Cultural Schemas: What They Are, How to Find Them, and What to Do Once You've Caught One*

Boutyline, Andrei
University of Michigan

Soter, Laura
University of Michigan

WORKING PAPER

May 4, 2020

Keywords: cultural schemas, cognitive sociology, automatic cognition, connectionism, Marr's levels

* We would like to thank Rachel Best, Margaret Frye, Alina Arseniev-Koehler, Fabiana Silva, and participants at the University of Michigan Measuring and Modeling Culture Workshop and the ASA Culture and Cognition Network Roundtable for feedback on this manuscript. We are also grateful to Stephen Vaisey for urging us to think of this project as a "how-to guide". Direct all correspondence to Andrei Boutyline at Department of Sociology, University of Michigan, Room 3115 LSA Building, 500 S. State Street, Ann Arbor, MI 48109-1382 Email: aboutyl@umich.edu.

ABSTRACT.

Cultural schemas are a central cognitive mechanism through which culture affects action. In this manuscript, we develop a theoretical model of cultural schemas that is better able to support empirical work, including inferential, sensitizing, and operational uses. We propose a multilevel framework centered on a high-level definition of cultural schemas that is sufficiently broad to capture its major sociological uses but still sufficiently narrow to identify a set of cognitive phenomena with key functional properties in common: *cultural schemas are socially shared heuristic representations deployable in automatic cognition*. We use this conception to elaborate the main theoretical properties of cultural schemas, and to provide clear criteria that distinguish them from other cultural or cognitive elements. We then propose a series of concrete tests empirical scholarship can use to determine if these properties apply. We also demonstrate how this approach can identify potentially faulty theoretical inferences present in existing work. Then, moving to a lower level of analysis, we elaborate how cultural schemas can be algorithmically conceptualized in terms of their building blocks. This leads us to recommend improvements to methods for measuring cultural schemas. We conclude by proposing fruitful sensitizing questions for future scholarship.

Cultural Schemas: What They Are, How to Find Them, and What to Do Once You've Caught One

Following DiMaggio's (1997) immensely influential article about "Culture and Cognition," sociologists have increasingly turned to a cognitively-minded model of how culture influences behavior. Central to this model is the concept of cultural schemas, or "knowledge structures that represent objects or events and provide default assumptions about their characteristics, relationships, and entailments under conditions of incomplete information" (DiMaggio 1997:269). In the decades since DiMaggio's work, this cognitive science-derived understanding[1] of cultural schemas has played a primary role within cognitive sociology of culture, figuring centrally in some of the most prominent studies published within this subfield (e.g., Frye 2017; Hunzaker 2016; Hunzaker and Valentino 2019; Vaisey 2009). Meanwhile, the concept has been used in a vast and growing range of applications across the discipline as a whole, including in the sociologies of religion, gender, race, organizations, and demography, among many other subfields (e.g., Cech and Blair-Loy 2014; Clawson and Gerstel 2014; Edgell 2012; Ray 2019; see overview in Hunzaker and Valentino 2019).

In recent years, engagements with cultural schemas within cognitive sociology have also become increasingly theoretically demanding, with a number of emerging projects focusing on operationalizing or formalizing the concept of cultural schemas, or attempting to discern general principles by which such schemas operate (e.g., Bachrach 2014; Boutyline 2017; Goldberg 2011; Hunzaker 2016; Hunzaker and Valentino 2019; Shaw 2015; Shepherd

[1] Following Wood et al. (2018), we distinguish between the cognitive conceptions of cultural schemas introduced by DiMaggio (1997) and Vaisey (2009) and the older non-cognitive conception from Sewell (1992). In cognitive conceptions, schemas are concrete mental structures internalized by individuals, while for Sewell they are "virtual" objects without "any particular location in space or time" (Sewell 1992:8), meaning they can be both mental structures and public representations. To demonstrate our analytic approach, we briefly discuss incompatibilities between the two perspectives at the end of Part II. We otherwise follow Wood et al. in restricting our scope to cognitive conceptions.

and Marshall 2018). This wellspring of interest brings the promise of a cumulative research enterprise capable of developing a body of general scientific knowledge regarding the structure, function, and distribution of cultural schemas, which has the potential to substantially benefit many empirical applications of the cultural schema concept in sociology, as well as creating a long-sought avenue by which sociological insights about culture can come to bear on the interdisciplinary study of cognition[2] (e.g., Cerulo 2010; Lamont et al. 2017; Lizardo 2014). This rise of demanding new uses of the cultural schemas concept, coupled with the continued importance of cultural schemas to sociology at large, makes this a fitting time to reexamine the theoretical foundation of this research program. That is the task we undertake in this paper.

Our Goals

Throughout our investigation, our primary perspective is pragmatic: we focus on how our theoretical conception can best support empirical work on cultural schemas, including inferential uses that draw on accumulated knowledge about schemas to make novel theoretical inferences about the analyst's domain; sensitizing uses that use the concept to pose empirical questions and provide an interpretive lens for observed phenomena; and operational uses that construct empirical measures aimed at finding schemas in observed data. We note the strengths and limitations of existing conceptualizations that affect their ability to carry out these functions. We then develop an alternate conceptualization aimed at best sustaining empirical research across these different uses. To further support this research, we draw out detailed implications of our conceptualization for empirical and methodological scholarship.

Perhaps the main analytical utility of "cultural schemas" for sociologists comes from their powerful relationship to cultural behavior. As we detail below, cultural schemas can fill in

[2] We detail this in Part I below.

unobserved or forgotten details with cultural assumptions; conceal "irrelevant" variability between objects or people by presenting them as instances of one shared category; and link situations to taken-for-granted normative prescriptions. Crucially, because internalized schemas can be invoked automatically, quickly, and implicitly, they can have these effects *outside an individual's control or even awareness*. They thus cannot be simply "turned off" like volitionally controlled "cultural tools" (Frye 2012; Vaisey 2009; Wood et al. 2018). This relatively direct connection to behavior is what lets cultural schemas serve as an important explanatory mechanism in contemporary accounts of culture's effects on action.

A major goal of our work is to guarantee that the cultural schemas concept can robustly play this explanatory role. This requires reasonably clear limits on what counts as a cultural schema (i.e., a clearly delimited denotation)—because if the bounds of the concept are excessively unclear, then it is also unclear what properties cultural schemas share. As we discuss below, scholars in other disciplines note that the concept of schemas has historically faced the problem of accumulating too many conflicting meanings to remain analytically useful (Ghosh and Gilboa 2014). One telltale sign that this problem may be affecting sociological applications is the frequent usage of "cultural schemas" as a generic term roughly synonymous with other loosely defined terms for "element of culture" like logics, models, or discourses.[3] If the concept can be used interchangeably with a long list of others, it contributes little to the work that uses it.

An overly ambiguous conceptualization can also lead to faulty theoretical inferences, where objects that are not cultural schemas are incorrectly inferred to possess the properties of schematic cognition. For example, we demonstrate below that the common sociological practice of using DiMaggio's (1997) and Sewell's (1992) conceptions of cultural schemas

[3] For a prominent example of this common practice, see Wilde (2004:581).

interchangeably (e.g., Ray 2019)—even though they describe different cultural constructs—can lead to incorrect conclusions. An important part of our contribution is thus the development of a conception of cultural schemas that has unambiguous bounds. We then elaborate this concept's theoretical entailments, and provide concrete empirical or conceptual tests that applied scholars can use to verify whether their objects of investigation are indeed cultural schemas. We also ascertain that our conception focuses analysts' attention to theoretically fruitful aspects of empirical phenomena. We further use our clarified conceptualization to investigate how schemas are algorithmically conceptualized in recent methodological work. We note important differences between conception and measurement and offer suggestions for how this operationalization can be improved. We conclude by proposing sensitizing questions for future empirical work.

Plan of Analysis

The conception we develop here uses two prominent existing views of schemas as its starting points. While most sociological work continues to use DiMaggio's "default assumptions" definition, a growing number of cognitive sociologists have instead begun to conceptualize cultural schemas as networks of *implicit associations* accrued through experience (see Table 1). This view was first prominently introduced to sociology by Vaisey (2009), who argues that it more accurately depicts the cognitive operation of schemas (see also Wood et al. 2018). Our work is the first to closely investigate the relationship between these two conceptions.

[Table 1 about here]

Following Vaisey (2009), many sociologists view the two conceptions as conflicting. However, when we trace the "implicit associations" view to its origins in formal connectionist modeling (Rumelhart, McClelland, and PDP Research Group 1986), we find that it describes the same stereotype-like cognitive structures referenced by the "default assumptions" conception,

but at different levels of analysis (Marr 1982). "Default assumptions" depicts cultural schemas on the functional level, whereas "implicit associations" depicts them on the algorithmic level. As we detail below, functional- and algorithmic-level descriptions best fit different kinds of investigations: the former answers the question "what do schemas do", and thus supports sensitizing and inferential uses, and the latter "what are schemas composed of," and thus enables operational uses that capture schemas by detecting their individual components. Since both sets of uses are key parts of the research program, the sociological conceptualization of cultural schemas needs to similarly operate on both levels of analysis.

The framework we propose below is such a multilevel conceptualization (see Figure 1). We develop it in three parts. In Part I, we describe the general state of schemas research in cognitive science. We then lay out Marr's (1982) "levels of analysis" framework to clarify the relationships between the different levels of our conceptualization.

[Figure 1 about here]

Part II is poised on the functional level. Here, we note that the "default assumptions" conception describes one important kind of schema (stereotype-like structures that operate through pattern completion,) but is not broad enough for many other kinds of cultural schemas analyzed by sociologists. We then propose a functional-level definition of cultural schemas that fits tightly around all of the major sociological uses and focuses attention on their theoretically central functional properties: *cultural schemas are socially shared heuristic representations deployable in automatic cognition.* We explain our choice of each term in this definition, and demonstrate how together they delineate a set of phenomena of special interest to sociologists. We also briefly outline some promising areas for empirical investigation suggested by this conception. We then provide a detailed guide to the properties of cultural schemas identified by this definition, and suggest specific empirical and conceptual tests scholars can directly

apply to their work. To demonstrate our approach, we conclude Part II by using some of these conceptual tests to interrogate Sewell's (1992) commodity schema, which Sewell used as a key example of his notion of "deep schemas." These analyses suggest that Sewell's cultural conception of schemas is partly incompatible with the cognitive science-derived conception, sometimes leading to faulty theoretical inferences in work that uses the two conceptions interchangeably.

