ANNUAL REVIEWS

Annual Review of Sociology

The Sociology of Interpretation

Amir Goldberg and Madison H. Singell

Graduate School of Business, Stanford University, Stanford, California, USA; email: amirgo@stanford.edu, msingell@stanford.edu

Annu. Rev. Sociol. 2024. 50:85–105

First published as a Review in Advance on May 1, 2024

The *Annual Review of Sociology* is online at soc.annualreviews.org

https://doi.org/10.1146/annurev-soc-020321-030515

Keywords

meaning, interpretation, culture, cognition, framing

Abstract

Recent years have seen a growing sociological interest in meaning. In fact, some argue that sociology cannot confront its foundational questions without addressing meaning. Yet sociologists mean many things when they talk about meaning. We propose a practical approach that conceptualizes meaning as an instance of an actor interpreting a stimulus. Reviewing existing literature, we find that most sociological accounts understand interpretation either as categorization or as semantic association. We show that an integrated approach is analytically useful for conceptualizing shared interpretation and the process by which people coordinate their interpretations. This provides a framework for addressing interpretative heterogeneity when studying attitudinal or behavioral variance. We conclude by highlighting how recent advances in computational linguistics have opened exciting new possibilities for the study of interpretation, and suggest several avenues for future research.

I want to be a part of the people that make meaning, not the thing that is made.

—*Barbie* (2023 movie)

1. INTRODUCTION

As a sociological term, "meaning" has—ironically—multiple meanings (Patterson 2014). Sociologists rarely define what they mean when they talk about meaning. In some instances, they refer to the interpretation of a single object (Childress & Friedkin 2012); in others, they refer to people's deep-seated theories of value and purpose, and how those shape their behavioral decisions (Christin 2020, Frye 2012, Zelizer 2010). Sociologists also appear to mean different things when they talk about meaning-making. They might refer to how an individual makes sense of her experiences (Cerulo 2018), or to the collective process by which people influence each other's perceptions of the world (Doering & McNeill 2020). Consequently, different researchers operationalize meaning using a variety of analytically distinct but overlapping constructs, from frames and schemas to beliefs, values, scripts and logics. This conceptual pluralism has led some to go so far as to conclude that "meaning" might be a meaningless sociological pursuit altogether (Wuthnow 1989).

Yet, despite these ambiguities, sociologists agree that meaning is important. Although they use different terminology and conceptual frameworks, they appear to be in consensus that culture inheres in the social processes that lead people to share meanings (Mohr et al. 2020, Tavory & Timmermans 2013). Meaning, moreover, is not solely the purview of cultural sociologists. As various observers have noted, one cannot fully understand how inequality emerges and persists without accounting for meaning (Lamont et al. 2014, Monk 2022).

A quarter of a century ago, Mohr (1998) argued that measuring meaning should be one of sociology's main objectives. A great deal has happened since. New data, computational methods, and findings from cognitive science ushered a wave of sociological work focused on measuring meaning and theorizing about its genesis and evolution. Indeed, the number of articles in leading sociological journals explicitly referencing meaning in their abstracts or titles has increased more than threefold in the past two decades, compared with the two decades prior (among articles published in the *American Sociological Review* and the *American Journal of Sociology* in 2003–2022 compared with 1983–2002).

Our objective is to offer an organizing framework for navigating this exciting but disjoint field. We have no intention of resolving epistemological debates about the meaning of meaning. Rather, we argue that a practical sociological perspective on meaning implies an actor doing meaning. We therefore suggest, as an analytical convenience, thinking about meaning as an instance of an actor interpreting a stimulus (Strauss & Quinn 1997). Reviewing existing literature, we argue that most sociological accounts effectively understand interpretation either as categorization or as semantic association. We integrate these two perspectives and introduce the categorization–association (CA) model of interpretation. We further show that our model is useful for addressing two foundational questions in the sociology of culture: what it means for people to share an interpretation, and the process by which interpretations are coordinated interpersonally. As we discuss, recent advances in computational linguistics have opened up exciting new possibilities for the study of interpretation. We conclude with a call for sociologists to address interpretative heterogeneity when studying attitudes or behaviors, and we suggest several avenues for future research.

2. FROM MEANING TO INTERPRETATION

Things are not meaningful in and of themselves. Consider medical face masks as an example. As the COVID-19 pandemic was unfolding, the act of wearing a mask became contentious. For

some, mask-wearing represented social responsibility. For others, it was a symbol of government oppression. In the United States, wearing a mask in public soon became a marker of liberal political identity. Consequently, liberals and conservatives adopted different health behaviors, resulting in different mortality rates by political ideology (Shepherd et al. 2020).

Meanings, in other words, are behaviorally consequential, even if these consequences are, luckily, mostly nonfatal. Cultural sociologists are therefore interested in understanding how things acquire meanings that are not directly derivative of these things' ontological features.

A long sociological tradition has been doing this by analyzing what Lizardo (2017) calls public culture, namely, objects such as texts or works of art (Mohr 1998). This approach has the advantage of exploring changes in meaning over time (e.g., Baumann 2007, Knight 2023) but the disadvantage of having to assume that objects have unitary meanings.

If meaning is made, however, rather than given, then there must be someone making it. We conjecture that most sociological accounts of meaning, whether explicitly or implicitly, assume an actor for whom something has meaning. We therefore follow Strauss & Quinn (1997) and suggest thinking about meaning as an instance of an actor interpreting a stimulus. By stimulus, we mean something in the world—an object, an action, a symbol—that is experienced subjectively by the actor.

Our actor-centric focus on interpretation does two things. First, it implies that meaning is not independent of the actor experiencing it. A given stimulus might mean different things to different people. Its interpretation, moreover, can change within a person over time or across different situations (e.g., Altomonte 2020).

Yet, not all acts of interpretation are cultural, even if different actors interpret the same stimulus differently. Different people, for example, might understand their feelings in different ways. This difference becomes sociologically interesting if people's interpretations are systematically shaped by their social experiences, for instance, by reading self-help books (Illouz 2008). We therefore define cultural meaning as an interpretative instance in which the actor's interpretation arose socially.[1] The social process leading people to converge on shared interpretations is often referred to as social construction.

Second, interpretation suggests a cognitive process. This, of course, opens a Pandora's box of its own, as what it means for someone to interpret something is a fraught and complex question. Nevertheless, as an analytical convenience, we define interpretation as the process by which a stimulus is associated with a cognitively represented concept (Hannan et al. 2019). We explore the implications of this practical conceptualization below.

Before we do so, however, it is important that we clearly specify what interpretation is not. Interpretation, as we define it here, is the cognitive process that determines what something is. That is analytically distinct from (*a*) one's attitude toward that thing, and (*b*) one's decision about how to act on it. Because interpretation is difficult to observe, sociologists often infer it by investigating people's attitudes or actions. Yet, people might express similar attitudes, or pursue similar courses of action, even if their interpretations are different. Two individuals, for example, might be discontented by the person wearing a mask next to them. One might feel so because she interprets this person as being sick, the other because she perceives this person as a liberal. **Table 1** lists various attitudinal and behavioral constructs commonly used in cultural analysis that are distinct from but often confused with interpretation.

Meaning: instance of an actor interpreting a stimulus

Cultural meaning: interpretation that arises through social process

Interpretation: the process of a stimulus activating a cognitive representation

[1]Note that we depart from common sociological approaches that define culture as shared interpretation. The fact that a group of people share an interpretation does not necessarily imply that culture is the culprit; they may have converged on similar interpretations due to the stimulus's semantically constraining features. Our definition, in contrast, is explicitly conditioned on a social process.

Table 1 Glossary of cultural constructs that are distinct from interpretation

Type	Construct	Definition and example
Attitudinal	Preference	Orientation toward a stimulus: "I prefer wearing a mask in public."
	Belief	Assumption about the empirical world: "COVID-19 was produced in a lab."
	Value	Abstract perception about what is desirable, often rooted in morality: "I prioritize others' well-being over my comfort."
Behavioral	Norm	Assumption about others' expectations of appropriate behavior: "People wear masks on the train."
	Script	Set of expected behaviors given a situation and one's role in it: "As a train employee, I should wear a mask even if passengers aren't."

How we interpret things often shapes our attitudes and actions. Whether consumers interpret a beer as commercial or craft, for example, impacts their enjoyment from it (Frake 2016). Researchers who fail to account for interpretation therefore risk misunderstanding the attitudinal and behavioral patterns they observe. This is especially the case when people have divergent interpretations (see the sidebar titled Values and Interpretative Heterogeneity). As various studies show, the relationship between sociodemographics and attitudes is often contingent on people's interpretations (e.g., Baldassarri & Goldberg 2014, Bonikowski & DiMaggio 2016).

