

Large Language Models and the Future of Organization Theory

Cornelissen, Joep; Höllerer, Markus A.; Boxenbaum, Eva; Faraj, Samer; Gehman, Joel

Document Version

Final published version

Published in:

Organization Theory

DOI:

[10.1177/26317877241239056](https://doi.org/10.1177/26317877241239056)

Publication date:

2024

License

CC BY-NC

Citation for published version (APA):

Cornelissen, J., Höllerer, M. A., Boxenbaum, E., Faraj, S., & Gehman, J. (2024). Large Language Models and the Future of Organization Theory. *Organization Theory*, 5(1). <https://doi.org/10.1177/26317877241239056>

[Link to publication in CBS Research Portal](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us (research.lib@cbs.dk) providing details, and we will remove access to the work immediately and investigate your claim.

Download date: 16. Jun. 2024





From the Editors

Large Language Models and the Future of Organization Theory

Organization Theory
Volume 5: 1–15
© The Author(s) 2024
Article reuse guidelines:
sagepub.com/journals-permissions
DOI: 10.1177/26317877241239056
journals.sagepub.com/home/ott



Joep Cornelissen^{1,2}, Markus A. Höllerer^{3,4,5} ,
Eva Boxenbaum⁶ , Samer Faraj⁷  and Joel Gehman⁸ 

Abstract

In this editorial essay, we explore the potential of large language models (LLMs) for conceptual work and for developing theory papers within the field of organization and management studies. We offer a technically informed, but at the same time accessible, analysis of the generative AI technology behind tools such as Bing Chat, ChatGPT, Claude and Gemini, to name the most prominent LLMs currently in use. Our aim in this essay is to go beyond prior work and to provide a more nuanced reflection on the possible application of such technology for the different activities and reasoning processes that constitute theorizing within our domain of scholarly inquiry. Specifically, we highlight ways in which LLMs might augment our theorizing, but we also point out the fundamental constraints in how contemporary LLMs ‘reason’, setting considerable limits to what such tools might produce as ‘conceptual’ or ‘theoretical’ outputs. Given worrisome trade-offs in their use, we urge authors to be careful and reflexive when they use LLMs to assist (parts of) their theorizing, and to transparently disclose this use in their manuscripts. We conclude the essay with a statement of *Organization Theory*’s editorial policy on the use of LLMs.

Keywords

AI technology, ChatGPT, large language models, LLM, organization theory

¹Erasmus University, Rotterdam, Netherlands

²University of Liverpool, United Kingdom

³UNSW Sydney, Sydney, NSW, Australia

⁴WU Vienna, Austria

⁵IAE Business School, Universidad Austral, Buenos Aires, Argentina

⁶Copenhagen Business School, Frederiksberg, Denmark

⁷McGill University, Montreal, QC, Canada

⁸George Washington University, Washington, DC, USA

Corresponding author:

Joep Cornelissen, Erasmus University, Burgemeester Oudlaan, Rotterdam, 3000 DR, Netherlands.

Email: cornelissen@rsm.nl



Creative Commons Non Commercial CC BY-NC: This article is distributed under the terms of the Creative Commons Attribution-NonCommercial 4.0 License (<https://creativecommons.org/licenses/by-nc/4.0/>) which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is attributed as specified on the Sage and Open Access pages (<https://us.sagepub.com/en-us/nam/open-access-at-sage>).

In recent years, many individuals, organizations and communities in society have experimented with the use of a growing array of large language models (LLMs), such as Bing Chat, ChatGPT, Claude and Gemini, to name but a few of the most prominent LLMs currently in use. As advanced artificial intelligence (AI) systems capable of mimicking human reasoning, LLMs have increasingly impacted the world of science and academia as well. For scientific research, the AI systems behind these advances have been billed as a ‘game-changer’ (Van Dis et al., 2023, p. 224), with many scientists and scholars anticipating that over time LLMs may replace many of the data analysis and theorizing processes now done by human beings. Indeed, the overriding belief is that such algorithms will reconstitute how scholars across the natural and social sciences form a theoretical understanding of their respective phenomena of interest (e.g. Beam et al., 2024; Heidt, 2023; Salvagno et al., 2023; Susarla et al., 2023; Thorp, 2023).

Considering these broader developments, the role of AI, and of LLMs specifically, is starting to be discussed within organization and management studies as well (e.g. Gatrell et al., 2024; Grimes et al., 2023). So far, the debate in the field around LLMs has – perhaps unsurprisingly – been divided between critics and evangelists: those who, from a normative standpoint, are staunchly against LLMs based on what it would undermine and take away from our inherently *human* scholarship (e.g. Lindebaum & Fleming, 2023), and those who, convinced of its technological prowess, see its productive advantages and argue that, when adopted, it will augment, rather than undercut, human reasoning, creativity and innovation (e.g. Dwivedi et al., 2023). These alternate positions are reflective of the emerging nature of the debate; one that furthermore is animated by a lot of uncertainty about what LLM tools such as ChatGPT can *actually* do; and how, guided by yet to be established standards and norms (Van Dis et al., 2023), they may come to be much more widely used by organization and management researchers. From our perspective, what may be needed

here as a further step in this emerging conversation is in fact a more nuanced look at how, based on what we now know, LLMs such as ChatGPT are likely to affect the core processes of theorizing that up until now have been largely the preserve of human organizational researchers.

In this editorial, we endeavour to provide such a perspective in relation to conceptual and theoretical papers. We do so by first sketching the acts of theorizing that routinely go into developing and writing such papers (Cornelissen et al., 2021); this sketch forms the backdrop to a dedicated discussion of the uses and possible impacts of LLMs. Our approach here essentially involves a focused thought experiment through which we consider, in a disciplined manner, likely scenarios of the impacts of LLMs on theorizing within our field.

Our inquiry is guided by the question: What does it mean, or will it mean, for the processes and products of our theorizing if they are automated or augmented with the assistance of generative AI systems? If some parts of theorizing, such as, say, summarizing or synthesizing a literature, defining a set of concepts or formulating a new proposition can be more quickly and effectively done by, for instance, Open AI’s ChatGPT or Google’s Gemini, will it indeed enhance our capacity to theoretically understand phenomena? Furthermore, what is potentially excluded or fundamentally changed in terms of the very human aspects of the theorizing process, such as creative and innovative conceptual thought, as well as the very human *existential* struggle to understand the world around us? Will we come to see the same kind of path-breaking conceptual innovations in our field, and indeed within the social sciences at large, as in the past, if our theorizing becomes to a greater extent than before automated or augmented by AI? Will future theorizing become even more a mere reflection of the past, given the algorithms defining contemporary LLMs? Or might AI technology allow us, as it further develops, to explore way beyond the current horizon in organization and management theory, and lead to new discoveries?