Finally, in Part III, we turn to the algorithmic level, which is most directly relevant to schematic measurement. We begin by closely examining the "implicit associations" conception of schemas. We show that, in its most common phrasings, this conception is so broad that it could be used to describe nearly any piece of learned cognitive contents. It is thus insufficiently precise to support its methodological and theoretical uses. However, we demonstrate that a seemingly similar connectionist conception of schemas as networks of *meaningful* entities interconnected by pairwise implicit associations has a substantially more specific denotation. We detail the formal properties of these networks to clarify the implications of this algorithmic-level account for schematic measurement. Our analysis reveals important differences between theory and measurement in existing work, and points to promising avenues for improvement.

We conclude our manuscript by proposing sensitizing questions that applied scholars may be able to productively investigate within their domains.

PART I: DRAWING ON COGNITIVE SCIENCE

Since the cognitive conception of cultural schemas has its origins outside sociology, it merits asking why we do not simply import the conceptualization from cognitive science instead of developing our own. The simple answer is that it would be far from obvious which conceptualization to import. Indeed, though the cognitive science literature on schemas has

been thoroughly reviewed for sociological audiences, existing sociological treatments have not adequately stressed the fact that this multidisciplinary research program never successfully developed a common theoretical core or a single agreed-upon definition of cognitive schemas. Following the concept's emergence near the birth of modern psychology (Bartlett 1932), its re-popularization in cognitive psychology and artificial intelligence in the 1970s (Mandler and Ritchey 1977; Rumelhart and Ortony 1977), and its subsequent booming popularity in linguistics, education, anthropology, and other fields, "schema" came to refer to so many different kinds of phenomena that Mandler (1984) noted that "The phrase [schema theory] itself is perhaps misleading [...] A more accurate phrase might be *schema framework*, since the principles subsumed under this view of the mind consist of very general beliefs about how this form of organization works" (Mandler 1984:1). A decade later, D'Andrade echoed Mandler, adding that "the term 'schema theory' is a little grandiose" (D'Andrade 1995:126).

More recent appraisals of the trajectory that schemas research took after Mandler's writing have been more pessimistic. As the number of different uses of the schemas concept continued to grow, there became fewer and fewer things of interest that one could say about *all* cognitive schemas. As a result, as van Kesteren and colleagues (2012:212) note, "Enthusiasm for schema research waned since the 1980s, partly because of the overextended definition of schema that arose from the explosion of interest". Ghosh and Gilboa (2014) concur, and interpret this decline as a cautionary tale against using a definition of schemas without a clearly delimited denotation. One motivation behind our present work is the desire to save schema research in cognitive sociology from falling victim to the same overextension.[4]

[4] Cognitive sociologists appear to generally be unaware of the decline of this research program. One reason for this may lie in the continued popularity of "schemas" as a pseudo-technical term that cognitive scientists use interchangeably with other generic concepts like "representations" or "mental models" (e.g., Tenenbaum et al. 2011).

While the 1980s saw an end to any unified research program into cognitive schemas writ large, there has continued to be plenty of productive research into different *kinds* of schemas across various subfields in cognitive science. Recent literature examines body schemas (D'Angelo et al. 2018), problem solving schemas (Hodnik Čadež and Manfreda Kolar 2015), memory schemas (Ghosh and Gilboa 2014), action schemas (Frijda 2010), image schemas (Mandler and Cánovas 2014), and narrative/story schemas (Kahan 2015), among many others. These research programs are diverse enough to share little in common beyond their recognition of various kinds of general, abstracted, flexible mental representations. What enables each of these "schema subtypes" to serve as focal concepts for successful research enterprises is a *clear delineation of each field's specific object of study*. The research program into cultural schemas in cognitive sociology has reached a point of maturity where it requires the same kind of delineation.

A search through major cognitive science journals reveals recent work on many different types of *cognitive* schemas, but a near-complete absence of work on *cultural* schemas. Cultural schemas received little systematic attention in cognitive science after the decline of the 'cultural models school' of cognitive anthropology (e.g., D'Andrade 1995; Quinn 2011; Strauss and Quinn 1997). Cognitive sociologists may thus be doubly amiss to defer to a 'schema theory' from cognitive science: first, because by treating these disconnected research programs as a single coherent source of knowledge, sociological work risks recapitulating the excessive breadth that doomed the original research program into cognitive schemas—a perennial threat to work that uses this concept—and second, because this practice conceals an opening in the interdisciplinary cognitive science that cognitive sociologists are well-positioned to fill. Our treatment is aimed in part at developing a conception of cultural schemas that would best support such sociological research.

Cognitive Meta-Theory: Levels of Analysis

We organize our treatment around Marr's (1982) concept of "levels of analysis", which is an influential meta-theoretical statement on cognitive explanation (see also Foster 2018; Brighton and Gigerenzer 2008). The crux of this framework is the idea that it is possible to provide qualitatively distinct descriptions of the same cognitive process that are simultaneously correct but useful for different tasks because they answer different questions about the process.

At the highest level are *functional* or *behavioral* explanations[5] in terms of the consequences of a process for the organism's behavior (or for "downstream" cognitive processes). We can illustrate this level by analogy to an article printed in an old-fashioned physical newspaper. On this level, a newspaper article is a vehicle allowing a small number of observers (journalists) to quickly convey information about important events to a vastly larger public without requiring them to be present in the same place and time. Most sociological research takes place on this level.

Next is the *representational/algorithmic* level, which is focused on providing parsimonious, analytically fruitful models of the computational task underlying the process. This level usually involves an analytic separation between cognitive contents (*representations*), and the operations performed on them (*algorithms*.) On this level, newspaper articles are narrative descriptions of events presented as written text. The two main algorithms that allow these representations to fulfill their functions are (i) writing, which encodes a narrative into written text; and (ii) reading, decodes the text back into a comprehensible narrative. Connectionist models of cognition and radial (prototype- or exemplar-based) views of categories are two examples of algorithmic-level models frequently referenced in cognitive

[5] Marr confusingly termed this level "computational." Like Brighton and Gigerenzer (2008), we instead adopt more intuitively interpretable labels for the levels.

sociology. The majority of cognitive science research has been poised at this level (Thagard 2005).

Finally, it is possible to describe the same cognitive system in terms of its *biological* or *physical* implementation. On this level, the newspaper article consists of ink on wood-pulp paper. Explanations in terms of brain regions or neural wiring are poised at this level, which is most closely associated with neuroscience (and most distant from sociology). For simplicity, we will often refer to the three levels by the shorthand *functional, algorithmic,* and *biological*, respectively.

This framework can help us make sense of which sociological applications the two conceptualizations of schemas best support. Many applied users are primarily interested in answers to two questions: *what schemas do* — i.e., how they affect behavior — and *what schematic cognition looks like* — i.e., how schemas can be observed in an empirical setting. These questions stand firmly on the functional level. The conception introduced by DiMaggio (1997) and elaborated by Cerulo (2010) offers ready answers to these questions because it is based around a functional-level definition of schemas as "providing default assumptions" (DiMaggio 1997:269) that "allow us to infer what [places, people, or objects] do, where they fit, and what to expect of them" (Cerulo 2010:117). Since this conceptualization is poised on the level of human behavior, it clearly suggests how schemas can be spotted "in the wild"—for example, when a waiter chooses whether to hand the bill to the male or female diner. This direct applicability to many sociological uses may help explain why this conception made cultural schemas such a ubiquitous part of the sociological lexicon.

In contrast, the newer conception of schemas as "largely unconscious networks of neural associations" (Vaisey 2009:1686) that are "developed via repeated embodied experience (i.e., perceptual, sensorimotor, interactional) and stored in long term memory"

(Wood et al. 2018:246) is poised at the algorithmic/representational level, and thus does not provide straightforward answers to these functional-level questions. For example, it is far from clear how these neural associations could be identified when observing a social setting. However, as a description of an algorithmic/representational scheme, it instead answers the lower-level, cognitive-scientific question about *what schemas consist of.* This makes it usable as the basis of various measurement techniques that aim to capture schemas by detecting their building blocks (Hunzaker and Valentino 2019; Shepherd and Marshall 2018; Boutyline 2017; Goldberg 2011; Schröder and Thagard 2013; Tsoukalas 2006). Such operational uses are a fundamentally important advance towards a robust research program into cultural schemas.

Neither current conception thus fulfills all the needs of sociological research if taken alone. To meet both sets of needs, the conceptualization we develop here is poised on both levels. On each level, we begin with an existing conception of schemas, and ask how it can be improved to best serve the needs present at that level of analysis. On the functional level, we develop a conception that provides clear guidance about what schemas are, what they do, and how they can be spotted "in the wild." On the algorithmic level, we instead aim for a conception that clarifies which components of schemas should be measured. On each level, we leverage our conception to produce concrete suggestions for empirical and methodological work.

PART II: FUNCTIONAL LEVEL

Existing sociological uses

The most commonly used conceptualization of cultural schemas in sociology is DiMaggio's (1997) functional-level treatment. Drawing on an interdisciplinary research program on cognitive knowledge organization and higher-order memory structures in cognitive psychology, cognitive linguistics, artificial intelligence, and—most directly—cognitive anthropology (D'Andrade 1995), DiMaggio defines schemas as "knowledge structures that

represent objects or events and provide default assumptions about their characteristics, relationships, and entailments under conditions of incomplete information" (DiMaggio 1997:269; see also Cerulo 2010). We refer to this as the "default assumptions" conception of schemas (see Table 1.)

The "default assumptions" conceptualization is a clear and parsimonious depiction of schemas as *pattern completion engines* that fill in missing pieces of knowledge with culturally learned defaults. The following story segment from Rumelhart (1980:43) demonstrates such schemas in action:

> "Business has been slow since the oil crisis. Nobody seemed to want anything really elegant anymore. Suddenly the door opened and a well-dressed man entered the showroom floor. John put on his friendliest and most sincere expression and walked towards the man."