2.1. Interpretation as a Sociological Concept

If interpretation is the association of a stimulus with a cognitively represented concept, then a theory of interpretation is necessarily a cognitive theory. Indeed, since DiMaggio's (1997) agenda-setting article, American sociology has oriented toward understanding culture as an inherently cognitive construct. Among the various psychological concepts introduced into sociology by this cognitive turn, none has been as influential as schema: a mental structure that represents a concept and its relationships.

Schemas have been immensely helpful conceptual building blocks. Consistent with dominant tool kit (Swidler 1986) and practice (Bourdieu 1990) theories, they allow sociologists to explain how cultural knowledge is stored nondeclaratively (Lizardo 2017), activated automatically (Vaisey 2009), and deployed across varied contexts in inconsistent ways (Wood et al. 2018). Indeed, recent

VALUES AND INTERPRETATIVE HETEROGENEITY

Recent sociological debates on values exemplify the usefulness of analytically distinguishing interpretation from other constructs. Whereas some argue that values are organizing principles structuring behavior across domains (e.g., Miles 2015), others dismiss them as post hoc justifications or analytically imprecise reifications (e.g., Martin & Lembo 2020). These accounts typically assume that values have logically consistent behavioral implications. Vaisey (2009), for example, shows that respondents who espouse a theistic morality are least likely to report engaging in deviant behaviors such as smoking marijuana. Yet, what one deems deviant behavior another might perceive as bravery. If people hold different interpretations, they may pursue different actions even if they espouse similar values. Whereas Americans overwhelmingly value freedom, for example, they vary dramatically in their interpretations of what this abstract ideal implies (Gorski & Perry 2022). Nevertheless, cross-sectional studies on values rarely tap into respondents' interpretations. Consequently, what might seem like empirical support for, or refutation of, the behavioral implications of values might instead be evidence for interpretative heterogeneity.

sociological work on culture has focused on understanding how schematic knowledge is acquired, stored, and retrieved (Cerulo et al. 2021).

But as much as they are a blessing, schemas are also a curse. Sociologists use them as high-level constructs, often conflating multiple aspects of cultural cognition (Leschziner & Brett 2021). In some accounts, schemas are described as tools of implicit categorization. In others, they are used to explain how people understand situations and the behavioral scripts such understandings imply. Attempting to provide a general definition, Boutyline & Soter (2021, p. 737) suggest conceptualizing schemas as "shared representations" that carry "interpretable information." Yet, such a definition inevitably leads one to hit a tautological wall when trying to use it to explain what it means for someone to interpret something.

Categorization: occurs when a stimulus activates a concept

Concept: a prototypical cognitive representation

Category: the set of stimuli that activate a concept

Imagine observing a person wearing a face mask in public. Interpretation occurs when this situation activates a schema. But what does that mean? Your schematic perception might lead you to infer this person is a liberal. Depending on your political ideology, such a determination might further lead you to think this behavior is considerate of others, or is a form of inauthentic virtue-signaling. If the former, you might perceive it positively and be inclined to wear a mask yourself. If the latter, you might refuse to wear one in scorning defiance. All of these different reactions fall under the wide umbrella of schematic cognition.

To eschew this analytical conundrum, we propose a different route. Rather than cognitive fidelity, we prioritize analytical usefulness. We draw on a set of related but mostly independent streams of sociological research to identify two dominant sociological approaches to interpretation: One understands interpretation through the prism of categorization, the second through the prism of associative semantics.

2.2. Interpretation as Categorization

When we experience something in the world, we immediately categorize it. Imagine observing a whiskered, four-legged animal during your morning stroll. Whether that animal is a domesticated cat or a cougar will significantly affect how you react.

Sociologists have long been interested in categorization because what something is perceived to be affects its subsequent treatment. This large body of work understands categorization as the foundation of interpretation, seeing categories as classificatory sieves that people use to "cut out islands of meaning" (Zerubavel 1999, p. 67). It explores how individuals' categorizations affect their evaluations of the people and objects they are categorizing (Delmestri & Greenwood 2016, Flores & Schachter 2018, Lena 2012).

Hannan et al. (2019) distinguish between a concept, which is a cognitive representation, and a category, which encompasses the things in the world that are associated with a concept. How, exactly, humans represent concepts is debated. Nevertheless, it is generally agreed that people store prototypical images of concepts (Murphy 2004). When they categorize a stimulus, they compare its observable features to prototype features. Prototypical perception has two sociological implications. First, it implies that categorical boundaries are porous and that category systems are flexible. Second, it suggests that stimuli vary considerably in how easily classifiable they are.

Two lines of work have studied these implications. The first focuses on symbolic boundaries (Lamont & Molnár 2002). This research stream explores how different people or objects are categorized and how such conceptual distinctions produce, reinforce, or subvert inequalities between groups. The classification of different US ethnic groups as white (e.g., Fox & Guglielmo 2012) or nonwhite (e.g., Maghbouleh 2017), for example, affects whether they are evaluated positively or negatively. These evaluations, in turn, influence group members' access to material resources (Massey 2007).

Semantic association: cognitive association between concepts

A second body of work explores the consequences of crossing categorical boundaries. This research demonstrates that actors who defy categorical codes are evaluated negatively. Whether these are individuals mixing conventionally perceived masculine and feminine behaviors (Lagos 2019, Meadow 2018) or organizations offering services that are typical of different kinds of businesses (Zuckerman 1999), such categorical noncompliance is difficult to interpret. This leads observers to respond to it negatively.

Some actors have more latitude to cross categorical boundaries than others, however (Rao et al. 2005). When high-status individuals defy categorical codes, their behaviors are often perceived as competent and creative, rather than unintelligible (Sgourev & Althuizen 2014). Because they are the greatest benefactors of the status quo, these actors also tend to react most negatively to others' categorically noncompliant behaviors (Goldberg et al. 2016). Yet they also enjoy the power to change it. Thus, when categorical boundaries shift, existing inequalities are often reinforced.

2.3. Interpretation as Semantic Association

How people categorize a stimulus does not fully capture how they interpret it. Liberals and conservatives, as we noted earlier, agree that the protective cloth Americans wore over their faces during the COVID-19 pandemic is a mask. But they disagree about whether this behavior is socially responsible or coercive. This deeper sense of interpretation is what sociologists often refer to when they discuss meaning. Tavory & Swidler (2009), for example, argue that Malawians' reluctance to use condoms in sexual relationships is related to their interpretation of such use as an indication that the relationship is less intimate. Condom use is associated with instrumentality.

We refer to this conceptualization of interpretation as semantic association.[2] Dominant cognitive theories essentially agree that people represent semantic knowledge as networks of association and dissociation between concepts (Patterson et al. 2007). The meaning of a concept is therefore a function of the other concepts with which it is associated or those it is opposed to.

Recent sociological work has paid considerable attention to the cognitive structures underpinning semantic knowledge, debating, for example, whether semantic association should be inferred from automatic or deliberative cognition. This level of detail is unnecessary for our purposes, however. Rather, we note that, irrespective of one's specific cognitive assumptions, the semantic association approach to interpretation is pervasive in sociological research, even if mostly implicitly. For example, a large body of work, using a variety of methods, demonstrates that race (Freeman et al. 2011), leisure activities (Kozlowski et al. 2019), and even scents (Cerulo 2018) and sounds (Schwarz 2015) are infused with class connotations. When qualitative researchers analyze meaning, moreover, they are typically describing semantic associations (Pugh 2013, Tavory & Timmermans 2013). Zubrzycki (2022), for example, argues that those promoting a Jewish revival in Poland are motivated by a desire to break the association between Polishness and Catholicism.

Despite the prevalence of this theoretical orientation, formal analyses of semantic association are relatively uncommon. The cognitive turn in cultural sociology recently catalyzed the use of new methods specifically intended for measuring associative cognition. Two approaches dominate. The first focuses on the semantic representation of a specific concept. Work in this vein often uses the implicit association test to elicit semantic associations (Miles et al. 2019, Shepherd 2011, Srivastava & Banaji 2011).

A second approach maps the network of semantic associations between multiple concepts. Inspired by work by cognitive anthropologists (D'Andrade 1995, Strauss & Quinn 1997), this

[2] We use the term "semantic" to denote cognitive representation relating to meaning, but we do not assume that it is necessarily linguistically encoded.

line of work measures individuals' constellations of semantic associations. Hunzaker & Valentino (2019), for example, explicitly ask respondents to report whether pairs of concepts are associated or not. Goldberg's (2011) relational class analysis infers pairwise conceptual associations from respondents' attitudes.

A network approach to semantic association, as opposed to a focus on a single concept, provides two vantage points. First, it affords mapping the overarching axes structuring people's interpretations. Brensinger & Sotoudeh (2022), for example, demonstrate that some Americans' attitudes toward members of other social groups are organized along partisan lines. Second, it provides an overview of the extent to which concepts are interdependent (DellaPosta 2020). Attitudes toward concepts that are associatively meshed with other concepts are less likely to change (Rawlings 2020).