Attempting to address these questions is crucial, we believe, as the increasing use of LLMs is inevitable and seems likely to become commonplace across our field. Banning such technology, therefore, might not work and is no option. Rather, it is imperative that we engage in a debate about the implications of this potentially disruptive technology for our theorizing and for the kinds of knowledge that we in turn create about our various phenomena of interest. We offer our perspective on these questions based on a deep reading of the technology underlying LLMs and informed by fundamental insights about the conceptual processes underlying different forms of creative, inferential thought (Fauconnier & Turner, 2008). The answers and views that we provide thus take the form of an informed perspective, while recognizing that LLMs may quickly improve on the specific points that we make. They furthermore reflect our interest in fostering and maintaining a healthy pluralism in different knowledge interests and forms of theorizing, as an important ideal that we have previously advocated (see Cornelissen et al., 2021; Reinecke et al., 2022). Finally, having provided our reading of the potential and limitations of LLMs for acts of theorizing, we conclude the paper by outlining our editorial policy on the use of LLMs in submissions made to *Organization Theory*.

Acts of Theorizing

Before looking in detail at the way in which advanced LLMs are already affecting our theorizing processes, and how we think they are likely to do so in the future, it may be helpful to first give a brief account of the processes, or acts, of theorizing through which we usually produce knowledge. Figure 1 offers a simplified representation of the way in which theoretical and conceptual papers (those that are not primarily based on empirical data) produce and contribute theoretical knowledge to help better understand certain key phenomena. All three parameters (priors, theorizing, phenomena) are connected, with the acts of theorizing that are done by researchers being the central conduit

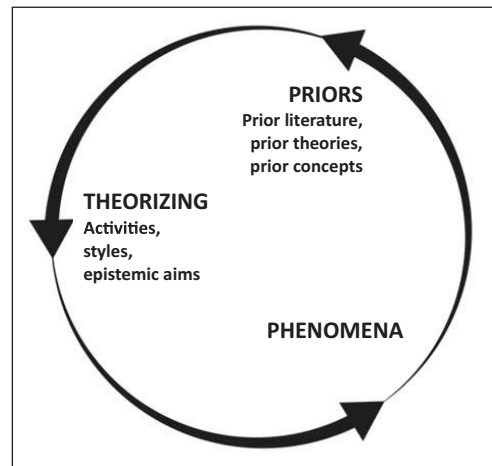


Figure 1. Acts of theorizing.

through which priors (prior literature or record of knowledge, prior theories, prior concepts, and so on) are transformed into a heightened or enlightened interpretation of a phenomenon under study.

Phenomena, as we will use the term here, are the recurring and general features of the real or experienced world of individuals and organizations; they are analytically separate from, and prior, to the researchers' theorizing and attempts at understanding them; they constitute things such as, for example, the socially constructed understanding of a particular workplace, or the diffusion and widespread use of certain technologies across organizations. Phenomena are ultimately the objects (or targets) which organizational researchers seek to better conceptualize and understand.

Priors involve the prior stock of knowledge that has been built up and recorded in the academic literature on a particular phenomenon (or related phenomena), often encapsulated in the form of *prior* 'theories', including existing theoretical concepts or constructs, relationships, processes and the like. Such priors are in fact the outcome or remnants of preceding acts of theorizing that have produced such knowledge, and on which researchers in turn build when they aim to further refine or enlighten our scholarly understanding – as a novel act of *theorizing*

that aims to contribute yet further to our body of knowledge on the phenomenon.

Acts of *theorizing*, as they are commonly understood (see Cornelissen et al., 2021), involve a set of *typical activities* that researchers engage in, such as reviewing and synthesizing prior literature, problematizing current understanding, defining key terms, conceptualizing (i.e. abstracting out presupposed features of) a phenomenon, crafting a line of argument around such a conceptualization, and producing formative outputs that distil the main contribution to knowledge (such as a new or revised concept, theoretical model, typology, and so on) (Weick, 1995). These activities are furthermore given shape, and in different ways, depending on the *theoretical style of reasoning* (propositional, interpretative, emancipatory, and so on) that is used by a researcher. We here employ the term ‘styles’ (in the plural) to highlight the fact that theorizing is not confined to one form but features multiple established forms of theoretically reasoning *about* phenomena (see Cornelissen et al., 2021). Each of these styles, such as, say, the propositional one or an intersectional feminist critique, effectively comes with a distinct grammar to reason with, and is furthermore guided by different *epistemic goals*, or knowledge interests. Linked to such goals, styles enable researchers – each in their own way – to configure their argument with the aim of furthering or changing our understanding; whether that is, say, to further explain or predict a phenomenon, conceptually frame and interpret it differently, or establish a way to provoke a more critical understanding of, and engagement with, a given phenomenon as their main knowledge interest (Cornelissen et al., 2021).

The notion of theorizing styles itself suggests that all three parameters (phenomena, priors, theorizing) are, in the practice of theorizing, intimately interconnected (and thus hard to separate) through the webs of implicit and explicit inferences that researchers make between them and that support them in developing theoretical claims from priors to better understand phenomena – whatever the qualifier ‘better’ might mean to different communities of researchers. As part

of this process, virtually all styles, are, as forms of reasoning, *ampliative* in the sense that the conclusion that is reached in this way (i.e. a novel theoretical claim about a phenomenon) has additional content beyond the priors. In each case, a researcher *creatively* draws in additional substantive theoretical assumptions and ideas that, configured into an inferential line of reasoning, go beyond the codified priors and thus expand and refine our knowledge of phenomena.

Set out in this way, Figure 1 offers a rather simple and stylistic depiction of the theorizing process. Actual acts of theorizing, as embodied by researchers, may no doubt be much more messy, iterative (with many false starts, or back and forth movements) and complex. The picture is probably also by itself not exhaustive nor complete and it presupposes a role for a pluralism in theorizing styles and epistemic goals to expand understanding and keep thought moving (as opposed to being, say, more narrowly focused on explanation and prediction only). But our main objective here is to offer a basic scaffolding, however imperfect, for reflecting on our current acts of theorizing – in terms of basic *activities* that we generally engage in when we theorize, such as synthesizing a literature; specific *inferential styles of reasoning* that we routinely use, such as formulating propositions on the basis of prior theory or knowledge; and specific *epistemic goals* that we have, as researchers, and are guided by in our scholarly pursuits. Scaffolding the theorizing process in this way provides a backdrop against which to explore the potential uses and impact of LLMs on acts of theorizing, broadly conceived. Specifically, if we want to figure out how LLMs as a generative form of AI might impact theorizing, we need a sense of the ways in which they might be *generative* and what in turn, in terms of our acts of theorizing, they might be *generative for*.