Although this passage did not explicitly reference automobiles, many readers effortlessly recognize that it is set in a car dealership (which may sell large U.S.-made sedans); John is a salesman; and the well-dressed man is a customer. The relevant schemas thus fill in missing information, letting us perceive the disjoint sentences that make up this vignette as parts of a meaningful gestalt whole (D'Andrade 1995). This pattern completion closely fits the operation of stereotypes, which are perhaps the most sociologically central type of schemas.

However, when considered against the breadth of sociological uses of the schemas concept, this conception appears too narrow. Indeed, even in his original treatment, some examples DiMaggio (1997) uses to illustrate "default assumptions" do not fit the literal text of the definition: e.g., schemas used for categorization, heuristic reasoning, or routine action. Sociological work also frequently focuses on schemas which we demonstrate below do not fit the "default assumptions" conceptualization—for example, category schemas used in implicit classification into product categories, genders, or races (Kovács and Hannan 2010; Lewis 2003;

Ridgeway and Correll 2004); and heuristic schemas that provide automatic rules of thumb for appropriate behavior in certain situations, whether as normative heuristics used in determining correct, moral, appropriate, or ideologically prescribed behavior, or as cultural "recipes" for achieving certain goals (Blair-Loy 2001; Frye 2012; Martin and Desmond 2010; Vaisey 2009; Vaisey and Lizardo 2010).

The mismatch between these examples and the "default assumptions" definition can be observed by noting that both categorization and heuristics have key properties that cannot be reduced to pattern completion without introducing undue overextension. A defining feature of heuristics is an asymmetrical structure linking *cues* to *judgements* or *behaviors* that should take place if the situation matches the cue (Kruglanski and Gigerenzer 2011; Evans 2008:261), e.g., *cue → judgement,* or *situation → appropriate action*. With moral schemas, this can be observed in the following vignette:

> "Jennifer works in a medical school pathology lab as a research assistant. The lab prepares human cadavers that are used to teach medical students about anatomy. The cadavers come from people who had donated their body to science for research. One night Jennifer is leaving the lab when she sees a body that is going to be discarded the next day […] When she saw a body about to be cremated, she thought it was irrational to waste perfectly edible meat. So she cut off a piece of flesh, and took it home and cooked it" (Haidt, Björklund, and Murphy 2000:20).

For many readers, this vignette triggers immediate revulsion, i.e., an emotional experience of this action being "wrong." This revulsion is not an unobserved aspect of the event. Rather, it is an automatically triggered response—a prescription for the perceiver rather than a description of the perceived phenomenon. The moral schema this vignette activates thus cannot be accurately described as "filling in incomplete information with default assumptions."

Implicit categorization similarly does not work through default assumptions. While the purpose of stereotype-like schemas is *filling in* missing information, the first function of category schemas is *ignoring* information that is irrelevant: placing objects in categories

simplifies the perceived world because it allows different objects to be treated as equivalent (Goldstone, Kersten, and Carvalho 2018; Rosch 1978; Zerubavel 1996). This process operates primarily by holistically matching the stimulus to cognitive representations of prototypical category members (Murphy and Hoffman 2012). The activation of an implicit category may eventually result in the activation of a stereotype; however, it could also simply lead to a stimulus being lumped together with other observations and thus largely ignored. Categorization is thus a substantively different process in both function and algorithm.

These examples suggest that there is a need for a broader functional-level conception of schemas of which "default assumptions" would be a subset. But just how broad should this definition be? Clearly, it would be unproductive to make the definition be coterminous with everything that sociologists have described by this concept, partly because some scholars use the term "cultural schema" as a shorthand for any piece of personal culture. A concept this broad, however, would carry little theoretical weight because the class of phenomena it identifies would have few interesting theoretical properties in common.

Luckily, the bulk of sociological uses of schemas do not require such overextension. Specifically, automatic stereotypes, normative heuristics (including social norms, cultural prescriptions, and ideological rules of thumb), implicit categories, and most other cultural schemas empirically investigated by sociologists *do* have key properties in common that set them apart from (i) other pieces of culture, (ii) other types of cognitive schemas, and (iii) other types of non-declarative culture. First, automatic stereotypes, normative heuristics, and implicit categories—which we will refer as our "central examples"—differ from discourses, narratives, and other consciously deployable cultural "tools" in that they are pieces of non-declarative (automatic) culture (Vaisey 2009; Lizardo 2017; Lizardo and Strand 2010; Wood et al. 2018).

Second, unlike many other kinds of learned cognitive contents, these central examples are widely shared. Because many sociologists take the socially shared character of cultural schemas as given, they may not appreciate how this distinguishes them from the schemas at the focus of other cognitive investigations. For example, consider the memory schemas studied in cognitive neuroscience of memory. In their programmatic statement, Ghosh and Gilboa define memory schemas partly through their adaptability, as memory structures that are "constantly developing, affected by every incoming sensory experience" (2014:108). For this reason, they explicitly leave out cultural constructs such as concepts, categories, and cultural scripts from their definition of memory schemas because they remain relatively fixed through usage.

In contrast, social categories, stereotypes, and cultural scripts are central examples of the kinds of schemas sociologists study. Their durability often makes them *more* rather than *less* sociologically interesting (e.g., Vaisey 2009; Vaisey and Lizardo 2010). More broadly, a cultural schema can become shared across a large population only if its key parts remain fixed across time periods and transmission events (Sperber 2006). To our knowledge, the specific focus of cognitive sociologists on durable, externalizable schemas distinguishes our field of investigation from that of investigations into other types of schemas in other areas of cognitive science.

Finally, these central examples are instances of "higher-order mental representations" (Evans 2008:259), as opposed to structures used in "lower-order" cognition (see also Railton 2017; Ryder 2009; Thagard 2012). This distinction becomes clear when we contrast them with other forms of non-declarative learning that occur during shared socialization, many of which are non-representational. For example, the autonomic neuroscience literature on acclimatization to cold climates documents learned changes in metabolic heat production, skin

blood flow, resting oxygen consumption and hormonal production, among other changes (Castellani and Young 2016). This kind of directly embodied learning is clearly different from higher-order learning in the more abstract "representation-hungry" (Clark and Toribio 1994) cognitive domains where stereotypes, categories and heuristics reside. A key difference between such lower-level learning and these higher-order cultural constructs is that the latter, but not the former, *carry interpretable information about the world*, which qualifies them as having *representational content* (Sperber 2006). We discuss this concept in detail below.

These commonalities indicate there is room for a productive conceptualization of schemas that is broader than the letter of the "default assumptions" definition but substantially narrower than "component of non-declarative culture." We thus propose a twofold change to the way sociologists conceptualize cultural schemas on the functional level. First, while the letter of DiMaggio's "default assumptions" definition does not need to change, sociologists should read it strictly and literally, i.e., as referring to stereotype-like schemas that operate through automatic pattern completion. Second, in the next section, we propose a new parsimonious functional-level conceptualization of schemas that is broad enough to cover the three major types of uses central to sociologists, but still restrictive enough that identifying something as a schema carries inferential weight.

Our Functional-Level Conceptualization

In light of these considerations, we propose this functional-level definition:

cultural schemas are socially shared heuristic representations

deployable in automatic cognition.

Table 2 provides a diagrammatic guide to the three criteria contained in this definition: sharedness, automaticity, and representational character.

[Table 2 about here]

We will now explain each term in this definition, and detail what its inclusion accomplishes. We then develop further implications of this conception in the following sections. To help ground our discussion, we begin with an imagined scenario of schematic cognition, featuring an adolescent masculinity schema that could be roughly phrased as "boys shouldn't work hard for their grades" (cf. DiPrete and Buchmann 2013):

> Billy, an eighth-grade student, is sitting at his desk while a teacher is returning graded homework essays. As the teacher hands him his assignment, she exclaims, "Great job, Billy! This must have taken a lot of work!" Billy quickly tells the teacher that she is mistaken and that it took him only a few hours. Although Billy had, in fact, spent all weekend working on the assignment, *he is intuitively aware that this hard work is in violation of the image of indifference* he and the other boys try to project in their relations with the school. To Billy's horror, the teacher responds with "No way! This is really well written, and you even wrote three pages more than you were required. You must have spent all weekend on this. You should be proud!" before turning her back to him to move on to the next student, oblivious to the quiet snickering that her praise elicited in the boy sitting next to Billy.

"Automatic cognition". In this scenario, the schema in question is first activated when Billy hears the teacher praise him for his hard work. The speed with which Billy understands the teacher's praise as a violation of masculine norms, and the fact that Billy spots it without intentionally searching for such violations are both characteristics of *automatic cognition*[6]. As Railton (2006) points out, because fluent conversation is highly taxing on resources required for declarative cognition, many important interactional norms are likely largely automatic in character.

Like other recent sociological treatments[7], we make automaticity central to our definition because it accounts for one of the primary reasons cultural schemas are important to

---

[6] Although verbal behavior is paradigmatic of deliberative cognition, conversation is also backed by many fast automatic processes (e.g., Baars and Franklin 2003).

[7] While automaticity is central to most recent conceptions of cultural schemas, this is not the case for many older accounts. For example, in D'Andrade's (1995:122-149) classic 27-page treatment of cultural schemas, automatic cognition is not referenced until the 23rd page—and even then it is unclear whether D'Andrade is describing *all* cultural schemas as automatic, or only a portion. Shore's (1998) treatment of "foundational schemas" is similarly ambiguous. Automaticity does not feature in Shore's definition (1998:53). While Shore appears to suggest that some

sociologists: schemas' close relationship to action. Cultural schemas can influence behavior in at least two relatively direct ways: (i) as implicit categories, stereotypes, and other sets of automatic "default assumptions", they influence what people perceive, remember, or understand others to be saying; and, (ii) as internalized norms, moral intuitions, and habitual cultural prescriptions, they directly produce action through reflex-like cognitive processes that occur without individuals' conscious awareness and control (e.g., DiMaggio 1997; Frye 2017; Gorman 2005; Hsu and Grodal 2015; Lizardo 2017; Sperber 1996; Vaisey 2009; Wood et al. 2018). Cultural schemas thus provide an avenue by which culture can (i) alter the inputs to conscious thought, or (ii) produce behavior without any conscious decision making. This makes them more closely linked to action than discourses, justifications, memorized facts, and other consciously deployable "cultural tools" that influence behavior only by "providing resources" for action (Swidler 1986:273; Vaisey 2009). This is exactly the distinction between Type I "automatic" or "implicit" cognitive processes that are fast, reflexive, unconscious, and nondeclarative, and Type II "explicit" cognitive processes that are slow, reflective, controlled, conscious, and declarative (Evans 2008; for extensive sociological treatments, see Lizardo et al. 2016; Lizardo 2017).[8] As we detail below, the Type I character of cultural schemas has far-reaching implications for what they can and cannot do.