3. SYNTHESIS: THE CATEGORIZATION–ASSOCIATION MODEL OF INTERPRETATION

The two dominant sociological approaches we identified above—one focused on categorization, the other on semantic association—are not competing conceptualizations of interpretation. Rather, they are complementary, each centering on a different cognitive aspect of interpretation. The former relates to inferences about what is a stimulus, the second to what it is like.

Integrating these two approaches, we propose the categorization–association (CA) model of interpretation. Our dual-phase model understands the act of interpretation as the activation of a cognitively represented focal concept and the set of other concepts with which the focal concept is cognitively associated. As we illustrate in **Figure 1*a***, this concept can be formalized as a node in a network of cognitive relationships. In the first phase, the actor determines what concept the stimulus is an instance of. In the second phase, the focal concept activates a set of associated concepts. These two phases are analytical rather than chronological moments. In reality, categorization and association occur simultaneously.

Because cognitive conceptual systems are continuous rather than discrete, the activation of a cognitive concept can be thought of as a probabilistic event (Hannan et al. 2019). Categorization

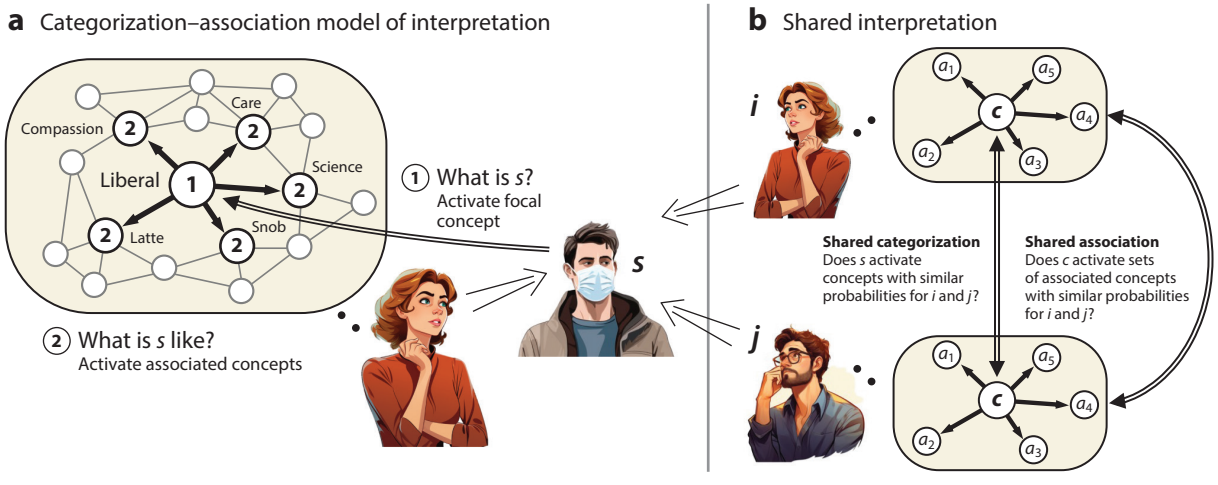

Figure 1

(*a*) The categorization–association model of interpretation of a stimulus *s* (using the example of someone wearing a mask in public), and (*b*) its implication for shared interpretation between individuals *i* and *j*, with respect to the activation of focal concept *c* and the set of associated concepts *a*.

can therefore be construed as a conditional probability P($c|s$), where c is a concept and s is a stimulus. Cognitive associations are similarly probabilistic ties of activation. Association, therefore, is the conditional probability P($a|c$), where c is a focal concept and $a = \{a_1, a_2, \ldots, a_n\}$ is a set of associated concepts. Interpretation, in other words, can be understood as the probability that a stimulus will activate a concept, and the corresponding probability that the activated concept will further activate other associated concepts.

3.1. Implications

Let us return to the example of observing someone wearing a mask in public (as illustrated in **Figure 1*a***). How is this person interpreted? In the categorization phase you would associate this person with a concept. In prepandemic times this concept might be sickness. If this were happening during the height of the COVID-19 pandemic, however, you may have categorized this person as a liberal. In the second phase your categorical determination would have activated additional associations. Depending on your own political identity, you might associate liberalism with compassion or with political correctness.

Our example highlights several important implications. First, the CA model admittedly does reductive violence to the complexities of human cognition. It is likely that categorization and semantic association are not experienced as separate subjective processes. Indeed, from a schematic cognition point of view, the various concepts in our example—mask, liberalism, compassion, political correctness—are all intricately interconnected. They are simultaneously categorical features and semantic associations. However, it is analytically useful, even if empirically reductive, to distinguish the focal concept being activated from the concepts with which it is associated. These relate to different aspects of interpretation.

Second, while our example simplistically illustrates how one focal concept is activated, in reality the same stimulus can activate multiple concepts. It is interpreted at different levels of construal (Trope & Liberman 2010). We may categorize the mask-wearing person as sick, or more specifically as suffering from a respiratory disease. We may also categorize them in terms of gender, race, age and other categorical distinctions that are observable and salient. While these concepts are all schematically interconnected, it is analytically useful to distinguish these different levels and dimensions of construal.

Third, as our example illustrates, categorization and semantic association are related but distinct. Two people might similarly categorize the person as liberal, but semantically associate liberalism with different concepts. Conversely, their disagreement may be rooted in different categorizations, one perceiving the person as sick, the other as performing a political identity. We return to this observation when we discuss the implications for shared interpretation in Section 4.1.

3.2. Application: Unequal Childhoods

To illustrate the CA model, we apply it to Lareau's (2011) influential ethnography of class-based parental cultures. Observing her subjects' daily routines, Lareau finds that while middle-class parents orchestrate their children's schedules and conversationally elicit their thoughts and feelings, working-class parenting is characterized by flexible time management and an emphasis on authority and instruction. These two cultural logics, as she terms them, are ultimately rooted in divergent interpretations of child-rearing.

Working- and middle-class parents differ in their interpretations of parenting in terms of both categorization and semantic association. Whereas middle-class parents categorize leisure time

as a parental responsibility, working-class parents do not. Moreover, while the former associate parenting with reasoning and emotion, the latter associate it with rules and boundaries.

Thinking about these different parenting styles through the prism of interpretation helps one to see why, despite wanting similar things for their children, parents of different class backgrounds adopt persistently different child-rearing practices, resulting in different life outcomes for their children. From a middle-class parent's point of view, failing to oversee leisure is interpreted as neglectful parenting; from the working-class parent's vantage point, however, supervising leisure is perceived as helicopter parenting. Similarly, these parents interpret their counterparts' practices as deficient because they do not associate authority or reasoning, respectively, with parenting. Even if they observe each others' behaviors, their divergent interpretations might lead them to entrench behaviorally. Thus, child-rearing practices are also infused with class connotations, becoming symbolic markers of class identity.

4. SHARED INTERPRETATIONS

So far we have described interpretation at the individual level. But sociologists are, ultimately, interested in understanding it as a collective property. When they describe culture as shared meanings, they refer to people's convergent or divergent interpretations. The CA model provides a framework for analytically describing this property. In this section we draw on the distinction between categorization and semantic association to define shared interpretation, and then to describe how sociological research conceptualizes interpretative heterogeneity.

4.1. What Is Shared Interpretation?

The CA model suggests thinking of shared interpretation as two forms of similarity between individuals: shared categorization and shared association. We illustrate these two constructs in **Figure 1*b*** and elaborate on them below.

4.1.1. Shared categorization. Shared categorization is essential for interaction. Organizational members, for example, need to develop a similar categorization scheme to effectively coordinate when performing tasks (Weber & Camerer 2003). People operating in the same cultural context tend to have similar categorical definitions for most concepts, otherwise they would be unable to interact productively. This is why some of the most compelling evidence for categorization incongruities comes from cross-cultural research (Malt 1995). Nevertheless, categorization disagreements occur frequently within the same cultural contexts. Zuckerman (1999) shows, for instance, that securities analysts disagree on how certain firms should be categorized.

We define shared categorization as whether two individuals categorize a given stimulus in the same way (see **Figure 1*b***). Probabilistically, shared categorization is the extent to which a stimulus activates concepts with identical probabilities for two individuals, denoted i and j. Let $P_i(c|s)$ denote the conditional probability distribution of stimulus s over categories $c \in C$, for individual i. Shared categorization is the similarity between the two individuals' probability distributions, $S(P_i(c|s), P_j(c|s))$.[3]

C denotes a set of concepts. It can refer to the universe of all concepts, or it can be defined as a sociologically relevant subset of concepts. This subset is determined by the context in which interpretation occurs. If, for example, s is a song, the context is music, which can imply that C is the subset of all musical genres. Two individuals will increasingly share categorization in the music

[3] S denotes a similarity function. A common method for measuring similarity between probability distributions is the Jensen–Shannon divergence.