Large Language Models

ChatGPT was the first product to bring large language models to the world’s attention.

Launched in November 2022, within two months it had garnered more than 100 million active users. OpenAI's success also precipitated the release of further generative AI technologies and applications in this space. For instance, Google released Gemini in December 2023, a multimodal LLM that can synthesize and reason about textual, numerical and visual data and which the company claims will be a game-changer for various domains of science. In the same month (but announced with less bravado), Clarivate, the owner of Web of Science (WoS), launched a built-in AI tool based on LLM technology – the Web of Science Research Assistant – to partly automate the review process of the contents of journals indexed by WoS. Cumulatively, these developments have led to LLMs being increasingly prominent and available for research purposes – well before academics were able to properly debate the technology, let alone put standards and norms for its responsible uses in place (Van Dis et al., 2023). And yet, surprisingly many of our university colleagues have not even attempted to experiment with this new toolbox.

A noted concern which bears repeating is that these LLM technologies are proprietary products of a small number of companies with commercial interests (Van Dis et al., 2023). For the most part, the companies behind these products do not share details on the inner workings of their technologies, or indeed on the underlying training sets on which they are based. This lack of transparency and its challenge to responsible, open science is a key concern (Ahmed et al., 2023). Without researchers being able to publicly and openly retrace how, for instance, a review or synthesis of a given literature is formed with the help of an AI assistant, it is hard to uncover the inferential strands that lie at its base. In the absence of open code, we may, to a large extent, 'sail blind' when we use such AI technology, even when we might establish, through human verification, that what such AI systems generate as outputs is intelligible or accurate enough to the task at hand.

Despite this general lack of transparency, we can still form a picture of how LLMs work,

creating a sense of how, with increasing use, they are likely to have a bearing on theorizing in management and organization studies. Although OpenAI, the company behind ChatGPT, has not released any data on its current version (GPT-4), it did openly share data on an earlier version (GPT-3). We briefly highlight what we know about this version and use that to extrapolate to more powerful versions of LLMs now and in the future (including GPT-4 and Google's Gemini) as a way of deciphering how LLMs may come to be more broadly used for research and, as is our focus here, specifically for theorizing purposes.

How do large language models work?

LLMs are deep learning, transformer-based AI models that are trained on large amounts of 'text' that makes them capable of understanding and generating natural language and other types of content (such as images, models, or figures). They are based on a neural network architecture that was first introduced by Google (Vaswani et al., 2017), which effectively models language processing as a process of transforming inputs, such as a question the chatbot is asked, into generated text that it offers as an answer. The architecture itself is built around an extensive, multidimensional space, where each word in a corpus is represented, and those words with similar meanings are placed closer together (Mikolov et al., 2013). Google's word2vec project (Mikolov et al., 2013) laid the basis for these vector representations, powerfully showing how word vectors allowed LLMs to combine semantic and contextual information about words to accurately predict which words tend to appear together in similar sentences. But the project also showed that with the same vector-based architecture, the LLM could 'reason' about problems or tasks using vector arithmetic, be able to understand and manipulate natural images, as well as creatively form analogies and metaphors (in language and imagery) as one of the pinnacles of creative human thought.

LLMs effectively employ such vector spaces across multiple layers through which it computes what, as it is processing, is likely to come next. Each of its layers takes a sequence of vectors as inputs – one vector for each word in the input text – and then adds information to help clarify the meaning of that word in context, and to better predict which word might come next. By sharing information across the vector space, the model calculates what a likely and thus useful response to the initial prompt could be. This means that ChatGPT is always aiming to produce a ‘reasonable continuation of whatever text it has already written, where reasonable implies what one might expect someone to write after seeing what people have written on billions of webpages’ (Wolfram, 2023).

Designed in this way, each layer acts as a so-called ‘transformer’, a computational agent or algorithm making predictions and feeding its input to the next layer. Across the vector space, there is in turn a division of labour between layers and the network connections that are made (i.e. transformed) in this way. Predictions made by earlier layers can be read and modified by later layers, allowing the LLM to sharpen its understanding and help it produce a seemingly more reasonable and intelligible text in return. Across the layers, the network thus combines both attention and information sharing mechanisms with feed-forward mechanisms with layers ‘thinking’ about the input given, and trying to predict the likely probability of the next word and alongside the text that is, in this way, composed as a whole. These transformations happen in parallel using massive computing power and, because of the model’s training, having a reach that is well beyond the cognitive abilities of any human being. Specifically, GPT-3 featured 12,288-dimensional word vectors and 96 layers, for a total of 175 billion parameters.

Some of this reasoning capability is built into LLMs, such as the engineered and trained ability of LLMs to operate on the specifics of words (allowing it to resolve ambiguities in context) alongside the representation (paragraph, text, corpus) as a whole (Tenney et al.,

2019). But, LLMs have, as a result of the transformer technology, also developed a capacity to generate intelligible outputs in ways that were not deliberately engineered into these models (Kosinski, 2023) and which suggests that rather than ‘stochastic parrots’ they may have formed an ‘intelligence’ and ‘ability’ to reason on their own (Bubeck et al., 2023; Kosinski, 2023) and in ways that, operating over billions of parameters, may for sizeable tasks and large datasets easily outperform humans. These capabilities provide LLMs with a unique ability to rival specific forms of human reasoning, including abstract and creative forms of reasoning – something that makes this technology potentially so disruptive.

Using Large Language Models as Part of Theorizing

At the same time, there are important differences between LLMs and human reasoning (Cristianini, 2023; Wolfram, 2023) that have a bearing on how such technology may be used for acts of theorizing (see Figure 1 above). As a foray into illustrating these differences, we asked ChatGPT (GPT-4) if it could come up with a new metaphor for organizations, as an abstract and creative reasoning task. It offered as a response to think of organizations as:

intricate ecosystems where each department or team functions like a unique species contributing to the overall biodiversity. Just as in nature, a delicate balance and collaboration among these entities are crucial for the health and success of the entire ecosystem.