---

foundational schemas are automatic—or at least lack a conscious reality (1998:70)—he does not appear to believe that they must *always* be automatic (e.g., 1998:368). Some of his central examples of foundational schemas are also incompatible with automaticity. For example, in Shore's analysis of "walkabout" narratives that underlie coming-of-age rituals and other cultural rites, he claims that the *narratives themselves* "serve as foundational schemas" (1998:212). Yet as we discuss below, narratives are a central example *non*-automatic cognition. More broadly, a fundamental property of foundational schemas – the fact that they underlie or organize a wide range of different cultural domains – is itself in conflict with many treatments of automatic cognition, which (as we detail below) is often characterized as domain-specific.

[8] Importantly, the Type I character of schemas implies that all schemas *can* function without a person's control and awareness, but it does not imply that all schemas *must always* function in this manner. Many kinds of automatic cognition are available to conscious introspection (Evans 2008). Additionally, all Type II processes involve the backing of some Type I processes, which means some Type I processes can be invoked volitionally. The automatic character thus does not necessarily imply a complete absence of awareness or control, although it implies that schemas should not *require* it to function.

"Socially shared". Social sharedness of schemas refers to their acquisition from other people—be it directly from other living human beings, or indirectly via some form of media. We can surmise that Billy's model of masculinity is *socially shared* because the norm violation is readily recognized by both Billy and the snickering student. In fact, it appears likely that anything that can properly be called a social norm must be socially shared. Sharedness is central to our definition because it delineates *cultural* schemas from non-cultural ones (Foster 2018; Sperber 1996). Categories, stereotypes, and heuristics become cultural when they are replicated from person to person, e.g., through interaction or from mass media (for example, stereotypes for groups one has never met). This sharedness is why they can diffuse over macro-scale populations and can outlive the person or people with whom they originated.

Our focus on social sharedness departs from the recent treatment of schemas by Wood et al. (2019). While we agree with Wood and colleagues that most schemas are likely to some extent shared, we disagree that this renders the term "cultural schema" meaningless. Schemas can vary greatly in the extent of their sharedness (Sperber 1996:25) and the importance of *social* learning to their diffusion. For example, contrast the basic skill of walking with the greater set of cultural behaviors surrounding it. While parents often try to help their infants learn to walk, the great bulk of actual learning comes from infants' individual practice and experimentation: for example, infants can discover the correct motor strategy for walking down slopes solely by repeatedly trying, falling down, adjusting, and trying again (Gill, Adolph, and Vereijken 2009). The social aspect of learning to walk thus seemingly serves to encourage an acquisition of skills that come primarily from the physical affordances of our bodies and environment.[9]

[9] There are also edge cases where this judgement is less clear—e.g., when cultural influences partly alter the results of otherwise non-cultural learning. We thus follow Foster (2018:146) and Sperber (1996:82) in thinking of "cultural-

In contrast to these basic motor schemas, the greater complex of behaviors *surrounding* walking has clearly socially shared components. Consider an example of a child raised in a zero-gravity space station. In the absence of gravity, no amount of direct parental instruction may be sufficient to teach her to walk like children raised on earth. However, she could likely still learn to be instinctively judgmental of men who walk in too feminine a fashion (for example, from watching Hollywood films)—even if she has never herself walked, and even if she has never seen a person walk in real life. Some learned cognitive elements such as social norms and shared stereotypes are thus fundamentally more tied to the transfer of information between people than others. This connection to social learning distinguishes cultural schemas from non-cultural cognitive schemas (see Table 2).

This perspective lets us reexamine the status of "image schemas" (Lakoff 1990): cognitive representations that arise from intuitive human bodily experiences of the physical world, and are then used to understand abstract concepts by metaphorically mapping "the experiential structure from the 'imagistic' realms of sensory-motor experience to non-imagistic ('abstract') ones" (Hampe 2005:1–2). Wood and colleagues (2018) presented image schemas as a key example of cultural schemas. However, image schemas originate in immediate bodily experience rather than social learning, and may in fact be human universals grounded in common brain structures (Dodge and Lakoff 2008). Thus, while we agree with Wood and colleagues that image schemas are a promising conceptual tool for analyzing culture, we disagree that they are themselves *cultural* schemas.

"Representations". Stereotypes, categories, and normative heuristics carry information about the world: stereotypes describe the default characteristics of social groups; categories

---

ness" as a gradient: the more closely a schema is tied to the transfer of representations between people, the more cultural it is.

convey how people, objects, or events should be grouped; and normative heuristics prescribe appropriate judgements or actions. All three are thus fundamentally *representational* concepts (Pitt 2020; Ryder 2009). Following Sperber (2006:25) and Foster (2018), we hold that something is representational when it carries *interpretable information* about real or imagined states of the world. In other words, representations are meaningful—or, more strictly, *semantically or morally evaluable*: if X is a representation, it should be possible to ask questions like "what does X describe?" "is it true that X?" or "do we find it morally acceptable that X?" (Ryder 2009). This "minimalist" conceptualization lets us sidestep some thorny theoretical debates around the term "representation"[10].

In the example above, Billy's stereotype clearly carries information about the world: that boys don't and shouldn't work hard in school. This licenses questions such as: is it *true* that boys don't work hard in school? Does Billy (or ought we) *endorse* the idea that boys should behave this way? Does this stereotype *refer to* all forms of school-relevant effort, or are some learning situations exempt? What matters here is not that these questions are answered in any particular way, but that we can meaningfully ask them about Billy's stereotype.

This representational character distinguishes cultural schemas from lower-order cognitive constructs like motor programs, which make up a substantial portion of Type I learning (e.g., acclimatization, walking). Further, it distinguishes representations from *mere associations* in the technical connectionist usage, where "association" strictly refers to the strength of connection between two elements[11] (Pitt 2020). For example, due to the title of the

[10] Ryder notes that the "minimalist notion" that some mental contents are semantically evaluable is uncontroversial: "On this minimalist notion, only a radical eliminativist would deny that there are mental representations. It is not clear that there are any such radical eliminativists" (2009:234). Although our treatment requires only this minimalist notion, we note that recent research on model-based learning provides compelling reasons to believe that schemas may also be representational in "thicker" senses of the term—see (Railton 2017).

[11] As we detail in Part III, in connectionist models, simple associations are the building blocks of representations (which exist at a higher level of organization), but they are not representations themselves.

U.S. national anthem, many Americans may perceive the word "spangled" to have a patriotic connotation that is entirely missing from "sprinkled," "glittery," or "covered in spangles"—despite the near synonymy of these terms. Although the association between "spangled" and patriotism may be widely shared and is clearly socially learned, we cannot meaningfully ask whether the content of this association is true: it may be true *that* "spangled" and patriotism are associated, but the content of this association cannot be true or false. We can at most inquire as to why or how it is that the two concepts came to be associated, how strong the association is, or what further consequences the association has.

Given the sorts of cultural constructs generally investigated by sociologists, we suspect that applied users of the "cultural schemas" concept will rarely run afoul of this representational criterion. However, the criterion is theoretically and methodologically vital for delimiting the concept of cultural schemas: without it, the concept of "cultural schemas" would be coterminous with all of Type I culture (see Table 2), and which would leave it too overextended to add much value to empirical work that uses it.

Conceiving of schemas as a form of representation is also advantageous for the concept's sensitizing uses because it highlights the schemas' *mode-independent representational contents*, which are a promising area of investigation for future sociological work. These contents can be *externalized* into public expression and *internalized* into personal culture—although, in their public form, they would be public representations rather than schemas (Lizardo 2017; Sperber 1996). As Sperber points out, a person does "not discover the world unaided, and then make public her privately developed representations of it; rather, a great many of her representations of the world are acquired vicariously" (1996:78-79). The social origins of automatically deployable cognitive representations are a natural topic for cognitive sociology to pursue.

The representational contents of a schema may also be able to persist across the divide between automatic and deliberative cognition. For example, consider an experienced driver who successfully makes a left turn while engrossed in a demanding conversation. Without consciously thinking about her driving, this driver could still automatically turn on her blinker, move to the leftmost lane, stop at the red light, and complete her turn after it turns green. These rules are shared, automatically-deployable representations—but, when she first learned them from a DMV handbook, they initially required substantial conscious effort to follow correctly. Thus, while the rules eventually became routinized into cultural schemas, she had first internalized them as explicit deliberative knowledge. What persisted across the three modalities — public culture in the DMV handbook, personal declarative culture, and finally personal nondeclarative culture — are the schema's representational contents.

Other pieces of culture may follow similar paths. For example, as Kruglanski and Gigerenzer note, some "rules" of social life may similarly be first learned deliberatively, and then transformed into intuitively accessible structures by routinization. Vaisey and Frye (2019) make a similar point about explicitly learned physical skills. And, as Foster (2018) points out, narratives may be learned as explicit declarative culture, but then be used to produce implicit nondeclarative culture by the learner "replaying" the narratives in their mind.

Our account thus points to a question: what routes do the representational contents of cultural schemas take across this divide between different cultural modalities? Lizardo et al. (2016), Lizardo (2017), and Wood et al. (2018) suggest that the route from public to nondeclarative personal culture happens primarily through implicit slow learning enabled by long-term sustained exposure. We agree that this is likely the route taken by many cultural schemas; however, the above examples suggest that cultural schemas may also be first rapidly learned as explicit representations rather than through repeated personal experience, and then

made implicit through repeated application[12]. We conjecture that different cultural schemas follow different pathways—and, if the explicit-to-implicit pathway enables schemas to diffuse more effectively via mass media than the schemas that can only be learned implicitly, this may greatly advantage the spread of some cultural schemas over others[13]. We thus believe that the different routes that representational contents take across the "cultural triangle" (Lizardo 2017) of public culture, personal declarative culture, and personal nondeclarative culture would be a fascinating topic for future research.