AMBIGUITY AND CONTESTATION

When people disagree about categorization, this can be the result of the stimulus being ambiguous or contested. An ambiguous stimulus is one that falls on the boundaries between concepts and is therefore uniformly difficult to categorize (McMahan & Evans 2018). For example, the singer Taylor Swift, who began her career as a country musician, confused audiences when she first transitioned into pop. A contested stimulus, in contrast, is one where categorization disagreement is rooted in conceptual incongruity between people (Murray 2010). Some conservatives, for instance, may not consider same-sex matrimony marriage because, unlike liberals, their cognitive representation of marriage requires an opposite-sex couple. Here disagreement is rooted in different conceptual representations, not in stimulus ambiguity. At the population level, ambiguous and contested stimuli might exhibit very similar patterns. Imagine a stimulus that has only two plausible categorizations, c_1 and c_2. If it is ambiguous, all people would sometimes categorize it as c_1, and other times as c_2. If it is contested, however, some people will always categorize it as c_1 and others will always categorize it as c_2. In both cases, the stimulus would overall have similar probabilities of being categorized as either c_1 or c_2.

context the more they have identical probability distributions over genres. Shared categorization can be aggregated to the group level as the mean pairwise similarity in categorization between group members, $\frac{2}{N(N-1)} \sum_i \sum_{j \neq i} S(\mathrm{P}_i(c|s), \mathrm{P}_j(c|s))$ (where N denotes the group's size).

Though straightforward, this approach provides limited insight into whether two individuals are in conceptual agreement—namely, whether their cognitive representation of concepts is similar or not. We can only observe the categorical labels people attach to stimuli, not their cognition (see the sidebar titled Ambiguity and Contestation). Ideally, we could infer conceptual agreement if we could observe people's categorization of the full universe of possible stimuli. This, of course, is impractical.

Most research focuses on agreement, or disagreement, on single concepts. Work on American nationalism, for example, finds that Americans disagree on whether being Christian is a typical feature of being American (Bonikowski et al. 2021, Gorski & Perry 2022). Yet shared categorization with respect to a single concept tells us little about the extent to which individuals overall employ similar or different categorization systems. Measuring whether people share categorization systems is complicated by the fact that categorization system disagreement can be the result of people having different cognitive definitions for the same concepts, or different conceptual inventories altogether (Hannan et al. 2019). Some people, for example, might see transgender as a unique category of person, whereas others might perceive transgender individuals through a sex-binary categorical prism. Sociological research comparing shared categorization at the system level, rather than on a single concept, is rare. Recent work has employed creative methods for inferring domain-level categorical boundaries from behavioral data, for instance, by examining the structure of mobility between occupations (e.g., Cheng & Park 2020). Yet this work, by construction, measures systemic structures shared across all actors. As Lynn & Ellerbach (2017) show, people vary systemically in how they categorically perceive occupations.

Shared categorization: the extent to which a stimulus activates concepts with similar probabilities for two or more individuals

A systemic approach to shared categorization implies comparing individuals in terms of not only how they delineate reality into categories but also the extent to which they perceive categorical boundaries as porous or impermeable (Monk 2022). This property is known as categorical contrast (Hannan et al. 2019). Disagreements on categorical contrast often relate to social struggles over access to resources (DiMaggio 1987, Gieryn 1983). Those whose social positions afford them advantages have more incentives to reinforce categorical boundaries sustaining their positions (Goldberg et al. 2016).

4.1.2. Shared association. Shared association is the extent to which a stimulus, conditional on its categorization, activates the same set of associated concepts for different individuals. Let i and j be two individuals, c the concept activated by stimulus s, and a a set of associated concepts. We can define individual i's semantic association as a probability distribution $\mathrm{P}_i(a|c)$ over all possible sets of associated concepts. Shared association between i and j is defined as the similarity in their probability distributions, $S(\mathrm{P}_i(a|c), P_j(a|c))$. This can be aggregated to the group level as the mean pairwise similarity in association between members, $\frac{2}{N(N-1)} \sum_i \sum_{j \neq i} S(\mathrm{P}_i(a|c), \mathrm{P}_j(a|c))$ (where N denotes the group's size). Shared categorization and shared association are analytically distinct. Two individuals might apply the same concept to a stimulus but have different semantic associations for that concept. Conversely, two individuals might use different concepts to categorize a stimulus, but these concepts might have very similar sets of semantically associated concepts.

Shared association: the extent to which a focal concept activates associated concepts with similar probabilities for two or more individuals

Interpretative heterogeneity: the extent to which members of a group have divergent interpretations

Shared association is a more complex construct than shared categorization because it involves comparing sets of multiple concepts. Like categorization, shared association can be conceptualized at the single concept or system level. At the single concept level, it is the degree of overlap between the associated sets activated by a stimulus for two individuals. Guilbeault et al. (2023), for example, use surveys to elicit sets of concepts associated with the concept of leadership. They demonstrate that managers are more likely to be influenced by others who interpret leadership differently from them, as long as they are unaware of that difference. The fact that two individuals have identical associated sets for a focal concept does not mean that their representations of these associated concepts are identical. Measuring shared association is therefore a complex empirical task.

Semantic association can be thought of as a network of associations between concepts. System-level shared association can therefore be operationalized as the similarity between two individuals' networks of association. Measuring similarity between networks is not straightforward. Hunzaker & Valentino (2019) use maximum likelihood consensus analysis to accomplish this task. Relational class analysis compares respondents' preferences using relationality (for a review of relational class methods, see Sotoudeh & DiMaggio 2021).

4.2. Interpretative Heterogeneity

A framework for operationalizing shared interpretation allows us to identify groups of interpretatively like-minded people, what Zerubavel (1999) calls "thought communities." While Zerubavel's account is predominantly focused on categorical perception, the CA model suggests that interpretative communities should be delineated along two dimensions: shared categorization and shared association. Categorization and association communities are not necessarily coterminous. Americans, for example, might agree on how to classify music into different genres but employ competing associative frameworks for interpreting the social significance of music (Goldberg 2011). This is an important distinction. Categorization disagreements are about how to classify something, whereas disagreements about semantic association are about what something is like. There is a difference between disagreeing on whether someone is a scientist, for example, and disagreeing on the cultural authority of science (Gauchat & Andrews 2018).

There are broadly two approaches for mapping the contours of interpretative heterogeneity. A deductive approach presupposes the existence of communities, or groups, usually along ideological (e.g., Hunzaker & Valentino 2019) or ascriptive (Lynn & Ellerbach 2017) dimensions. It evaluates either the extent to which groups share interpretation or the extent to which there is interpretative variation within groups. An inductive approach, in contrast, infers the structure of interpretative groups on the basis of patterns of similarity and difference between individuals. Such an approach necessarily requires employing clustering or dimensionality reduction methods, a methodological challenge warranting its own, separate, discussion. Sociological work overwhelmingly assumes

that the boundaries between interpretative groups follow social boundaries (Pachucki & Breiger 2010). Inductive approaches can identify when this assumption holds and when it does not. They can also evaluate the extent to which there exists patterned heterogeneity within society and when members broadly subscribe to a single interpretative framework (e.g., Boutyline & Vaisey 2017).

Interpretative coordination: the process of interpretative alignment between two or more individuals

Categorical coordination: the process of alignment in conceptual representation between two or more individuals

Associative coordination: the process of alignment in semantic association between two or more individuals

Framing: influencing another person's interpretation of a stimulus

Categorical framing: influencing another person's categorization of a stimulus

Associative framing: influencing another person's semantic association of a concept

5. COORDINATING INTERPRETATION

A common explanation for why people have shared interpretations is that their experiences of the world are shaped by the same environmental constraints. Given our human neurophysiology, for example, it is not surprising that there are universal patterns in color naming shared by all cultures. At the same time, however, differences in color naming across cultures appear to be associated with differences in color perception (Kay & Regier 2006). Explaining how such socially constructed interpretations become shared is perhaps cultural sociology's most important objective. We refer to this process as interpretative coordination and, drawing on the CA model, distinguish between categorical and associative coordination. In this section we survey how sociologists conceptualize these different modes of interpretative coordination at the dyadic and population levels.

5.1. Dyadic Interpretative Coordination

Sociologists call the process by which people affect each other's perceptions social influence. Social influence is commonly conceptualized and modeled as a simple process whereby one person's attitude on a single issue changes in response to another person's behavior. Yet, because people's conceptual knowledge is relational, not discrete, interpretative influence is more complex than is implied by such a model. As Berger & Luckmann (1967) noted half a century ago, people align their interpretations through ongoing interaction. Consistent with the CA model, we propose that people influence each other's interpretations by shaping their conceptual representations and their semantic associations. This leads them to share categorizations and associations.