The metaphor it offered is in effect a simple simile that is somewhat banal and uninspiring, drawing a basic correspondence with a familiar and salient domain (ecosystems) that is bolted onto a basic, rudimentary understanding of organizations (as a coordinated effort). It is, as a creative product, quite flat, based on how the technology has leveraged large data regularities to project a probabilistic rendering from what it already ‘knows’.

Of course, this kind of output does not come close to the kinds of creative, generative metaphors that have been developed by organizational scholars (including in the pages of *Organization Theory*) and that progress our understanding of organizations, such as, for example, ground-breaking metaphors of ‘surveillance capitalism’ (Zuboff, 2022) or ‘institutional isomorphism’ (Powell & DiMaggio, 2023). We asked ChatGPT in turn why it had suggested this metaphor. The chatbot answered that in its view the metaphor ‘captures the interconnectedness, diversity, and interdependence within a workplace’. In other words, it could only affirm its answer, based on vector arithmetic (determining its likelihood and match), but it could not, in a meta-representational sense, reflect in any meaningful way on what it had produced including any strengths as well as limitations or biases to this image.

The example is telling in multiple ways. It shows the obvious strength of LLMs as well as their evident limitations when they are used fully as a means of theorizing and not just as an editing tool or heuristic input for a particular task. Its obvious strength is that, as a feed-forward system, it can generate intelligible enough conceptual output (such as the ecosystems simile), and can do so in ways that, depending on the task at hand, might be useful to researchers – for example, as a concept to work with and extend into a more full-blown conceptualization. ChatGPT is particularly good at reaching forward to a ‘new’ answer based on what, operating over large tracts of language data (as *priors*), it knows and can leverage into a likely new candidate inference. In this way it can indeed ‘reason’, but in ways that are guided by the kinds of probability assessments that are built into its transformer capability. Such capability makes LLMs good at *basic* reasoning tasks, as *acts* of theorizing (see Figure 1), such as, for example, providing a review and basic synthesis of a large literature or dataset, or offering a novel candidate hypothesis on the back of previously confirmed hypotheses and evidence (Belikov et al., 2022). Particularly when the *prior* literature (see Figure 1) is large

enough, it allows the vector arithmetic to recognize and compute correspondences across data points into an abstracted, representationally accurate set of simple surface-level inferences over such a corpus (Vapnik, 1999). It can do so much faster and with more precision than individual researchers or research teams. Not surprisingly, therefore, this kind of basic inferential, or theorizing, activity is where its potential uses for scientific research have been demonstrated already.

Specifically, within fields such as the biomedical sciences, LLMs have been assisting scientists in synthesizing a large knowledge base, extracting patterns from doing so, and recommending fruitful new areas of inquiry and new hypotheses to accelerate cumulative scientific advance – and in ways that have garnered more success than the models and prescriptions that human researchers had previously been working from (Belikov et al., 2022). But, importantly, this success merely shows its strength for a particular set of automatable basic ‘reasoning’ or ‘theorizing’ tasks (see Figure 1), such as summarizing a literature, and in service of a particular model of scientific research (Grimes et al., 2023). Indeed, the tasks that it is currently particularly good at completing aligns particularly well with what we have previously labelled as a propositional style of theorizing (Cornelissen et al., 2021) which involves organizational scholars synthesizing prior research and producing a new set of propositions or hypotheses (as similes) on the back of prior evidence and theorizing about a phenomenon.

Abstract reasoning and creativity

In this context, the million-dollar question is whether LLMs can now, or in the near future, be truly creative beyond such basic tasks, and in ways that might capture or mimic human levels of creativity in acts of theorizing. It is perhaps a bit premature here to give a conclusive answer, but classic work on creativity (Koestler, 1964; Fauconnier & Turner, 2008) might already provide some important clues. Creativity generally involves making novel connections between

different ideas, as frames or ‘thought matrices’ (Koestler, 1964), that are joined together and elaborated into a new combined frame or reframing that offers a new set of candidate inferences for thinking about and understanding a particular phenomenon (see Figure 1). Such a reframing may be more incremental, requiring lower levels of imagination and creativity, as well as more radical, with bold imaginative leaps generating completely new ways of seeing and understanding a phenomenon (Boxenbaum & Rouleau, 2011).

Fauconnier and Turner (2008) formalized these differences by distinguishing between creative combinations that, on the one hand, are simple variations of the input frames (e.g. managers *as* leaders) or mirror images of one another (e.g. similes such as organizations as ecosystems) versus those that, on the other hand, involve dynamically elaborated recombinations, or ‘blends’, that draw intricate connections across multiple, often dissimilar frames, and expand thought in that way. An example of the latter (see Oswick et al., 2011) is Foucault’s (1977) panopticon metaphor which creatively and counter-intuitively reconfigured our understanding of how a subtle and all-encompassing disciplining power is exercised by institutions in free and democratic societies involving techniques derived from penal incarceration.

When LLMs are used as inputs for creative conceptual thought, they are likely to offer outputs that are similar to what has gone before. Because of its prior training and the probability-based algorithms on which it operates, an LLM is prone to draw connections across similar registers or ‘matrices’ of words and concepts (Koestler, 1964) which it will ‘interpolate’ into a novel, but largely familiar output (Cristianini, 2023; Wolfram, 2023). The LLM may conceivably be ‘helped’ by researchers to overcome this constraint, such that, with further prompts and questions, it may be pushed to draw more unusual connections across more distantly related sets of ideas (such as one of us tried by asking ChatGPT to develop a Deleuze and Guattari-inspired take on various subjects).

Even when the LLM is helped this way, however, the mathematized connections that it makes across distant sets of ideas will largely produce ‘reasonable’ new correspondences, such that it is bound to end up with syntactic structures (new concepts, categories, types, propositions, and so on) that in many cases will be stale intersectional representations that do not capture the rich inferential machinery that a worked-through ‘blend’ such as the panopticon provides. Of course, whatever the LLM produced might still be useful as a starting point for further thinking and theorizing, as an interim output on the way to potentially a novel theoretical insight.