"Heuristic". We specify that schemas are *"heuristic"* representations to clarify that they generally depict processes or phenomena in rough outline rather than in elaborate specific detail. This term is meant to stress that the "representations" contained in schemas are not verbal, "quasi-linguistic" (Lizardo 2017:97), explicit, or exact. This approximate character of schematic representations is already implied by their use within automatic cognition, which itself is often characterized as non-verbal, "heuristic", "fuzzy", or "gist"-like (e.g., Evans 2008; Gawronski and Bodenhausen 2006; Lizardo and Strand 2010; Reyna and Brainerd 2011; see also Martin 2010). It is also related to the important observation that people may hold cultural schemas without being able to verbalize their contents (Lizardo 2017; Vaisey 2009). We do not intend the term "heuristic" to place any restrictions on our concept's denotation not already implied by "automatic cognition." We instead include this term to assist sensitizing uses of the concept by making sure that our definition can serve as an accurate intuitive gloss of the underlying theory[14].

---

[12] Evans (2008:259) points out a similar split between different dual process models, where some conceive Type I processes as "automated" from Type II processes, and others treat them as different forms of learning.

[13] If more complex forms of cultural knowledge are more reliant on Type I enculturation (Lizardo et al. 2016), then this pathway may advantage the spread of simpler schemas that are able to cross from declarative to nondeclarative cognition. See also Sperber (1996) on "cultural attractors."

[14] This is also why we avoid referring to schemas as "associative"—see Kruglanski and Gigerenzer (2011).

In the next section, we discuss implications of our framework for what kinds of constructs are likely or unlikely to be cultural schemas. We use these implications to propose a series of conceptual and empirical tests. To illustrate how our approach can be productively applied to empirical cases, we then use it to investigate Sewell's (1992) conception of commodity schemas.

Implications of functional-level account

We noted above that our goal was to develop a conception of cultural schemas that fits tightly around the three central sociological examples: implicit categories (e.g., those used to automatically categorize a person by race or gender); default assumptions (e.g., stereotypes); and internalized cultural heuristics (e.g., those used to effortlessly recognize behavior as appropriate or inappropriate; or to evaluate agreement or disagreement with political claims.) As intended, all three constructs clearly fit under the above conceptualization: they are automatically accessible, widely shared, and heuristically representational in character.

Given the perennial threat of overextension faced by research programs into cognitive schemas, the opposite question is equally important: what *isn't* a cultural schema? We cannot, of course, provide an exhaustive list of things that are not schemas. However, our definition lends itself to a number of straightforward empirical and conceptual tests to answer whether any given cultural construct could be a schema. We discuss these below. We organize our account around the keywords in our conceptualization (representation, sharedness, automaticity), and additional consequences of their automatic character (conscious control, resource limits, nonverbal character and domain specificity.)

Representation. First, something could fail to be a schema if it is not representational. As we noted above, to qualify as a representation, a schema must carry information about the world. To test this, we may ask ourselves whether we can meaningfully posit questions like "is

X true," "do we endorse X," "when does X apply," "what does X refer to," "do people believe X" (Ryder 2009)? Lower-order cognitive constructs and mere associations will fail this test; stereotypes, heuristics, categorizations, and other higher-order constructs will pass.

Sharedness. To qualify as a cultural schema, an automatic representation must be widely (but not universally) shared. We could thus imagine an automatic mental representation that is schematic, but is too idiosyncratic to qualify as cultural (for example, if a girl's experiences with her father and brother lead her to develop the idiosyncratic stereotype "men love tomatoes," it would not be a *cultural* schema). Accordingly, the critical test question is this: is this schema common to some social group—e.g., occupants of a certain social position or geographical area, members of an organization, or participants in some social activity?

Scholars may be able to answer this question by looking for taken-for-granted character or intelligibility. For example, Rumelhart's "oil crisis" vignette described above is only easily comprehensible to readers who possess cognitive schemas for car showrooms and salespeople. If that passage had instead come from an interview transcript, it could be evidence that the respondent assumes that the listener shares the schemas in question. For other types of schemas—especially those with normative contents—it may be possible to observe whether the violation of the schema is understandable to others. For example, if someone violates an unstated behavioral prescription, would this violation be clear to others in the relevant community, without the need for explanation?

Automaticity. Finally, something could fail to be a cultural schema because it is not automatic, i.e., it cannot operate within Type I cognition. Although the Type I character of cultural schemas has been noted in many sociological treatments, we go beyond this existing work by proposing that empirical scholars use the distinctive characteristics of Type I and II cognition to judge whether something could be a schema. Type I cognition is characterized as

unconscious, implicit, automatic, fast, high capacity, domain- or context-specific, and nondeclarative, while Type II is conscious, explicit, controlled, high effort, low capacity, domain- or context-general, and declarative (Evans 2008; Lizardo 2017; Lizardo et al. 2016; Smith and DeCoster 2000). The following tests use these distinctions to identify whether particular cultural-cognitive construct could plausibly be a schema.

Conscious control. Since cultural schemas must be deployable automatically, i.e., without conscious control, any cognitive mechanism that we can *always* easily "turn off"—or choose not to engage—is not a schema. For example, many social norms and habits—e.g., our tendency to turn our bodies fully towards someone with more authority than ourselves, or a habit of touching the shoulder of a subordinate (Railton 2006)—are extremely difficult to turn off. In contrast, it is often easy to forget to follow the social norms of an unfamiliar culture, which we have not yet internalized deeply enough for automaticity. Put simply, intentionally controlled pieces of culture take effort to deploy; automatic ones take effort *not* to deploy. Scholars should ask which category the cultural construct they are examining falls into. If it is always deployed intentionally, it is likely an element of declarative culture rather than a schema (see Table 2).

Resource limits. Relatedly, Type I processes can be deployed in parallel: many can run at the same time, and they can operate even while we do other demanding tasks. This is not the case for Type II processes, which are dependent on access to a single conscious, limited-capacity central working memory system (Evans 2008). Because this working memory system can only handle one substantive conscious task at a time, Type II cognitive processes are similarly limited. Accordingly, something is unlikely to be a schema if it cannot be successfully deployed while performing other demanding cognitive tasks (e.g., being engrossed in a debate).

These differences are used as the basis for various *cognitive load* exercises (Feldon 2007). For example, to examine whether some focal task is automatic, a "digit memory"

exercise requires one group of subjects to remember a long string of random numbers, which necessitates active rehearsal (e.g., Wegner, Erber, and Zanakos 1993). If their performance on the focal task suffers compared to "un-loaded" participants, then the task requires working memory resources. A simpler but less robust test simply requires participants to complete the focal task under severe time pressure (e.g., Payne, Lambert, and Jacoby 2002), which may deny Type II processes the time and effort they require.

While behavioral experiments provide the best test of automaticity, they are not always realistically feasible. In such cases, a researcher could conceptually test a possible schema for automaticity by asking themselves whether it could successfully be deployed simultaneously with some other cognitively demanding task. Everyday tasks that mimic heavy cognitive load exercises include trying to remember a grocery list, engaging in a heated but intellectually challenging argument, or speaking in a foreign language one has only recently begun to learn. If it seems implausible that someone could invoke a representation concurrently with uninterrupted participation in such concurrent tasks, that representation is unlikely to be a cultural schema.

With many schemas, automaticity can be tested by constructing vignettes that require this schema to be effortlessly understood. For example, Rumelhart's (1980) car dealership example quoted above describes a car dealership without ever referencing cars and dealers directly (mentioning only related terms like "oil crisis" and "showroom floor"). If understanding this passage is so undemanding that population members can do so even under severe time pressure or while otherwise distracted, then their mental representation of a car dealership may qualify as a schema. We provide a more detailed example of this conceptual exercise below, where we apply this approach to Sewell's (1992) "commodification" schema.

The tests of automaticity we detailed in this section are applicable to all scholars whose approaches to observing schemas empirically do not themselves robustly guarantee automaticity. For example, recent techniques for surfacing schemas from survey data (e.g., Boutyline 2017; Goldberg 2011) locate organizing schemas by searching for certain kinds of similarity between different respondents' answers—e.g., responses that are rescaled, vertically shifted, and/or inverted versions of one another may be classified as following one schema. Since the data for these analyses are ordinary survey responses (e.g., where individuals report their musical tastes—a largely conscious, Type II activity), these methods cannot possibly guarantee that any latent pattern of attitudes or tastes they identify corresponds to an automatically deployable representation. Hunzaker and Valentino's innovative approach for measuring associations similarly taps respondents' conscious Type II processing (see Hunzaker and Valentino 2019:25). Users of these methods should thus apply the conceptual or empirical tests we described here to verify that the patterns they find indeed qualify as cultural schemas—or, failing that, they could instead use a more neutral term such as "cultural models", which does not carry the connotations of automaticity.

Nonverbal character and domain specificity. Type I cognition is primarily nonverbal, while Type II cognition is characteristically linguistic or discursive (Evans 2008:259). Accordingly, certain phenomena might fail to be schemas because they involve complex verbiage. For example, discourses, complete narratives, or highly detailed ideologies are unlikely to be schemas—although understanding and producing them likely involves applying many smaller, more specific schemas, e.g., narrative structure schemas or core ideological heuristics. This is related to our assertion that schemas are heuristic representations: while slow discursive cognition can support precise and intricate descriptions, the fast nonverbal

character of schemas suggests that they are instead relatively simple representations of gist (see also Martin 2010; Lizardo and Strand 2010:205–6).