5.1.1. Framing. To affect others' cognitive representations, people need to be able to influence others' interpretations of a given stimulus. Sociologists commonly refer to this intersubjective process as framing (Benford & Snow 2000, Polletta & Ho 2006). As Wood et al. (2018) note, framing occurs when a communicative act influences the activation of a recipient's schema. Framing, in other words, is a form of public culture, while interpretation is private. Extending this definition in accordance with the CA model, we conceptualize framing as a two-dimensional process whereby the actions of an actor shape, in the eyes of a recipient, a stimulus's categorization (categorical framing), semantic association (associative framing), or both. Importantly, framing need not be conscious or intentional. Whether or not we are purposefully trying to do so, our actions invariably influence others' perceptions.

As an example of such a process, consider Donald Trump's discourse on political correctness. As Schwarz (2023) compellingly demonstrates, Trump reframed this concept both categorically and associatively. His rhetoric broadly expanded the concept's categorical definition when he publicly repudiated acts of norm-compliance, rule-following, and decorum as political correctness. Associatively, his rhetoric construed political correctness as an inauthentic display of moral superiority. By associating political correctness with mask-wearing during the COVID-19 pandemic, for instance, he was able to frame it as a form of strategic virtue signaling rather than true concern for others. This led to interpretative coordination among listeners who gradually adjusted their cognitive representations, both conceptual and associative, accordingly.

Framing does not exclusively occur linguistically. Rather, frames are "assemblages of material objects" (Wood et al. 2018, p. 251). These objects can materialize verbally when people

Table 2 Interpretative coordination: verbal and nonverbal dimensions of categorical and associative framing with examples

Medium	Categorization	Semantic association
Verbal	Labeling: Americans of Latin-American descent are called "Hispanic" (Mora 2014).	Analogy/metaphor: Corporations are analogized to the human body (Knight 2023).
Nonverbal	Relational inclusion: Collaborators of blacklisted artists are assumed to be Communist (Pontikes et al. 2010).	Co-occurrence: Latte drinking co-occurs with eating organic food (DellaPosta et al. 2015).

communicate linguistically, or nonverbally through their behaviors. For example, gift-giving by nightclub promoters frames attractive partygoers' regular attendance as leisure rather than labor (Mears 2015).

Framing, therefore, is a multilayered interpersonal process. It occurs through two different mediums (verbal and nonverbal) and along two interpretative modes (categorization and semantic association). It leads to interpretative coordination when it affects recipients' conceptual and associative cognition. We summarize this process in **Table 2**.

As **Table 2** illustrates, framing happens through four types of actions. Categorization occurs verbally through labeling, namely, when a linguistic category is applied to a stimulus. When people add gendered pronouns to their email signatures, for example, they are engaging in labeling. Nonverbal categorization occurs when stimuli are included or excluded relationally. As structural sociologists have long observed, categorical identities emerge from bounded patterns of interaction (Fuhse 2021). The categorical identity of an object—be it a firm (Podolny 2001) or an artist (Crossley 2015)—is often inferred on the basis of the other objects with which it interacts. As Macy et al. (2019) show experimentally, for example, people learn to categorize unfamiliar policies as Democratic or Republican by observing the partisan identities of others who endorse or reject them.

Semantic association also occurs in two distinct ways. Verbally, association is invoked through connotative rhetoric using analogies and metaphors. Legal language analogizing corporations with the human body, for instance, was pivotal in construing corporations as natural entities independent of the state (Knight 2023). Associations are invoked nonverbally through patterned co-occurrence (Goldberg & Stein 2018). When we observe that people who drink lattes also tend to wear face masks, for example, we learn to associate latte-drinking with liberalism (DellaPosta et al. 2015).

5.1.2. Situational and dispositional coordination. Framing leads to interpretative coordination on two levels. On the first level, it influences the situational interpretation of a given stimulus. Men gaslighting their female partners, for example, often manipulate them into accepting that their emotional expressions are a form of insanity (Sweet 2019). On the second level, it results in long-lasting changes in the recipient's cognitive conceptual and associative representations. Women repeatedly subjected to gaslighting eventually conceptualize insanity to include emotions otherwise categorized as normal. Following Rawlings & Childress (2019), we refer to these levels as situational coordination and dispositional coordination, respectively.

Situational coordination: alignment between two or more individuals' interpretation of a stimulus

Dispositional coordination: alignment between two or more individuals' cognitive representation

5.2. The Diffusion of Interpretation

Framing is the dyadic act through which people coordinate their interpretations. What explains population-level diffusion of interpretations? Sociological research on diffusion has generally disregarded interpretation (Fuhse 2021), analogizing cultural diffusion to viral contagion. Such

an approach, by design, cannot account for the complex nature of interpretation. Recent years have seen an increase in sociological interest in how meanings spread. This work has generally offered two types of answers to the question of interpretative diffusion.

The first takes an ecological approach, demonstrating that interpretative coordination is more likely to occur the more it is resonant (McDonnell et al. 2017). Such resonance can be the product of logical constraints (e.g., Friedkin et al. 2016, Greve et al. 2022) or of consonance with the recipient's existing system of representation (e.g., Cheng et al. 2023, Koopmans & Olzak 2004). In a serial transmission study, for example, Hunzaker (2016) shows that people gradually alter the stories they transmit to increase their cultural consistency.

A second line of work focuses on the role of network structure in shaping patterns of interpretative diffusion. Group experiments in which participants label ambiguous stimuli demonstrate that the speed of categorical coordination increases with network connectivity (Centola & Baronchelli 2015) and that participants converge on similar categorical systems as the size of their network grows (Guilbeault et al. 2021). Network studies of associative coordination are rare. One exception is the associative diffusion model, which explicitly assumes that people influence others' perceived semantic associations (Goldberg & Stein 2018). It shows that populations converge associatively, but diverge in preferences, the more random their underlying network of interaction.

6. MEASURING INTERPRETATION

Sociologists have long recognized that measuring meaning is a central sociological endeavor (Mohr 1998, Mohr et al. 2020). Our actor-centric focus on interpretation implies that it should be measured directly by eliciting people's cognition. We have thus far already reviewed a variety of survey-based and experimental methods designed to elicit categorical and associative perception (e.g., Guilbeault et al. 2023, Hunzaker & Valentino 2019). Space limitations do not permit us to comprehensively review this methodological terrain. Rather, we highlight two recent developments in methods and data availability that have dramatically impacted sociologists' ability to measure interpretation at scale, without having to probe respondents' cognition directly using surveys.

6.1. Categorization: Online Platforms

The first methodological advancement was ushered in by the proliferation of online digital platforms, especially those that serve as marketplaces. Sociologists have mostly used these data for the purpose of understanding categorization processes. This line of work examines how the classification of objects relates to their appeal and popularity (Hsu et al. 2009, Silver et al. 2022). This approach conventionally draws on the classification schemes formulated by platform administrators, effectively assuming interpretative homogeneity among users. A different approach peers more deeply into objects' features. Work on musical consumption, for example, measures song categorical conventionality as the extent to which it includes typical feature combinations (Askin & Mauskapf 2017, Negro et al. 2022). Whether the stimuli in question are products or people, this body of work generally explores the conditions under which atypicality is penalized or rewarded (e.g., Ferguson & Hasan 2013, Leung 2014, Zuckerman et al. 2003).

6.2. Semantic Association: Natural Language Processing

A second, and more dramatic, innovation was made possible by recent advancements in the field of natural language processing. These have enabled sociologists to explore semantic association in ways that hitherto necessitated human-based interpretative analysis. Given the centrality of language for cultural exchange, sociologists were quick to adopt these innovations (Evans & Aceves

2016). Perhaps most important among them was the introduction of word-embedding models (Mikolov et al. 2013). These models are trained on large text corpora and rely on neural network architectures to represent words. Their exact technical implementation is beyond the scope of this review. Importantly, a word-embedding model represents each word contained in the text corpus as a coordinate in a hyperdimensional space. As various applications demonstrated, the distribution of words in this space appears to correspond strongly to how humans represent semantic knowledge (Arseniev-Koehler & Foster 2022). Embedding space, in other words, can be thought of as a representation of conceptual space (Aceves & Evans 2023, Stoltz & Taylor 2021). The distance between words in this space reflects the strength of semantic association between the concepts they correspond to.