However, the general point to note here is that besides struggling to make previously unforeseen, improbable connections across frames or ideas (as Foucault did), LLMs also cannot reproduce the active process of thought experimentation (Kornberger & Mantere, 2020) through which a researcher goes to produce a coherent and compelling new conceptualization. Such thought experimentation highlights that for virtually all *theorizing styles* (with the exception perhaps of the mentioned style of simply extending a proposition or hypothesis on the back of prior literature), the new conceptualization and generated inferences rest on the connections that researchers actively elaborate and work through in their minds, and which often, rather than extending past knowledge (as LLMs do), involves them revisiting, questioning and problematizing prior literature and established understandings. The crux here, of course, is that vector-based neural networks, like computers in general, are ultimately just dealing with data; any newly generated data points (symbolically converted into outputs such as new concepts, categories, types, propositions, and so on) cannot substitute, at least not fully, for the creative, ampliative reasoning that humans do and that underlies such connections (see Figure 1).

In the future, then, might it be conceivable that LLMs develop the capability to produce, say, a reasonable enough process theory or a

sufficiently well-developed post-colonial critique for a given subject? It might happen, but when it does it will be in a largely syntactic way by resembling the form of what such *styles of theorizing* look like, analogous to what we discussed for propositional and metaphorical reasoning. The fundamental idea of the neural, vector-based architecture of LLMs is that it provides these systems with a flexible and plastic ‘computing fabric’ that can be incrementally modified and extended when it learns further examples (Wolfram, 2023). As such, researchers in the organization and management domain could in the future train an LLM, or their own modified version of an LLM, with batches of, say, process model papers so that the AI can learn to recognize those cases and examples, and can leverage those into reasonable enough ‘new’ outputs. But, as we have argued, the outputs generated in this way will feed forward extensions of what it already knows and will be devoid of the human ingenuity and deep thought that for process models as well will be more often than not the basis for their substance and unique insight.

Furthermore, even when a form of a *theorizing style* can be copied and can generate reasonably creative outputs to start with, many of these theorizing styles presuppose *human reflexivity* (Lindebaum & Fleming, 2023), the meta-representational ability of human beings to reflect on and interrogate the assumptions, biases and (consequential) privileges associated with particular styles of reasoning and knowing. ChatGPT has, as we have highlighted, no ability to reflect on, let alone question for itself, what it produces, or, for instance, its ethical implications. In addition, because its vectors are built from the way in which humans generally use words (as embedded in its training sets), LLMs end up reflecting many of the biases that are present in human language data (Caliskan et al., 2017) and which in turn are perpetuated in what it produces. For instance, a recent experimental study found that when ChatGPT-4 was used as an assistant in different clinical situations it ‘did not appropriately model the demographic diversity of medical

conditions, consistently producing clinical vignettes that stereotype demographic presentations’ (Zack et al., 2024, p. e12), including races, ethnicities and genders, with the potential for great harm when it is integrated into clinical decision-making.

Automating, augmenting, or displacing human theorizing?

Based on our deconstruction of how LLMs work and what they produce, what is our take on their possible uses, now and in the future, as part of acts of theorizing? A first conclusion is that LLMs, as an algorithmic form of forwarding new ‘theory’ in the form of a new synthesis, proposition or hypothesis, metaphor, or anything else ‘conceptual’, cannot directly – i.e. *without human intervention or modification* – offer a viable substitute for the kinds of abstract and creative forms of reasoning that make up acts of theorizing (see Figure 1 above). LLMs come perhaps closest to the propositional style, but even then, they only work in a very delimited way to the extent they can leverage large datasets and do not require any further inferential work (such as, say, integrating different theories as the basis for a novel set of integrative propositions or hypotheses). In short, on current evidence, acts of theorizing cannot be fully mimicked, nor therefore *automated*, by the vector arithmetic of LLMs; nor therefore is it likely that whatever it produces as theory will in and of itself – that is, without any further human modification – stand the test of making a viable theoretical contribution (that is, at least in terms of how we have defined such contributions here: as essentially worked-through forms of ampliative reasoning in support of theoretical claims that foster our understanding of phenomena).

Alternatively, one might argue that if LLMs cannot copy (nor therefore automate) the theorizing process fully, the technology may still come in handy as an ‘assistant’. LLMs can be used as a prod, as a heuristic, or as a general thinking or writing device to support and in turn ‘augment’ the theoretical reasoning done by

academics. From this perspective, LLMs may, as we have shown above, for example, aid in reviewing a large literature or knowledge base, provide a first textual input for a more extended human reasoning process, or through further prods by the researcher generate a more creative or detailed conceptual representation (e.g. a model or typology) which can then potentially form the basis for a meaningful theoretical contribution.

This augmentation line of thinking is reflective of much theorizing around emerging technologies in the organizational field, which suggests that technologies do not fully determine how they are being used (technological determinism) but set in motion new ‘co-constitutive’ or ‘relational’ technology–human constellations which often spawn new functionalities (hence, the ‘augmentation’ label) that serve human ends (e.g. Bailey et al., 2022). Applied to the context of theorizing, this suggests that LLMs, when they are fed into theorizing processes, may come to boost and expand the ways in which we theorize, in any of our common theorizing styles (Cornelissen et al., 2021), and, by doing so, ‘may improve existing [processes and] functions or create new ones’ (Bailey et al., 2022, p. 9).

Such augmentation logic is persuasive and may well be the overriding logic that many organization and management scholars will choose to follow; in fact, many have already begun incorporating the use of LLMs into their theorizing and writing (Grimes et al., 2023). As far as its logic goes, it is simultaneously comforting and affirming. We inevitably find a measure of comfort in that a new emerging technology, such as LLMs, will, like many previous technologies (printing press, computers, internet), not replace human ingenuity but rather unlock new possibilities for us (Bailey et al., 2022, p. 2). It is affirming in that it suggests that, rather than blindly following the technology (technological determinism) or the rhetoric of those pushing it upon us (e.g. Google, OpenAI, Clarivate), we have our wits about us and will figure out among ourselves (Bailey et al., 2022) how LLMs might responsibly and

productively be factored into our theorizing processes such that they support our (and not the AI’s!) goals and criteria of producing valuable knowledge about phenomena.

Notwithstanding its appeal, the augmentation logic comes with a catch. Its central underlying assumption is that crucial parts of the reasoning that we, as human researchers, do as part of our theorizing can now or in the (immediate) future be transferred and offloaded onto LLMs, while at the same time unlocking new possibilities for us. Such a presumed singularity might then *in extremis* mean that, in time, as the technology improves in its reasoning capability and becomes increasingly used:

The role of human academics may shift from traditional knowledge production to knowledge verification, oversight, and translation, and ensuring impact from that knowledge. The exclusivity of the academic profession might decrease as journals are more likely to accept and trust the knowledge production and synthesis from accessible AI tools. (Grimes et al., 2023, p. 1623)

As demonstrated by this quote, the broader concern here is that although we might use LLMs to more efficiently or effectively produce new knowledge and theory, our increasing reliance on the technology might at the same time fundamentally reconfigure scholarship and, in time, trigger the demise or ‘displacement’ of the institution of *human* knowledge production altogether (Orlikowski & Scott, 2023).