A related quality of Type I cognition is *domain specificity*. Type II knowledge is context-independent and abstract—qualities enabled in part via the powerful expressive characteristics of language— and tied to general-purpose intelligence and abstract reasoning (e.g., logical or statistical inferences). In contrast, Type I knowledge is often *domain specific*—thoroughly tied to, and specifically functioning within, contexts closely resembling the one in which it was learned (Evans 2006, 2008). While Type II knowledge—e.g., abstract logic or rhetorical tools—can be transposed with relative ease across diverse contexts, the principles that underly Type I inferences may not be transferrable to other domains without the help of Type II processes. For related reasons, Type I skills and habits may be helpful in the familiar social contexts in which they were developed, but may prove useless or counterproductive in other environments—a characteristic that Foster (2018) links to Bourdieu's notions of *ontological complicity* and *hysteresis*.

Some classical research paradigms illustrate this difference. For example, in the Wason Selection Task, participants decide which of four cards—displaying D, 3, B, and 7—they must flip over to assess whether the following statement is true: "if there is a D on any side of a card, then there is a 3 on the other side." The logically correct answer is to flip the D to assess *modus ponens* (P implies Q; P; therefore Q), and the 7 to assess *modus tollens* (P implies Q; not-Q; therefore, not-P). Participants generally correctly pick out the "D", but most fail to select the "7". Yet when a structurally identical task was placed in a specific, familiar domain, participants fared much better: given the rule "if a person is drinking beer, then he must be over 20 years old" and presented with cards reading "drinking beer," "drinking coke," "25 years old," and "16 years old," participants generally successfully identified "drinking beer" and "16 years old" as

the cards to flip (Cosmides 1989). While the subjects' intuitions concerning the commonly invoked social rule about drinking age enabled them to correctly answer the contextualized "underage drinker" question, they could not transpose these intuitions to the decontextualized domain of formal logic, especially when Type II reasoning was restricted via time pressure (Roberts and Newton 2001).[15] Relatedly, scholars have repeatedly documented that people may implicitly learn to follow a statistical principle within a familiar domain, but be unable to successfully translate these intuitions to unfamiliar contexts (Evans 2008; Fong, Krantz, and Nisbett 1986; Nisbett and Krantz 1983).

While the domain-specific character of automatic cognition does not point to any specific tests, it suggests that a full empirical account of a cultural schema should aim to specify the domains in which that schema is invoked (we return to this in the concluding discussion). Additionally, it highlights a conflict between "cultural schemas" in the cognitive sense of the term—e.g., stereotypes, automatic categories, and implicit norms—and the term's older non-cognitive incarnation in cultural sociology articulated by Sewell (1992), who used the term to refer to a substantially broader set of elements of both personal and public culture, and conceived of them as applicable across highly varied contexts. We illustrate this in the following section.

Example: Sewell's "deep schemas"

To provide an example of how sociologists can conceptually examine whether a proposed construct could indeed be a cultural schema, we apply our approach to Sewell's (1992) commodity (or commodification) schema[16], which he used as a primary example of his notion

[15] Also see research on belief-biased reasoning (Evans and Curtis-Holmes 2005).
[16] Sewell terms this transposable schema "commodification." However, since "commodification" is the act of turning non-commodities into commodities, we believe it is more accurate to term the *transposable* schema "commodity." Commodification is then the *transposition* of the commodity schema to new domains.

of highly transposable "deep schemas." Our analysis will also illustrate how using the cognitive and cultural conceptions of cultural schemas interchangeably (as sociologists often do) can lead to faulty inferences.

To examine whether the commodity schema could indeed be a cultural schema in our cognitive sense of the term, we start by noting that this schema can be phrased roughly as *"if X is useful, then X can be bought and sold for money."* This construct is obviously representational and socially shared. We thus focus on verifying whether it could plausibly be deployed in implicit, automatic cognition. We construct a passage which could only be effortlessly interpreted with such a schema: *"I really wanted one too, but when I got to the store, I realized my credit card was already maxed out."* We can conceptually verify that this is a plausible schema by noting that we effortlessly interpreted this passage, and recognized—without even realizing that we are filling in blanks—that "I really wanted one too" probably refers to an object which can be bought at the store in question. This effortlessness and lack of conscious awareness suggest that we indeed possess a commodity schema.

Sewell, however, argues this commodity schema "is exceptionally transposable. It knows no natural limits; it can be applied not only to cloth, tobacco, or cooking pans, but to land, housework, bread, sex, advertising, emotions, or knowledge, each of which can be converted into any other by means of money" (1992:25). The preceding passage, on the other hand, references a thoroughly familiar setting closely associated with commodities: the sale of goods in a store. To examine whether we can indeed effortlessly transpose it to less familiar settings, we next construct a passage set in a context that saw many recent innovations in commodification: *"I really wanted to get to my airplane seat quickly, but when I got to my boarding gate, I realized my credit card was already maxed out."* To seamlessly interpret this passage, we would need to recognize that, since boarding early is a useful service, it is possible

to purchase it with money (specifically, by buying an upgrade to an earlier boarding group.) Unlike the first passage, we suspect that most readers could not *effortlessly* interpret this second passage. Moreover, we suspect that readers had this difficulty even if they knew that such boarding group upgrades exist. This suggests that the automatic transposition of this schema to *truly* novel domains is even less likely (e.g., if the example concerned the novel practice of paying your neighbor to say "hi" rather than look the other way when they pass you by on the street.)

We similarly expect that airline management did not effortlessly and unintentionally invent the idea of charging money for earlier boarding (or for the privilege of choosing one's own seat rather than having one assigned, etc.) Rather, it is more likely that these novel instances of commodification were *consciously* and *intentionally* devised by airline employees whose job required them to seek out new revenue streams for the airlines. Thus, while we can readily apply the commodity schema to known settings, it appears that we may not be able to transpose it to new domains without Type II cognition—i.e., unless we deploy it not as a schema, but as a consciously accessible, declarative cultural construct (Lizardo 2017).

This conclusion fits with the literature on domain specificity we reviewed above. It also does not present a problem for Sewell's (1992) own account, since Sewell does not require cultural schemas to be implicit. However, it demonstrates that the cognitive and non-cognitive conceptions are sometimes incompatible, which can be a problem for scholars who use them interchangeably. One prominent recent example of this common practice is found in Ray's (2019) theory of racial schemas. Ray follows Sewell in arguing that "racial schemas, like the commodity form, are easily applied to new organizational resources" (2019:32). At the same time, Ray argues that cultural schemas can be used outside of actors' intention or awareness—for example, "Placing broadly shared racial schemas at the center of a structural theory of race

renders conscious coordination unnecessary. […] In novel situations, people transpose existing racialized schemas to a new set of organizational resources. This transposition need not be conscious or intentional" (Ray 2019:35). Our reasoning suggests that this inference may be incorrect. While we agree with Ray that schematic cognition may explain situations where people unconsciously apply existing schemas to new situations *that have the same key features as familiar situations,* we suspect that racial schemas, like the commodity form, likely come to be applied to *truly* novel situations primarily via conscious, intentional transposition.

In the preceding sections, we developed our functional-level conception of schemas, and demonstrated that it has a clearly delimited denotation which covers the main types of schemas of interest to sociologists but excludes pieces of culture that do not share important characteristics with them. We also linked our definition to concrete tests—some empirical, others conceptual—that scholars can apply to their cases to decide whether the construct they are investigating is likely to be a cultural schema. We then applied this approach to Sewell's commodity schema to demonstrate it in action. We will now proceed to the algorithmic-level conceptualization.

PART III: ALGORITHMIC LEVEL

Schemas as Networks of Implicit Associations

Over the last decade, the "default assumptions" conception of schemas in sociology has been gradually giving way to a newer, lower-level definition, which we term the *"networks of implicit associations*". This view was first prominently articulated in sociology by Vaisey (2009), whose influence on how sociologists conceptualize schemas is an underappreciated aspect of this otherwise-famous work. Drawing on cognitive anthropologists Claudia Strauss and Naomi Quinn (1997), who in turn based their definition on the work of Rumelhart, McClelland, and

colleagues (1986), Vaisey (2009) talks of schemas as "cultural-cognitive structures [that] are built up out of experience and allow a person to respond to stimuli in ways that are automatically generated by the weighted connections between the elements of the inputs at hand," and "deep, largely unconscious networks of neural associations that facilitate perception, interpretation, and action" (Vaisey 2009:1686). Table 1 also includes related conceptualizations from Wood et al. (2018) and Hunzaker and Valentino (2019).

The implicit associations conception can be interpreted as either a representational/algorithmic-level or a biological-level claim. First, because the definition references neurons, some scholars may read it as a biological-level statement about brain wiring—i.e., the network of axons interconnecting literal neurons. But under this interpretation, this conception simply states that schemas are structures of interconnected neurons developed through experience, which can describe *any* product of individual learning. This is clearly too broad a conception to carry analytical weight.

We therefore focus on the second reading of this definition, which is implied by the language of "spreading activation" (and is likely the intended meaning). In this reading, it is a *representational/algorithmic* claim schemas can be modeled as "neural networks"—a class of computational models also called "connectionist." This is a broad class of dynamic network models loosely inspired by the biological structure of the brain (Thagard 2012), which represent cognitive processes as the diffusion of activation states across weighted networks of "neurons" (i.e., nodes). Connectionist models capture some important aspects of brain biology: the brain is indeed made up partly of processing units (neurons) and interconnections between them (synapses, dendrites, axons, etc.) These models abstract away other important aspects of brain function, such as the role of neurotransmitters and hormones and the complexity and variation in how different types of neurons function (Buckner and Garson 2019).