A common application of word-embedding models takes advantage of the geometric properties of the embedding space to infer word meanings by projecting them onto specific semantic dimensions. In an early and powerful demonstration of this approach, Garg et al. (2018) inferred occupations' gender connotations by measuring their relative distance to feminine and masculine clusters of words. They demonstrate that occupations' gender connotations trace their historical demographic gender compositions. Taking a similar approach, Kozlowski et al. (2019) show that leisure activities have strong class connotations.

Sociological applications using word embeddings typically train the model on a large corpus of public texts, such as books or speeches. Although they are trained on public culture, they are interpreted as reflections of private semantic cognition. Such an approach comes with significant limitations. It necessarily assumes that all documents contained in the corpus were produced by the same language model—namely, that all the people who produced these texts have identical semantic associations. Such an assumption cannot account for interpretative variability across people. To overcome this limitation, some studies train separate models on different subcorpora. Jones et al. (2020), for example, train separate models on different historical periods to demonstrate a secular decline in gender stereotypical associations since the early nineteenth century. Other studies use word embeddings to describe variability in individuals' semantic combinations (Gouvard et al. 2023, Hofstra et al. 2020, Zhou 2022). More recent approaches, such as that proposed by Rodriguez et al. (2023), adjust word embeddings to different contexts, allowing researchers to identify variability in individuals' semantic association. Such applications hold the promise of measuring interpretative heterogeneity.

Recent advances in computational linguistics also afford the possibility of measuring, at scale, the process of associative interpretative coordination. Card et al. (2022), for example, use contextual embeddings to identify metaphors that dehumanized immigrants or characterized them as immoral in congressional speeches. Text parsers provide the possibility of inferring dependency relationships between words in specific sentences (Stuhler 2022). Knight (2022) uses such an approach to demonstrate an increasing prevalence of agentic metaphors applied to corporate actors in the media.

Most recently, the introduction of large language models—such as GPT-4 (generative pretrained transformer 4)—has opened unprecedented opportunities for measuring interpretative coordination in naturally occurring text. In fact, such models appear to exhibit analogical reasoning (Webb et al. 2023). Although these methods' sociological implications are still in their infancy, they hold the promise of distinguishing categorization from semantic association both as communicative acts of framing and as manifestations of private cognition.

7. CONCLUSION

Thinking about interpretation through the lens of categorization and semantic association distinguishes it from other constructs—such as attitudes and behaviors (see **Table 1**)—all of which

WHAT IS MEANING-MAKING?

Sociologists often discuss meaning-making. Yet, the absence of a consensual definition for meaning leads to confusion about who is doing meaning and what is being made. In some instances, sociologists use the term meaning-making to denote an individual mental process. In others, it is used to describe a social process, denoting either the communicative act of construing a stimulus in a particular way, or the collective interactive process by which two or more individuals converge interpretatively. As we discuss in the main text, we find it useful to distinguish these three different processes with the labels interpretation, framing, and interpretative coordination, respectively. Sociologists, moreover, often use the terms meaning-making and sense-making interchangeably. In Weick's terminology, sense-making is an individual verbal process in which people retrospectively rationalize their experiences by imposing a coherent ordered logic on them (Weick et al. 2005). We suggest that sociologists use sense-making in a manner consistent with Weick's definition.

fall under the wide sociological umbrella of meaning-making (see the sidebar titled What Is Meaning-Making?). Although such a formulation does violence to the messy and complex nature of human cognition, it provides analytical clarity and sketches operationalizable empirical implications. We expect that computational linguistic methods will continue to open up exciting avenues for realizing these implications, influencing sociological work in two fundamental ways.

First, new developments will hopefully facilitate integrating the study of interpretation into sociological fields that have traditionally eschewed it. In particular, we believe sociologists should pay greater attention to interpretative heterogeneity. Consider Americans' attitudes about diversity. As Abascal et al. (2021) show, people reach divergent conclusions about when to categorize a neighborhood as diverse. They might all express support for enhancing diversity but, because they disagree about what this concept means, reach different conclusions about how to obtain it. Alternatively, people may agree about what diversity is but disagree about its desirability. More broadly, it is important that sociologists distinguish when variation in attitudes or behaviors is the result of interpretative heterogeneity and when it is not. This heterogeneity can be conceptualized and operationalized along the dimensions of categorization and association.

Second, the CA model provides a framework for addressing fundamental questions related to cultural change. These include the following:

1. Categorical change: How do new concepts emerge and become conventionally associated with a label? When is a new label mere relabeling of old concepts, and when is it indicative of conceptual change?
2. Associative change: Semantic associations are embedded in complex relational networks. Changes in one association are constrained by and affect other associative links. How do these interdependencies shape the process of associative change?
3. Interpretative coordination: When does situational coordination result in dispositional coordination, and when does it only affect the interpretation of a stimulus without influencing long-term representation? How does this process play out interactionally?
4. Interpretative diffusion: Why are some acts of framing more successful than others in affecting recipients' representations? Does it matter whether this framing is categorical or associative, or are both types of interpretative diffusion subject to the same dynamics?

We look forward to seeing how new methods will be brought to bear on these important questions.

SUMMARY POINTS

1. We suggest thinking about meaning as an actor interpreting a stimulus.
2. Interpretation occurs when a stimulus activates a cognitively represented concept (categorization) and other concepts with which the focal concept is associated (semantic association).
3. Shared interpretation is the extent to which a stimulus activates focal concepts and associated concepts for two or more actors with similar probabilities.
4. Shared interpretation emerges when actors coordinate their categorizations and/or semantic associations through framing.
5. It is useful to distinguish interpretation from attitudes and behaviors.
6. It is important that sociological work on attitudes and behaviors account for interpretative heterogeneity.

DISCLOSURE STATEMENT

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

ACKNOWLEDGMENTS

We thank Jon Atwell, Bart Bonikowski, Julien Clement, Paul DiMaggio, Özgecan Koçak, Danyang Li, Daniel McFarland, Jesper Sørensen, Sameer Srivastava, Paul Vicinanza, and Lara Yang for their feedback on earlier drafts.

LITERATURE CITED

Abascal M, Xu J, Baldassarri D. 2021. People use both heterogeneity and minority representation to evaluate diversity. *Sci. Adv.* 7(11):eabf2507

Aceves P, Evans JA. 2023. Mobilizing conceptual spaces: how word embedding models can inform measurement and theory within organization science. *Organ. Sci.* In press

Altomonte G. 2020. Exploiting ambiguity: a moral polysemy approach to variation in economic practices. *Am. Sociol. Rev.* 85(1):76–105

Arseniev-Koehler A, Foster JG. 2022. Machine learning as a model for cultural learning: teaching an algorithm what it means to be fat. *Sociol. Methods Res.* 51(4):1484–539

Askin N, Mauskapf M. 2017. What makes popular culture popular? Product features and optimal differentiation in music. *Am. Sociol. Rev.* 82(5):910–44

Baldassarri D, Goldberg A. 2014. Neither ideologues nor agnostics: alternative voters' belief system in an age of partisan politics. *Am. J. Sociol.* 120(1):45–95

Baumann S. 2007. *Hollywood Highbrow: From Entertainment to Art*. Princeton, NJ: Princeton Univ. Press

Benford RD, Snow DA. 2000. Framing processes and social movements: an overview and assessment. *Annu. Rev. Sociol.* 26:611–39

Berger PL, Luckmann T. 1967. *The Social Construction of Reality: A Treatise in the Sociology of Knowledge*. New York: Anchor

Bonikowski B, DiMaggio P. 2016. Varieties of American popular nationalism. *Am. Sociol. Rev.* 81(5):949–80

Bonikowski B, Feinstein Y, Bock S. 2021. The partisan sorting of "America": how nationalist cleavages shaped the 2016 U.S. presidential election. *Am. J. Sociol.* 127(2):492–561

Bourdieu P. 1990. *The Logic of Practice*. Stanford, CA: Stanford Univ. Press

Boutyline A, Soter LK. 2021. Cultural schemas: what they are, how to find them, and what to do once you've caught one. *Am. Sociol. Rev.* 86(4):728–58

Boutyline A, Vaisey S. 2017. Belief network analysis: a relational approach to understanding the structure of attitudes. *Am. J. Sociol.* 122(5):1371–447

Brensinger J, Sotoudeh R. 2022. Party, race, and neutrality: investigating the interdependence of attitudes toward social groups. *Am. Sociol. Rev.* 87(6):1049–93

Card D, Chang S, Becker C, Mendelsohn J, Voigt R, et al. 2022. Computational analysis of 140 years of US political speeches reveals more positive but increasingly polarized framing of immigration. *PNAS* 119(31):e2120510119

Centola D, Baronchelli A. 2015. The spontaneous emergence of conventions: an experimental study of cultural evolution. *PNAS* 112(7):1989–94

Cerulo KA. 2018. Scents and sensibility: olfaction, sense-making, and meaning attribution. *Am. Sociol. Rev.* 83(2):361–89

Cerulo KA, Leschziner V, Shepherd H. 2021. Rethinking culture and cognition. *Annu. Rev. Sociol.* 47:63–85

Cheng M, Smith DS, Ren X, Cao H, Smith S, McFarland DA. 2023. How new ideas diffuse in science. *Am. Sociol. Rev.* 88(3):522–61

Cheng S, Park B. 2020. Flows and boundaries: a network approach to studying occupational mobility in the labor market. *Am. J. Sociol.* 126(3):577–631

Childress CC, Friedkin NE. 2012. Cultural reception and production: the social construction of meaning in book clubs. *Am. Sociol. Rev.* 77(1):45–68

Christin A. 2020. *Metrics at Work: Journalism and the Contested Meaning of Algorithms*. Princeton, NJ: Princeton Univ. Press

Crossley N. 2015. *Networks of Sound, Style and Subversion: The Punk and Post–Punk Worlds of Manchester, London, Liverpool and Sheffield, 1975–80*. Manchester, UK: Manchester Univ. Press. 1st ed.