Though it is perhaps too early to make that call (as we do not know yet whether LLM-augmented forms of theorizing might actually take off, become commonly used, let alone become the new normal), we do want to articulate our stance in relation to this broader concern. From our perspective, any form of augmented theorizing that uses LLMs will inevitably work from the algorithm-based operators or ‘transformers’ at its base and which, through repeated and extended use in the field, might thus perpetuate its underlying predictive logic and probability maxims as a new *episteme* in itself – as a kind of confirmatory, tautological

paradigm. In other words, increasing use of LLMs may come to (pre)structure and channel the theorizing that is done by scores of researchers in a specific direction; and potentially away from the rich variety of theorizing styles and epistemic aims of human theorizing that we have highlighted. The clear risk here is that LLM-augmented theorizing effectively comes to create a single but narrow conduit where a diversity of epistemic goals is levelled to a single epistemic goal of machine-based prediction regarding the contours (i.e. likely attributes and outcomes) of a phenomenon (as somewhat advocated for by Grimes et al., 2023).

In a recent essay in *The New Yorker*, the historian Jill Lepore (2023) referred to this epistemological capture as a predictive delusion. She captures this delusion with the metaphor of a cupboard with ‘drawers’ representing all human knowledge, with each drawer offering a different mode or epistemology of knowing and understanding the phenomenal world around us. The algorithmic prediction at the heart of LLMs (and any augmented forms of theorizing derived from them) is just one drawer, yet when it is elevated to the only or preferred one, as it nowadays often seems to be (see, e.g., Grimes et al., 2023), this is tantamount to keeping certain other drawers of epistemologies (such as, say, interpretive phenomenology and critical theory) shut or only minimally opened.

To be clear about our stance here, making *useful predictions* (Liu et al., 2023), such as in the human genome project where ‘smart’ algorithms helped map genetic sequences, is a valuable form of furthering our understanding of some phenomena. At the same time, it is not the only means of understanding, nor therefore the only – or preferred – *epistemic goal* (Habermas, 1972). Phenomena themselves are not limited to conditions, characteristics or effects to be predicted; and within organization and management studies, as well as within the social sciences in general, phenomena can, as the past record shows (e.g. Foucault, 1977), be usefully constructed, and in turn understood, in multiple and alternate ways. The predictive delusion, in other words, is one where algorithmic models

drive and limit the range of what we understand phenomena to be, and what in turn we can theoretically say about them, at the expense of *what else* we might otherwise theorize about organizations and organizational life. This would without a doubt constitute impoverished, certainly not enriched, scholarship.

What is more, the architecture on which LLMs run cannot capture the very human motivation to engage in theorizing and in understanding the world around us, as well as the profound struggles that we often experience while doing so. Years ago, the feminist scholar belle hooks described how she came to theory

because I was hurting – the pain within me was so intense that I could not go on living. I came to theory desperate, wanting to comprehend – to grasp what was happening around and within me. Most importantly, I wanted to make the hurt go away. (hooks, 1991, p. 1)

Although other researchers no doubt have other motivations and epistemic goals, each organizational researcher, alone or with others, brings their own self to acts of theorizing (Wright et al., 2023), thus contributing to our collective knowledge production and doing so in ways that are uniquely and distinctly human, and which cannot as a human drive or quest for understanding be automated or replaced by machines – nor should we perhaps even try.

Organization Theory’s Editorial Policy on the Use of Large Language Models

As a journal, *Organization Theory’s* mission is to support the publication and dissemination of the best theoretical and conceptual work in the field of organization and management studies. Its mandate is to do so in a pluralistic and inclusive manner, encompassing different styles and forms of theorizing, spanning the entire breadth of this domain of scholarly inquiry, and ‘crossing paradigms, subject areas, disciplines and geographical communities’ (Cornelissen &

Exhibit I SAGE Editorial Guidance on the Use of Large Language Models (LLMs) in Submissions to SAGE Journals.

Sage recognizes the value of large language models (LLMs) (e.g. ChatGPT) and generative AI as productivity tools that can help authors in preparing their article for submission; to generate initial ideas for a structure, for example, or when summarizing, paraphrasing, language polishing etc. However, it is important to note that all language models have limitations and are unable to replicate human creative and critical thinking. Human intervention with these tools is essential to ensure that content presented is accurate and appropriate to the reader. Sage therefore requires authors to be aware of the limitations of language models and to consider these in any use of LLMs in their submissions:

- **Objectivity:** Previously published content that contains racist, sexist or other biases can be present in LLM-generated text, and minority viewpoints may not be represented. Use of LLMs has the potential to perpetuate these biases because the information is decontextualized and harder to detect.
- **Accuracy:** LLMs can ‘hallucinate’, i.e. generate false content, especially when used outside of their domain or when dealing with complex or ambiguous topics. They can generate content that is linguistically but not scientifically plausible, they can get facts wrong, and they have been shown to generate citations that don’t exist. Some LLMs are only trained on content published before a particular date and therefore present an incomplete picture.
- **Contextual understanding:** LLMs cannot apply human understanding to the context of a piece of text, especially when dealing with idiomatic expressions, sarcasm, humour, or metaphorical language. This can lead to errors or misinterpretations in the generated content.
- **Training data:** LLMs require a large amount of high-quality training data to achieve optimal performance. However, in some domains or languages, such data may not be readily available, limiting the usefulness of the model.

Guidance for authors

Authors are required to:

1. **Clearly indicate the use of language models in the manuscript**, including which model was used and for what purpose. Please use the methods or acknowledgements section, as appropriate.
2. **Verify the accuracy, validity and appropriateness of the content** and any citations generated by language models and correct any errors or inconsistencies.
3. **Provide a list of sources used to generate content** and citations, including those generated by language models. Double-check citations to ensure they are accurate, and are properly referenced.
4. **Be conscious of the potential for plagiarism** where the LLM may have reproduced substantial text from other sources. Check the original sources to be sure you are not plagiarizing someone else’s work.
5. **Acknowledge the limitations of language models in the manuscript**, including the potential for bias, errors and gaps in knowledge.
6. Please note that AI bots such as ChatGPT **should not be listed as an author** on your submission.