Variations in the wiring architecture and functional forms for weight and activation computations produces a vast range of connectionist models useful for examining a similarly broad set of cognitive processes, including grammar, speech, concept learning, visual information processing, and spatial movement (Buckner and Garson 2019; Smolensky 1988). Advocates of a type of connectionism called "predictive coding" have gone as far as suggesting that it "provides a unified account of all cognitive phenomena, including perception, reasoning, planning and motor control" (Buckner and Garson 2019). Connectionist networks can also form Turing-complete systems, meaning they can carry out any arbitrary computation doable by a computer (Graves, Wayne, and Danihelka 2014). Thus, the broad claim that schemas can be represented as neural networks does not, on its own, create any more limits to its denotation than the biological-level reading of this definition: without further specification, it again leaves schemas roughly equivalent to all learned cognitive contents.[17]

However, a number of sociologists have used a subtly different third version of this conception (e.g., Frye 2017; see Table 1), which we will demonstrate is substantially more useful. In what we call the "implicit concept network" (ICN), the network of implicit associations is not between *generic* computational units, but between stimuli, concepts, ideas, attitudes, or other *meaningful* cognitive objects (hereafter, we will call them "concepts" as shorthand.) This may at first appear very similar to Vaisey's (2009) "implicit associations" conception, but the subtle difference here is important: a connectionist model where the associations are between semantically meaningful units is a *localist* network. In contrast, the majority of neural network models use *distributed representations*, where different concepts

[17] Some sociologists using the implicit associations conception may prefer it as a sensitizing concept partly because, unlike "default assumptions", neural networks do not sound like the kind of "tool-like" mental structure that one can consciously deploy in declarative cognition. But this is a misunderstanding, as connectionist networks can be used to model both declarative and nondeclarative cognition — a point first demonstrated by Rumelhart et al. (1986), which is the same article that first prominently theorized cognitive schemas via connectionist networks.

correspond not to distinct nodes (or separate sets of nodes) but rather to different *patterns of activation* across *intersecting sets* of nodes (Buckner and Garson 2019; Thagard 2012; Widrow and Lehr 1990). The fact that the ICN lacks high-dimensional interactions[18], has no specialized input or output units, is fully connected, and has no pre-set direction for activation flow further narrows down the set of possible neural network models: specifically, it implies that ICNs are *localist fully connected recurrent neural networks with bidirectional ties*.

A typical network of this sort is depicted in Figure 2, which resembles the example from Rumelhart et al.'s pioneering treatment (1986). Here, each node represents a piece of furniture. During training, the network is "shown" a series of different rooms by simultaneously activating sets of nodes corresponding to the furniture seen in that room. So, the network may be shown a "typical bedroom" by activating *lamp*, *dresser*, *bed,* and *nightstand*[19]. The network then learns by following the Hebbian principle of "neurons that fire together wire together", i.e., if two nodes are activated simultaneously, their pairwise connection strength increases.

After training, the network is able to make guesses about rooms. In this mode, after a set of nodes is activated, this activation is allowed to cascade between the nodes in rough proportion to the tie strengths between them. In a key findings anchoring 1980s connectionist scholarship, Hopfield (1982) showed that, if there are stable patterns in the environment, such networks will learn to reproduce them. That is, if the network had been exposed to enough "typical bedrooms" during training and is then shown only *lamp*, *dresser*, and *nightstand*, the network would then automatically activate *bed*. This pattern-completion property is why

[18] For example, Hunzaker and Valentino (2019) measure the ties between all pairs of concepts in the schema by asking how strongly each concept is associated with each other concept. The presence of hidden units between the concepts would mean that ties between pairs of concepts are not be reducible to such pairwise associations.
[19] With cultural learning, this could represent exposure to real rooms as well as, e.g., seeing depictions of rooms in film or imagining rooms described in a novel.

Rumelhart et al. (1986) first proposed that connectionist networks provide a good model for schemas.

[Figure 2 about here]

This elaborated model provides a clearer description of the algorithmic-level building blocks of schemas than prior sociological conceptions and has a number of important implications for recent methodological research into schematic measurement. We develop these suggestions below.

Implications of clarified networks-of-associations model

Relationship to functional level. First, ICNs learn schemas in exactly the sense that they internalize default assumptions about the learning environment. Thus, although sociologists have at times treated the "default assumptions" and "implicit associations" conceptions of schemas as contradictory, they in fact appear to describe one and the same phenomenon: automatic, stereotype-like schemas.[20] The two descriptions sit at different levels of analysis, and thus best support different kinds of investigations—but because they refer to identical sets of schemas, these investigations can still yield mutually compatible findings.

Associations vs. schemas. Second, while connectionism-inspired sociological work often uses "implicit associations" and "schemas" interchangeably, the ICN model highlights a subtle but important difference between them: as Rumelhart et al. note, schemas are "the maxima in the goodness of fit landscapes" (1986:31) of these networks. They are thus a *possible product* of associative learning, which arises only when these networks have repeatedly encountered the same patterns in their learning environment. At that point, the connectionist networks stop merely associating concepts and start successfully "filling in the blank" in partially observed

[20] Rumelhart et al. thus argue that the "language of schemata and schema theories should be considered an approximation to the language of PDP [i.e., connectionism]" (1985:21).

phenomena (see also Schröder and Thagard 2013:258)—a property that is closely related to our earlier discussion of schemas' representational character. Thus, while stereotype-like schemas can indeed be modelled as sets of associations between concepts, *not every set of associations is a schema*.

This has direct implications for recent sociological work that aims to uncover cultural schemas by measuring subjects' associations between the concepts in the domain of interest (Hunzaker and Valentino 2019). For example, to capture schemas of poverty and welfare, Hunzaker and Valentino measure respondent's associations among a set of conceptionally related ideas including "being a welfare recipient," "becoming a teen parent," and "racism." The above logic suggests that this general approach may indeed measure schemas, but it is not *guaranteed* to be measuring a schema in any instance. Instead, the fact that the measured ICN can function as a schema should be demonstrated empirically—for example, by using the estimated network structure to make testable predictions about how subjects will perform automatic pattern completion within the given domain.

Types of connections to measure. This model can also be used to clarify what kinds of ties need to be measured by the connectionist-inspired methodologies in sociology. Since each tie in an ICN is produced through Hebbian learning, the tie weight T(A,B) is determined by how often A and B tend to coincide in the learning environment. This tie then has a single function: it determines how much the activation of one node should increase or decrease in response to changes in state of the other node. Roughly speaking, if you think of concept A, T(A,B) determines how much more (or less) likely this makes you to think of concept B. The ties have no other meaning.

The weight T(A,B) can thus be measured by sequential priming, which tests how much faster (or more likely) a person is to recognize concept B after they are presented with concept

A (e.g., Shepherd and Marshall 2018). It is also measured by the Implicit Association Test, which captures the relative ease of keeping A and B activated in mind at the same time. In contrast, when attempting to measure the connectionist network behind individuals' schemas of poverty, Hunzaker and Valentino measure each tie T(A,B) by asking respondents to report whether A and B are "related '(1) because one causes the others or (2) because they commonly go together in this context [of poverty in the US] for some other reason' (i.e., the relationship between the concepts need not be causal)" (2019:956). This procedure thus captures individuals' *perception* of causal and/or correlational relationship between A and B. While it is plausible that this quantity may be related to T(A,B), this relationship is not guaranteed. Future work building on Hunzaker and Valentino's approach should thus either validate this tie strength measure—e.g., by examining its relationship to IAT scores—or switch to a measure aimed specifically at implicit associations.

Relational vs. connectionist models. Connectionism-inspired measurement methods in sociology often seek to construct network representations of their target domains. While we find this approach highly promising, we note that existing work elides important differences between these networks and connectionist models. Relational views of culture long pre-date connectionism. For example, in mid-20th-century structural linguistics and anthropology, "word meaning was reduced to just those features that define the differences between words" (Quinn 2011:32) in a conceptual domain. The final product was a static relational map of a domain that could be function akin to a dictionary, where the meaning of a word or symbol could be "read off" from its position vis-à-vis other words. Connectionist models advance beyond these earlier approaches by focusing not on the *position* of nodes in a network but rather the *flow of activation* between nodes. They examine cognitively realistic ways that ideas could be learned and deployed—e.g., the way that the ICN described above learns to fill in the

missing elements in partially observed "bedrooms." They are dynamic models rather than static images.

While contemporary sociological approaches to schematic measurement are often inspired by connectionism, they generally construct static network images of associations between cultural elements in their domain that lack the key dynamic element of connectionist models (e.g., Boutyline 2017; Goldberg 2011; Hunzaker and Valentino 2019; but see Arseniev-Koehler and Foster 2020). Because of this, unlike connectionist models, the static networks produced by these methods neither measure nor articulate which cognitive tasks the measured schemas can perform. The cognitive relevance of the measured schemas is assumed rather than demonstrated. They are representations without an algorithm. In contrast, connectionist networks are explicit models of both contents and process: they depict the representation and literally implement the algorithm. This enables a connectionist model to serve as its own proofs-of-concept, thus making its theoretical import far clearer, richer, and more convincing. We therefore urge this important line of sociological work to go beyond viewing schemas as static networks, and to prioritize capturing their crucial dynamic components.

Other algorithmic-level models. As we argued above, the language of "default assumptions" provides a functional-level description of one variety of cultural schemas, but there are others it does not capture. Since ICNs describe stereotype-like schemas at a lower level of analysis, we would expect them to have the same limitations. And indeed, prior work notes that models of associative processing can account for only a subset of implicit cognition (e.g., Evans 2008:261). One major limitation of ICNs concerns directed relationships. ICNs assume that ties between concepts are symmetric, whereas cultural heuristics and social norms are instead based on asymmetric relationships of implication: for example, the etiquette rule "if you are at a dinner party, then it is impolite to talk loudly" is distinctly different from "if it is

impolite to talk loudly, then you are at a dinner party." The ICN cannot distinguish between these two rules. It is thus not the right model for cultural heuristics that have important directional components. In this kind of situation, simple "if-then" rules may be a more appropriate algorithmic approach to consider (Kruglanski and Gigerenzer 2011; Roeder and Ashby 2016). Another relevant algorithmic approach is Bayesian modeling, which has successfully been used to model inferences in a broad range of domains and is a rich area of investigation in contemporary cognitive science (Tenenbaum et al. 2011; Tenenbaum, Griffiths, and Kemp 2006).

More broadly, the range of different algorithmic approaches that have been productively used to study schemas is vast. For example, scholars have used more complex types of neural networks to model aspects of schematic cognition that ICNs cannot accurately represent, such as performance of subjects with memory degradation (Botvinick and Plaut 2004). Even with stereotypes, ICNs are not the only game in town: recent work in both cognitive science and sociology has successfully used high-dimensional spatial embeddings to model the stereotypes suggested by patterns of word cooccurrence in written texts (Arseniev-Koehler and Foster 2020; Caliskan, Bryson, and Narayanan 2017). As a representational scheme, spatial embeddings have the advantage of being able to simultaneously model an enormous set of interrelated stereotypes evidenced by multimillion- or even billion-word corpora—something that ICNs could not successfully capture.