D'Andrade RG. 1995. *The Development of Cognitive Anthropology*. Cambridge, UK: Cambridge Univ. Press

DellaPosta D. 2020. Pluralistic collapse: the "oil spill" model of mass opinion polarization. *Am. Sociol. Rev.* 85(3):507–36

DellaPosta D, Shi Y, Macy M. 2015. Why do liberals drink lattes? *Am. J. Sociol.* 120(5):1473–511

Delmestri G, Greenwood R. 2016. How Cinderella became a queen: theorizing radical status change. *Adm. Sci. Q.* 61(4):507–50

DiMaggio P. 1987. Classification in art. *Am. Sociol. Rev.* 52(4):440–55

DiMaggio P. 1997. Culture and cognition. *Annu. Rev. Sociol.* 23:263–87

Doering LB, McNeill K. 2020. Elaborating on the abstract: group meaning-making in a Colombian microsavings program. *Am. Sociol. Rev.* 85(3):417–50

Evans JA, Aceves P. 2016. Machine translation: mining text for social theory. *Annu. Rev. Sociol.* 42:21–50

Ferguson JP, Hasan S. 2013. Specialization and career dynamics: evidence from the Indian Administrative Service. *Adm. Sci. Q.* 58(2):233–56

Flores RD, Schachter A. 2018. Who are the "illegals"? The social construction of illegality in the United States. *Am. Sociol. Rev.* 83(5):839–68

Fox C, Guglielmo TA. 2012. Defining America's racial boundaries: Blacks, Mexicans, and European immigrants, 1890–1945. *Am. J. Sociol.* 118(2):327–79

Frake J. 2016. Selling out: the inauthenticity discount in the craft beer industry. *Manag. Sci.* 63(11):3930–43

Freeman JB, Penner AM, Saperstein A, Scheutz M, Ambady N. 2011. Looking the part: social status cues shape race perception. *PLOS ONE* 6(9):e25107

Friedkin NE, Proskurnikov AV, Tempo R, Parsegov SE. 2016. Network science on belief system dynamics under logic constraints. *Science* 354(6310):321–26

Frye M. 2012. Bright futures in Malawi's new dawn: educational aspirations as assertions of identity. *Am. J. Sociol.* 117(6):1565–624

Fuhse J. 2021. *Social Networks of Meaning and Communication*. Oxford, UK: Oxford Univ. Press

Garg N, Schiebinger L, Jurafsky D, Zou J. 2018. Word embeddings quantify 100 years of gender and ethnic stereotypes. *PNAS* 115(16):E3635–44

Gauchat G, Andrews KT. 2018. The cultural-cognitive mapping of scientific professions. *Am. Sociol. Rev.* 83(3):567–95

Gieryn TF. 1983. Boundary-work and the demarcation of science from non-science: strains and interests in professional ideologies of scientists. *Am. Sociol. Rev.* 48(6):781–95
Goldberg A. 2011. Mapping shared understandings using relational class analysis: the case of the cultural omnivore reexamined. *Am. J. Sociol.* 116(5):1397–436
Goldberg A, Hannan MT, Kovács B. 2016. What does it mean to span cultural boundaries? Variety and atypicality in cultural consumption. *Am. Sociol. Rev.* 81(2):215–41
Goldberg A, Stein SK. 2018. Beyond social contagion: associative diffusion and the emergence of cultural variation. *Am. Sociol. Rev.* 83(5):897–932
Gorski PS, Perry SL. 2022. *The Flag and the Cross: White Christian Nationalism and the Threat to American Democracy*. Oxford, UK: Oxford Univ. Press
Gouvard P, Goldberg A, Srivastava SB. 2023. Doing organizational identity: earnings surprises and the performative atypicality premium. *Adm. Sci. Q.* 68(3):781–823
Greve HR, Rao H, Vicinanza P, Zhou EY. 2022. Online conspiracy groups: micro-bloggers, bots, and coronavirus conspiracy talk on Twitter. *Am. Sociol. Rev.* 87(6):919–49
Guilbeault D, Baronchelli A, Centola D. 2021. Experimental evidence for scale-induced category convergence across populations. *Nat. Commun.* 12(1):327
Guilbeault D, van Loon A, Lix K, Goldberg A, Srivastava S. 2023. Exposure to the views of opposing others with latent cognitive differences results in social influence–but only when those differences remain obscured. *Manag. Sci.* In press
Hannan MT, Mens GL, Hsu G, Kovács B, Negro G, et al. 2019. *Concepts and Categories: Foundations for Sociological and Cultural Analysis*. New York: Columbia Univ. Press
Hofstra B, Kulkarni VV, Galvez SMN, He B, Jurafsky D, McFarland DA. 2020. The diversity–innovation paradox in science. *PNAS* 117(17):9284–91
Hsu G, Hannan MT, Koçak Ö. 2009. Multiple category memberships in markets: an integrative theory and two empirical tests. *Am. Sociol. Rev.* 74(1):150–69
Hunzaker MBF. 2016. Cultural sentiments and schema-consistency bias in information transmission. *Am. Sociol. Rev.* 81(6):1223–50
Hunzaker MBF, Valentino L. 2019. Mapping cultural schemas: from theory to method. *Am. Sociol. Rev.* 84(5):950–81
Illouz E. 2008. *Saving the Modern Soul: Therapy, Emotions, and the Culture of Self-Help*. Berkeley: Univ. Calif. Press
Jones JJ, Amin MR, Kim J, Skiena S. 2020. Stereotypical gender associations in language have decreased over time. *Sociol. Sci.* 7(1):1–35
Kay P, Regier T. 2006. Language, thought and color: recent developments. *Trends Cogn. Sci.* 10(2):51–54
Knight C. 2022. When corporations are people: agent talk and the development of organizational actorhood, 1890–1934. *Sociol. Methods Res.* 51(4):1634–80
Knight CR. 2023. Classifying the corporation: the role of naturalizing analogies in American corporate development, 1870–1930. *Socio-Econ. Rev.* 21(3):1629–55
Koopmans R, Olzak S. 2004. Discursive opportunities and the evolution of right-wing violence in Germany. *Am. J. Sociol.* 110(1):198–230
Kozlowski AC, Taddy M, Evans JA. 2019. The geometry of culture: analyzing the meanings of class through word embeddings. *Am. Sociol. Rev.* 84(5):905–49
Lagos D. 2019. Hearing gender: voice-based gender classification processes and transgender health inequality. *Am. Sociol. Rev.* 84(5):801–27
Lamont M, Beljean S, Clair M. 2014. What is missing? Cultural processes and causal pathways to inequality. *Socio-Econ. Rev.* 12(3):573–608
Lamont M, Molnár V. 2002. The study of boundaries in the social sciences. *Annu. Rev. Sociol.* 28:167–95
Lareau A. 2011. *Unequal Childhoods: Class, Race, and Family Life, With an Update a Decade Later*. Berkeley: Univ. Calif. Press. 2nd ed.
Lena JC. 2012. *Banding Together: How Communities Create Genres in Popular Music*. Princeton, NJ: Princeton Univ. Press. 1st ed.
Leschziner V, Brett G. 2021. Have schemas been good to think with? *Sociol. Forum* 36(S1):1207–28