We will take appropriate corrective action where we identify published articles with undisclosed use of such tools.

Authors should check the guidelines of the journal they are submitting to for any specific policies that may be in place on that journal.

Höllerer, 2020, p. 2). The journal does not limit itself to only one form of theorizing, such as, for instance, formal theory or propositional work, nor does it presuppose, or privilege, a single mode or model of knowledge production over others.

In line with this editorial mission, we take, as an editorial team, a guarded and cautious

stance on the use of LLMs as part of organizational theorizing. We do so motivated by our mandate of fostering plural ways of theorizing and informed by the analysis that we have offered of the still limited capabilities of LLMs. This position leads us, in turn, to adopt a *pre-cautionary principle* (Jonas, 1984) when we

make editorial judgements on the use of LLMs in each manuscript submitted to the journal. We take the prudent and cautious view here that we need to judge its use in each instance in terms of whether it is compatible with a particular style of theorizing, supports the development of theory, and does not ultimately compromise the conditions for a pluralistic and inclusive system of knowledge production in the field. As editors, we reserve the right to reject papers on this basis, and such considerations could conceivably overrule other editorial criteria (such as whether a submission is complete, well written, or innovative in nature).

Concretely, what this means for authors is that we do not wish to declare an outright ban on the use of LLMs, nor do we wish to limit its future potential for innovation in knowledge production. However, considering our mandate and our reservations about LLMs, we stipulate two general requirements. First, authors need to openly and transparently declare any uses that they have made of LLMs. That is, in terms of their *authorship*, authors need to explicitly declare the specific uses they have made of LLMs in developing and writing their manuscript, including any extensive use they have made of LLMs for grammar or spellchecking (Van Dis et al., 2023). SAGE, as the publisher of *Organization Theory*, has provided a set of specific guidelines for this purpose (see Exhibit 1). We encourage authors to carefully follow these guidelines as they prepare their submissions with the aim of transparently disclosing and acknowledging, in their manuscript and in any supporting material (such as a cover letter), the use they have made of LLMs.

Second, authors are asked to spell out explicitly how LLMs feature in their chain of reasoning in support of any theoretical claims and contributions. In terms of substantive *scholarship*, LLMs cannot themselves be authors (see Exhibit 1), nor can they, as we have highlighted, produce fully formed theoretical outputs and contributions. As such, it is incumbent on each author to disclose in their manuscripts: (a) how an LLM was used to prime or augment their

theoretical reasoning; (b) how the algorithmic inferences that it provides sit alongside other theoretical assumptions and inferences made (for reference and guidance, see our discussion of ampliative reasoning around Figure 1 above); and (c) how they have themselves reflexively probed and considered the theoretical contribution and understanding that this LLM-augmented line of argument provides against alternative lines of reasoning and of knowing the same phenomenon. We encourage authors to heed these pragmatic principles as they construct and lay out their line of argument in support of any theoretical claims (culminating in, for instance, a process model, typology, integrative framework, provocation, or meta-synthesis), also because a careful consideration of the arguments made will anyhow enhance their scholarship.

Taken together, we hope that such a nuanced and balanced view, reflective of the opportunities mirrored by the augmentation argument as well as the considerable limitations and risks that stem from the fundamental constraints in how contemporary LLMs work, will help both authors and readers of *Organization Theory* manoeuvre the challenges that lie ahead for us as a scholarly community.

Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.


Funding


The author(s) received no financial support for the research, authorship, and/or publication of this article.

ORCID iDs

Markus A. Höllerer  <https://orcid.org/0000-0003-2509-2696>

Eva Boxenbaum  <https://orcid.org/0000-0002-7197-9805>

Samer Faraj  <https://orcid.org/0000-0002-6585-0351>

Joel Gehman  <https://orcid.org/0000-0003-0585-9351>

References

- Ahmed, N., Wahed, M., & Thompson, N. C. (2023). The growing influence of industry in AI research. *Science*, 379, 884–886.
- Bailey, D. E., Faraj, S., Hinds, P., Leonardi, P. M., & von Krogh, G. (2022). We are all theorists of technology now: A relational perspective on emerging technology and organizing. *Organization Science*, 33(1), 1–18.
- Beam, A. L., Drazen, J. M., Kohane, I. S., Leong, T.-Y., Manrai, A. K., & Rubin, E. J. (2024). Artificial intelligence in medicine. *New England Journal of Medicine*, 388(13), 1220–1221.
- Belikov, A. V., Rzhetsky, A., & Evans, J. (2022). Prediction of robust scientific facts from literature. *Nature Machine Intelligence*, 4, 445–454.
- Boxenbaum, E., & Rouleau, L. (2011). New knowledge products as bricolage: Metaphors and scripts in organizational theory. *Academy of Management Review*, 36(2), 272–296.
- Bubeck, S., Chandrasekaran, V., Eldan, R., & Zhang, Y. (2023). Sparks of artificial general intelligence: Early experiments with GPT-4. *arXiv*. <https://doi.org/10.48550/arXiv.2303.12712>
- Caliskan, A., Bryson, J. J., & Narayanan, A. (2017). Semantics derived automatically from language corpora contain human-like biases. *Science*, 356, 183–186.
- Cornelissen, J., & Höllerer, M. A. (2020). An open and inclusive space for theorizing: Introducing organization theory. *Organization Theory*, 1(1), 1–5. <https://doi.org/10.1177/2631787719887980>
- Cornelissen, J., Höllerer, M. A., & Seidl, D. (2021). What theory is and can be: Forms of theorizing in organizational scholarship. *Organization Theory*, 2(3), 1–19. <https://doi.org/10.1177/26317877211020328>
- Cristianini, N. (2023). *The shortcut: Why intelligent machines do not think like us*. CRC Press.
- Dwivedi, Y. K., Kshetri, N., Hughes, L., Slade, E. L., Jeyaraj, A., Kar, A. K., Baabdullah, A. M., Koohang, A., Raghavan, V., Ahuja, M., Albanna, H., Albashrawi, M. A., Al-Busaidi, A., S., Balakrishnan, J., Barlette, Y., Basu, S., Bose, I., Brooks, L., Buhalis, D., & ... Wright, R. (2023). Opinion paper: “So what if ChatGPT wrote it?” Multidisciplinary perspectives on opportunities, challenges and implications of generative conversational AI for research, practice and policy. *International Journal of Information Management*, 71, 102642.
- Fauconnier, G., & Turner, M. (2008). *The way we think: Conceptual blending and the mind's hidden complexities*. Basic Books.
- Foucault, M. (1977). *Discipline and punish: The birth of the prison*. Penguin.
- Gatrell, C., Muzio, D., Post, C., & Wickert, C. (2024). Here, there and everywhere: On the responsible use of artificial intelligence (AI) in management research and the peer-review process. *Journal of Management Studies*. Advance online publication. <https://doi.org/10.1111/joms.13045>
- Grimes, M., Von Krogh, G., Feuerriegel, S., Rink, F., & Gruber, M. (2023). From scarcity to abundance: Scholars and scholarship in an age of generative artificial intelligence. *Academy of Management Journal*, 66, 1617–1624.
- Habermas, J. (1972). *Knowledge and human interests*. Polity Press.
- Heidt, A. (2023). Artificial-intelligence search engines wrangle academic literature. *Nature*, 620, 456–457.
- hooks, b. (1991). Theory as liberatory practice. *Yale Journal of Law and Feminism*, 4(1), 1–12.
- Jonas, H. (1984). *The imperative of responsibility: In search of an ethics for the technological age*. University of Chicago Press.
- Koestler, A. (1964). *The act of creation*. Penguin Books.
- Kornberger, M., & Mantere, S. (2020). Thought experiments and philosophy in organizational research. *Organization Theory*, 1(3), 1–19. <https://doi.org/10.1177/2631787720942524>
- Kosinski, M. (2023). Theory of mind might have spontaneously emerged in large language models. *arXiv*. <https://arxiv.org/abs/2302.02083>
- Lepore, J. (2023, March 27). The data delusion. *New Yorker*. <https://www.newyorker.com/magazine/2023/04/03/the-data-delusion>
- Lindebaum, D., & Fleming, P. (2023). ChatGPT undermines human reflexivity, scientific responsibility and responsible management research. *British Journal of Management*. Advance online publication. <https://doi.org/10.1111/1467-8551.12781>
- Liu, L., Jones, B. F., Uzzi, B., & Wang, D. (2023). Data, measurement and empirical methods in the science of science. *Nature Human Behavior*, 7(7), 1046–1058.
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv*. <https://doi.org/10.48550/arXiv.1301.3781>

