level
- **Courts' emerging (reductionist) view**
 - No infringement unless a specific output or model layer is substantially similar
- **Lesson from privacy to copyright law**
 - Change in scale may imply a change in principle?

Causal dynamics

- **Causality in many legal domains**
 - Tort law: conduct and harm
 - Copyright law: input and output
- **Regulatory challenges from emergence**
 - Tort law:
 - Regulatory challenges with autonomous systems — a “victim without a perpetrator” problem
 - Copyright law:
 - Discussions on AI generated content — creative outputs without an author



What insights does complexity science offer to the field of law?

- **Mental modes**
 - System level analytical framework
 - Accounting for system wide (e.g. emergent) behaviors
- **Tools & methods for understanding & modelling complex systems**
 - Agent-based modelling
 - Network analysis

Mental modes

- **A system level analytical framework**
 - Ecosystem approach in competition law
 - Shift from regulating single markets to ecosystem of markets
 - New forms of power
 - Incorporate emergent thinking
 - Developments in the US regarding constitutional privacy
- **System level analytical framework in AI regulation**
 - Beyond single-model, single output logic
 - Consider MAS & system level testing in certain sectors?
 - Fitting institutional framework
 - Institutions with system-level view?

Industrial and Corporate Change

IssuesJEL ▾Advance articlesSubmit ▾PurchaseAbout ▾Industrial and Corp: ▾

Industrial and Corporate Change

Volume 30, Issue 5
October 2021

JOURNAL ARTICLE

Ecosystems and competition law in theory and practice

Michael G Jacobides, Ioannis Lianos

Industrial and Corporate Change, Volume 30, Issue 5, October 2021, Pages 1199–1229, <https://doi.org/10.1093/icc/dtab061>
Published: 18 November 2021Article history ▾

PDFSplit ViewCitePermissionsShare ▾

Article Contents

Abstract

1. Introduction

2. Understanding ecosystems: framing observations

3. Identifying the regulatory gap

4. Tackling the regulatory gap

5. Breaking new ground: the Greek proposal and its rationale as a case study

6. Challenges in putting forth an ecosystem regulatory framework

Funding

Footnotes

References

Author notes

< PreviousNext >

Abstract

One of the most profound changes in the industrial landscape in the last decade has been the growth of business ecosystems—groups of connected firms, drawing on (digital) platforms that leverage their complementors and lock in their customers, exploiting the “bottlenecks” that emerge in new industry architectures. This has created new asymmetries of power, where the “field” of competition is not the relevant product market, as is usually the case in competition law, but rather the ecosystem of various complementary products and associated complementor firms. These dynamics raise novel concerns over competition. After examining the foundational elements of the ecosystem concept, we review how ecosystems are addressed within the current scope of competition law and identify the gap in the existing framework of conventional competition law. We then move to a critical review of current efforts and proposals in the European Union for providing regulatory remedies for *ex ante* and *ex post* resolution of problems, focusing on the current (2020) proposals of the Market Act on *ex ante* regulation, with its particular focus on “gatekeepers” and on the role of business models and ecosystem as a regulatory framework before providing a deep dive into proposed Greek legislation explicitly focuses on ecosystem regulation. We conclude with our observations on the challenges in instituting and implementing a regulatory framework for ecosystems, drawing on research and our own engagement in the research process.

Article

The Right to Quantitative Privacy

David Gray[†] & Danielle Citron^{††}

Introduction 63

I. Quantitative Privacy: The Perils of Broad and Indiscriminate Surveillance 73

II. Quantitative Privacy and the Fourth Amendment 83

A. Qualitative Privacy: The Fourth Amendment Before *United States v. Jones* 83

B. A Fourth Amendment Foothold for Quantitative Privacy in *United States v. Jones* 87

C. The Fourth Amendment Foundations of

[†] Professor of Law, University of Maryland Francis King Carey School of Law.



^{††} Lois K. Macht Research Professor & Professor of Law, University of Maryland Francis King Carey School of Law, Affiliate Scholar, Stanford Center on Internet and Society, Affiliate Fellow, Yale Information Society Project. The authors thank everyone who generously commented on this work during presentations at Yale's Information Society Project, the Annual Meeting of the ABA/ALS Criminal Law Section, the University of North Carolina, North-western, Yale's Conference on Locational Privacy and Biometrics, Law and Society, the Privacy Law Scholars Conference, and the Computers, Freedom, and Privacy Conference, and during conversations at the American Law Institute Meeting on Information Privacy Law and Harvard Law Review's Symposium on Informational Privacy. Particular thanks go to Ronald Allen, Julia Angwin, Jack Balkin, Kevin Bankston, Steve Bellwin, Marc Bitt, Richard Boldt, Becky Bolin, Mary Bowman, Al Brophy, Andrew Chin, Bryan Choi, Thomas Clancy, Julie Cohen, Thomas Crocker, Nick Doty, LisaMarie Freitas, Susan Friedwald, Barry Friedman, Brandon Garrett, Bob Gellman, Dan Gifford, Mark Graber, John Grant, James Grimmelmann, Deborah Hellman, Leslie Meltzer Henry, Lance Hoffman, Rensé Hutchins, Camilla Hrdy, Orin Kerr, Joseph Kennedy, Catherine Kim, Anne Rinefelter, Avner Levin, Michael Manheim, Dan Markel, Christina Mulligan, Richard Myers, Neil Richards, Catherine Sabbe, Laurent Sacharoff, Paul Schwartz, Christopher Slobogin, Robert Smith, Dan Solove, Max Stearns, David Super, Harry Surden, Peter Swire, Peter Quint, Jason Weisstein, Arthur Weisburd, and Jonathan Witmer-Rich. We also thank Liz Clark Rinehart for her research assistance and Max Siegel for his insightful editing. Finally, we are grateful to Frank Lancaster for holding us together. Copyright © by David Gray and Danielle Citron.

BI

62

Tools & methods

- **Competition law leading the way**
 - ABM to uncover emergent collusive behavior between firms (AI agents)
 - Network analysis to map positional & data power
- **Opportunities in other legal domains**
 - ABM + network analysis for liability attribution
 - Use ABM to identify & assess emergent (systemic) risks in AI Act?



Final remarks

- **Complex systems are everywhere — tech law must evolve**
- **Challenges**
 - Law built on stability, predictability, and clear causal chains
- **Opportunities — a new way of seeing and governing**
 - Adopt system-level mental models
 - Experiment with new tools and methods
- **Ambition**
 - Not to make law complex — but to make it fit for a complex world

Therapeutic
Junk