Leung MD. 2014. Dilettante or renaissance person? How the order of job experiences affects hiring in an external labor market. *Am. Sociol. Rev.* 79(1):136–58
Lizardo O. 2017. Improving cultural analysis: considering personal culture in its declarative and nondeclarative modes. *Am. Sociol. Rev.* 82(1):88–115
Lynn FB, Ellerbach G. 2017. A position with a view: educational status and the construction of the occupational hierarchy. *Am. Sociol. Rev.* 82(1):32–58
Macy M, Deri S, Ruch A, Tong N. 2019. Opinion cascades and the unpredictability of partisan polarization. *Sci. Adv.* 5(8):eaax0754
Maghbouleh N. 2017. *The Limits of Whiteness: Iranian Americans and the Everyday Politics of Race*. Stanford, CA: Stanford Univ. Press
Malt B. 1995. Category coherence in cross-cultural perspective. *Cogn. Psychol.* 29(2):85–148
Martin JL, Lembo A. 2020. On the other side of values. *Am. J. Sociol.* 126(1):52–98
Massey DS. 2007. *Categorically Unequal: The American Stratification System*. New York: Russell Sage Found.
McDonnell TE, Bail CA, Tavory I. 2017. A theory of resonance. *Sociol. Theory* 35(1):1–14
McMahan P, Evans J. 2018. Ambiguity and engagement. *Am. J. Sociol.* 124(3):860–912
Meadow T. 2018. *Trans Kids: Being Gendered in the Twenty-First Century*. Berkeley: Univ. Calif. Press
Mears A. 2015. Working for free in the VIP: relational work and the production of consent. *Am. Sociol. Rev.* 80(6):1099–122
Mikolov T, Sutskever I, Chen K, Corrado GS, Dean J. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26*, ed. CJC Burges, L Bottou, M Welling, Z Ghahramani, KQ Weinberger, pp. 3111–19. Red Hook, NY: Curran
Miles A. 2015. The (re)genesis of values: examining the importance of values for action. *Am. Sociol. Rev.* 80(4):680–704
Miles A, Charron-Chénier R, Schleifer C. 2019. Measuring automatic cognition: advancing dual-process research in sociology. *Am. Sociol. Rev.* 84(2):308–33
Mohr JW. 1998. Measuring meaning structures. *Annu. Rev. Sociol.* 24:345–70
Mohr JW, Bail CA, Frye M, Lena JC, Lizardo O, et al. 2020. *Measuring Culture*. New York: Columbia Univ. Press
Monk EP. 2022. Inequality without groups: contemporary theories of categories, intersectional typicality, and the disaggregation of difference. *Sociol. Theory* 40(1):3–27
Mora GC. 2014. *Making Hispanics: How Activists, Bureaucrats, and Media Constructed a New American*. Chicago: Univ. Chicago Press
Murphy G. 2004. *The Big Book of Concepts*. Cambridge, MA: MIT Press. Revised ed.
Murray F. 2010. The oncomouse that roared: hybrid exchange strategies as a source of distinction at the boundary of overlapping institutions. *Am. J. Sociol.* 116(2):341–88
Negro G, Kovács B, Carroll GR. 2022. What's next? Artists' music after Grammy awards. *Am. Sociol. Rev.* 87(4):644–74
Pachucki MA, Breiger RL. 2010. Cultural holes: beyond relationality in social networks and culture. *Annu. Rev. Sociol.* 36:205–24
Patterson K, Nestor PJ, Rogers TT. 2007. Where do you know what you know? The representation of semantic knowledge in the human brain. *Nat. Rev. Neurosci.* 8(12):976–87
Patterson O. 2014. Making sense of culture. *Annu. Rev. Sociol.* 40:1–30
Podolny JM. 2001. Networks as the pipes and prisms of the market. *Am. J. Sociol.* 107(1):33–60
Polletta F, Ho MK. 2006. Frames and their consequences. In *The Oxford Handbook of Contextual Political Analysis*, ed. R Goodin, C Tilly, pp. 187–209. Oxford, UK: Oxford Univ. Press
Pontikes E, Negro G, Rao H. 2010. Stained red: a study of stigma by association to blacklisted artists during the "red scare" in Hollywood, 1945 to 1960. *Am. Sociol. Rev.* 75(3):456–78
Pugh AJ. 2013. What good are interviews for thinking about culture? Demystifying interpretive analysis. *Am. J. Cult. Sociol.* 1(1):42–68
Rao H, Monin P, Durand R. 2005. Border crossing: bricolage and the erosion of categorical boundaries in French gastronomy. *Am. Sociol. Rev.* 70(6):968–91
Rawlings CM. 2020. Cognitive authority and the constraint of attitude change in groups. *Am. Sociol. Rev.* 85(6):992–1021

Rawlings CM, Childress C. 2019. Emergent meanings: reconciling dispositional and situational accounts of meaning-making from cultural objects. *Am. J. Sociol.* 124(6):1763–809
Rodriguez PL, Spirling A, Stewart BM. 2023. Embedding regression: models for context-specific description and inference. *Am. Political Sci. Rev.* 117(4):1255–74
Schwarz O. 2015. The sound of stigmatization: sonic habitus, sonic styles, and boundary work in an urban slum. *Am. J. Sociol.* 121(1):205–42
Schwarz O. 2023. Why did Trump call prayers politically correct? The coevolution of the PC notion, the authenticity ethic, and the role of the sacred in public life. *Theory Soc.* 52:771–804
Sgourev SV, Althuizen N. 2014. "Notable" or "not able": When are acts of inconsistency rewarded? *Am. Sociol. Rev.* 79(2):282–302
Shepherd H. 2011. The cultural context of cognition: what the implicit association test tells us about how culture works. *Sociol. Forum* 26(1):121–43
Shepherd H, MacKendrick N, Mora GC. 2020. Pandemic politics: political worldviews and COVID-19 beliefs and practices in an unsettled time. *Socius* 6. **https://doi.org/10.1177/2378023120972575**
Silver D, Childress C, Lee M, Slez A, Dias F. 2022. Balancing categorical conventionality in music. *Am. J. Sociol.* 128(1):224–86
Sotoudeh R, DiMaggio P. 2021. Coping with plenitude: a computational approach to selecting the right algorithm. *Sociol. Methods Res.* 52(4):1838–82
Srivastava SB, Banaji MR. 2011. Culture, cognition, and collaborative networks in organizations. *Am. Sociol. Rev.* 76(2):207–33
Stoltz DS, Taylor MA. 2021. Cultural cartography with word embeddings. *Poetics* 88:101567
Strauss C, Quinn N. 1997. *A Cognitive Theory of Cultural Meaning*. Cambridge, UK: Cambridge Univ. Press
Stuhler O. 2022. Who does what to whom? Making text parsers work for sociological inquiry. *Sociol. Methods Res.* 51(4):1580–633
Sweet PL. 2019. The sociology of gaslighting. *Am. Sociol. Rev.* 84(5):851–75
Swidler A. 1986. Culture in action: symbols and strategies. *Am. Sociol. Rev.* 51(2):273–86
Tavory I, Swidler A. 2009. Condom semiotics: meaning and condom use in rural Malawi. *Am. Sociol. Rev.* 74(2):171–89
Tavory I, Timmermans S. 2013. A pragmatist approach to causality in ethnography. *Am. J. Sociol.* 119(3):682–714
Trope Y, Liberman N. 2010. Construal-level theory of psychological distance. *Psychol. Rev.* 117(2):440–63
Vaisey S. 2009. Motivation and justification: a dual-process model of culture in action. *Am. J. Sociol.* 114(6):1675–715
Webb T, Holyoak KJ, Lu H. 2023. Emergent analogical reasoning in large language models. *Nat. Hum. Behav.* 7:1526–41
Weber RA, Camerer CF. 2003. Cultural conflict and merger failure: an experimental approach. *Manag. Sci.* 49(4):400–15
Weick KE, Sutcliffe KM, Obstfeld D. 2005. Organizing and the process of sensemaking. *Organ. Sci.* 16(4):409–21
Wood ML, Stoltz DS, Van Ness J, Taylor MA. 2018. Schemas and frames. *Sociol. Theory* 36(3):244–61
Wuthnow R. 1989. *Meaning and Moral Order: Explorations in Cultural Analysis*. Berkeley: Univ. Calif. Press
Zelizer VA. 2010. *Economic Lives: How Culture Shapes the Economy*. Princeton, NJ: Princeton Univ. Press
Zerubavel E. 1999. *Social Mindscapes: An Invitation to Cognitive Sociology*. Cambridge, MA: Harvard Univ. Press. Reprint ed.
Zhou D. 2022. The elements of cultural power: novelty, emotion, status, and cultural capital. *Am. Sociol. Rev.* 87(5):750–81
Zubrzycki G. 2022. *Resurrecting the Jew: Nationalism, Philosemitism, and Poland's Jewish Revival*. Princeton, NJ: Princeton Univ. Press
Zuckerman EW. 1999. The categorical imperative: securities analysts and the illegitimacy discount. *Am. J. Sociol.* 104(5):1398–438
Zuckerman EW, Kim TY, Ukanwa K, von Rittmann J. 2003. Robust identities or nonentities? Typecasting in the feature-film labor market. *Am. J. Sociol.* 108(5):1018–73