- Orlikowski, W. J., & Scott, S. V. (2023). The digital under-tow and institutional displacement: A sociomaterial approach. *Organization Theory, 4*(2), 1–16. <https://doi.org/10.1177/26317877231180898>
- Oswick, C., Fleming, P., & Hanlon, G. (2011). From borrowing to blending: Rethinking the processes of organizational theory building. *Academy of Management Review, 36*, 318–337.
- Powell, W. W., & DiMaggio, P. J. (2023). The iron cage redux: Looking back and forward. *Organization Theory, 4*(4), 1–24. <https://doi.org/10.1177/263178772312215>
- Reinecke, J., Boxenbaum, E., & Gehman, J. (2022). Impactful theory: Pathways to mattering. *Organization Theory, 4*(3), 1–20. <https://doi.org/10.1177/26317877221131>
- Salvagno, M., Taccone, F. S., & Gerli, A. G. (2023). Can Artificial Intelligence help for scientific writing? *Critical Care, 27*(1), 1–5.
- Susarla, A., Gopal, R., Thatcher, J. B., & Sarker, S. (2023). The Janus effect of generative AI: Charting the path for responsible conduct of scholarly activities in information systems. *Information Systems Research, 34*(2), 399–408.
- Tenney, I., Das, D., & Pavlick, E. (2019). *BERT rediscovers the classical NLP pipeline*. <https://arxiv.org/pdf/1905.05950.pdf>
- Thorp, H. H. (2023). ChatGPT is fun, but not an author. *Science, 379*, 313.
- Van Dis, E. A. M., Bollen, J., Zuidema, W., Van Rooij, R., & Bockting, C. L. (2023). ChatGPT: Five priorities for research. *Nature, 614*(7947), 224–226.
- Vapnik, V. N. (1999). *The nature of statistical learning theory*. Springer.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. *arXiv*. <https://doi.org/10.48550/arXiv.1706.03762>
- Weick, K. E. (1995) What theory is not, theorizing is. *Administrative Science Quarterly* 40: 385–390.
- Wolfram, S. (2023). *What is ChatGPT doing . . . and why does it work?* <https://writings.stephenwolfram.com/2023/02/what-is-chatgpt-doing-and-why-does-it-work>
- Wright, A. L., Kent, D., Hällgren, M., & Rouleau, L. (2023). Theorizing as mode of engagement in and through extreme contexts research. *Organization Theory, 4*(4), 1–26. <https://doi.org/10.1177/26317877231217310>
- Zack, T., Lehman, E., Suzgun, M., Rodriguez, J. A., Celi, L. A., Gichoya, J., Jurafsky, D., Szolovits, P., Bates, D. W., Abdulnour, R.-E. E. Butte, A. J., & Alsentzer, E. (2024). Assessing the potential of GPT-4 to perpetuate racial and gender biases in health care: A model evaluation study. *The Lancet Digital Health, 6*(1), e12–e22.
- Zuboff, S. (2022). Surveillance capitalism or democracy? The death match of institutional orders and the politics of knowledge in our information civilization. *Organization Theory, 3*(3), 1–79. <https://doi.org/10.1177/263178772211292>

Author biographies

Joep Cornelissen is a Professor of Corporate Communication at Rotterdam School of Management in the Netherlands and Chair in Strategy and Organization (part-time) at Liverpool University Management School, United Kingdom. He is currently Co-Editor-in-Chief of *Organization Theory*.

Markus A. Höllerer is Professor in Organization and Management at UNSW Business School, Sydney, Australia, as well as Senior Research Fellow in urban management and governance at WU Vienna University of Economics and Business, Austria. He is currently Co-Editor-in-Chief of *Organization Theory*.

Eva Boxenbaum is Professor of Organization and Management Theory at Copenhagen Business School. She is an Associate Editor of *Organization Theory* and served in 2022–2023 as Chair of the Organization and Management Theory (OMT) Division of the Academy of Management.

Samer Faraj holds the Canada Research Chair in Technology, Innovation and Organizing at the Desautels Faculty of Management, McGill University. He is an Associate Editor of *Organization Theory*.

Joel Gehman is Professor of Strategic Management and Public Policy, and Lindner-Gambal Professor of Business Ethics at George Washington University. He is an Associate Editor of *Organization Theory*.