Thus, we believe there is likely no single "one-size-fits-all" solution appropriate for modeling all types of cultural schemas at the algorithmic level—and consequently no single general-purpose methodology that could successfully measure all types of cultural schemas. Rather, methodologists developing new approaches to measuring schemas should begin by determining which functional-level features are the most theoretically central to the schemas

they are investigating. They should then select an algorithmic-level model that clearly and parsimoniously expresses these relevant features.

For example, to study implicit category schemas used in implicit categorization, a good starting point is the theory of cognitive category structure. As Rosch's (1978) first demonstrated, implicit categories are generally not defined via necessary and sufficient conditions, but rather have a radial structure with a small number of *prototypes* (or *exemplars*) in the center, and a range of increasingly peripheral cases that belong to the category to the extent that they resemble one of the central cases (see overview in Murphy and Hoffman 2012). This algorithmic-level view thus points to two main targets for empirical measurement: first, the prototypes that define the centers of these categories, and second, the dimensions along which similarity of other cases to these prototypes is judged.[21] Well-developed measurement techniques exist to uncover these elements of category structure, for example by applying multidimensional scaling to a dataset of free pile sort results (Vanpaemel and Storms 2008). Thus, while we believe that the current focus of schematic measurement methods on ICNs is highly promising, we also urge methodologists not to close themselves off to the wide library of other measurement and modeling methods in cognitive science, which are more appropriate than ICNs for some types of sociologically interesting schemas.

CONCLUDING DISCUSSION: WHAT ARE SCHEMAS *FOR*?

In this article, our contributions were aimed at maximizing the usefulness of the "cultural schemas" concept across its ever-expanding range of sociological uses. A key part of our contributions thus consisted of a series of recommendations for this growing body of applied and methodological work. In this concluding discussion, we propose one final

[21] Radial category learning can also be captured with connectionist models (Clark 1993:98–100); however, these models are substantially more complex than ICNs.

recommendation, which concerns two sensitizing questions that could productively guide empirical work.

To arrive at this recommendation, we begin with a theoretical question: why do people acquire automatically deployable heuristic representations? One broadly applicable answer is that most such representations may enable fluent performance at frequently recurrent cognitive tasks. As Wilson (2002) notes, tasks that involve novel interactions with unfamiliar environments require people to construct new mental representations: a cognitively demanding process that often requires conscious, Type II cognition. Automatic cognition is instead "responsible for forming stable […] representations of the *typical properties of the environment*" (Smith and DeCoster 2000:110 emphasis added). Representations accessible within automatic cognition are then part of a cognitive strategy that "involves *exploiting predictability in the task situation* being automatized", wherein "prior experience allows whatever representations are necessary for task performance to be built up before the fact" (Wilson 2002:633, emphasis added). This automatization speeds up individuals' responses and frees up cognitive resources for simultaneous engagement in other tasks. If such representations are widely socially shared within a population, they appear likely to automatize the recurring features of cognitive tasks that population members repeatedly encounter in their environments, including imagined environments constructed by other public representations (e.g., stories or films).

The central examples of schemas we discussed here fit well with this behavioral function. For example, implicit social categories such as age or gender allow people to effortlessly classify others into these groupings. This is a task that people rapidly perform at the onset of social interactions to determine how to relate to others (Ridgeway and Correll 2004). Implicit stereotypes are also frequently deployed to form expectations of others'

performance and thus decide their position in task groups (Correll and Ridgeway 2006). Memory schemas like the car dealership schema discussed above allow people to construct rich representations of common situations based on inputs that are incomplete, vaguely phrased, or partly remembered (DiMaggio 1997). Highly internalized cultural prescriptions for proper decorum allow people to fluently follow these social norms while engaged in cognitively taxing social interactions (Railton 2006).

We believe that an explicit focus on these types of concrete functional roles could enrich many empirical investigations into specific cultural schemas. As a source of research questions, it could lead scholars towards more comprehensive and satisfying understandings of the schemas they study. As a sensitizing idea, it may help scholars spot schemas in their empirical investigations and may also discourage fruitless searches for cognitive structures that would be too complex or too general to plausibly play a role in automatic cognition. Similarly, as part of algorithmic conceptualizations, it can clarify to methodologists which properties of schemas are important for their novel measures to capture. We therefore conclude our treatment with two sensitizing questions that could be posed about specific schemas:

(1) *what common task does this schema enable its bearers to perform fluently and effortlessly?*

—or, conversely—

(2) *what kind of schemas could arise from the informationally redundant features of this common task?*

REFERENCES

Arseniev-Koehler, Alina, and Jacob G. Foster. 2020. *Machine Learning as a Model for Cultural Learning: Teaching an Algorithm What It Means to Be Fat*. *preprint*. SocArXiv.

Baars, Bernard J., and Stan Franklin. 2003. "How Conscious Experience and Working Memory Interact." *Trends in Cognitive Sciences* 7(4):166–72.

Bachrach, Christine A. 2014. "Culture and Demography: From Reluctant Bedfellows to Committed Partners." *Demography* 51(1):3–25.

Bartlett, Frederic. 1932. *Remembering: A Study in Experimental and Social Psychology*.

Blair-Loy, Mary. 2001. "Cultural Constructions of Family Schemas: The Case of Women Finance Executives." *Gender & Society* 15(5):687–709.

Botvinick, Matthew, and David C. Plaut. 2004. "Doing without Schema Hierarchies: A Recurrent Connectionist Approach to Normal and Impaired Routine Sequential Action." *Psychological Review* 111(2):395.

Boutyline, Andrei. 2017. "Improving the Measurement of Shared Cultural Schemas with Correlational Class Analysis: Theory and Method." *Sociological Science* 4:353–93.

Brighton, Henry, and Gerd Gigerenzer. 2008. "Bayesian Brains and Cognitive Mechanisms: Harmony or Dissonance." *The Probabilistic Mind: Prospects for Bayesian Cognitive Science, Ed. N. Chater & M. Oaksford* 189–208.

Buckner, Cameron, and James Garson. 2019. "Connectionism." in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. Metaphysics Research Lab, Stanford University.

Caliskan, Aylin, Joanna J. Bryson, and Arvind Narayanan. 2017. "Semantics Derived Automatically from Language Corpora Contain Human-like Biases." *Science* 356(6334):183–86.

Castellani, John W., and Andrew J. Young. 2016. "Human Physiological Responses to Cold Exposure: Acute Responses and Acclimatization to Prolonged Exposure." *Autonomic Neuroscience* 196:63–74.

Cech, Erin A., and Mary Blair-Loy. 2014. "Consequences of Flexibility Stigma Among Academic Scientists and Engineers." *Work and Occupations* 41(1):86–110.

Cerulo, Karen A. 2010. "Mining the Intersections of Cognitive Sociology and Neuroscience." *Poetics* 38(2):115–32.

Clark, Andy. 1993. *Associative Engines: Connectionism, Concepts, and Representational Change*. MIT Press.

Clark, Andy, and Josefa Toribio. 1994. "Doing without Representing?" *Synthese* 101(3):401–31.

Clawson, Dan, and Naomi Gerstel. 2014. *Unequal Time: Gender, Class, and Family in Employment Schedules*. Russell Sage Foundation.

Correll, Shelley J., and Cecilia L. Ridgeway. 2006. "Expectation States Theory." Pp. 29–51 in *Handbook of Social Psychology*, *Handbooks of Sociology and Social Research*, edited by J. Delamater. Boston, MA: Springer US.

Cosmides, Leda. 1989. "The Logic of Social Exchange: Has Natural Selection Shaped How Humans Reason? Studies with the Wason Selection Task." *Cognition* 31(3):187–276.

D'Andrade, Roy G. 1995. *The Development of Cognitive Anthropology*. Cambridge University Press.

D'Angelo, Mariano, Giuseppe di Pellegrino, Stefano Seriani, Paolo Gallina, and Francesca Frassinetti. 2018. "The Sense of Agency Shapes Body Schema and Peripersonal Space." *Scientific Reports* 8(1):1–11.

DiMaggio, Paul. 1997. "Culture and Cognition." *Annual Review of Sociology* 23:263–87.

DiPrete, Thomas A., and Claudia Buchmann. 2013. *The Rise of Women: The Growing Gender Gap in Education and What It Means for American Schools*. Russell Sage Foundation.

Dodge, Ellen, and George Lakoff. 2008. "Image Schemas: From Linguistic Analysis to Neural Grounding." in *From Perception to Meaning: Image Schemas in Cognitive Linguistics*. Walter de Gruyter.

Edgell, Penny. 2012. "A Cultural Sociology of Religion: New Directions." *Annual Review of Sociology* 38(1):247–65.

Evans, Jonathan. 2008. "Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition." *Annual Review of Psychology* 59(1):255–78.

Evans, Jonathan, and Jodie Curtis-Holmes. 2005. "Rapid Responding Increases Belief Bias: Evidence for the Dual-Process Theory of Reasoning." *Thinking & Reasoning* 11(4):382–89.

Evans, Jonathan St BT. 2006. "Dual System Theories of Cognition: Some Issues." in *Proceedings of the Annual Meeting of the Cognitive Science Society*. Vol. 28.

Feldon, David. 2007. "Cognitive Load and Classroom Teaching: The Double-Edged Sword of Automaticity." *Educational Psychologist* 42(3):123–37.

Foster, Jacob G. 2018. "Culture and Computation: Steps to a Probably Approximately Correct Theory of Culture." *Poetics* 68:144–54.

Frijda, Nico H. 2010. "Impulsive Action and Motivation." *Biological Psychology* 84(3):570–79.

Frye, Margaret. 2012. "Bright Futures in Malawi's New Dawn: Educational Aspirations as Assertions of Identity." *American Journal of Sociology* 117(6):1565–1624.

Frye, Margaret. 2017. "Cultural Meanings and the Aggregation of Actions: The Case of Sex and Schooling in Malawi." *American Sociological Review* 82(5):945–76.

Gawronski, Bertram, and Galen V. Bodenhausen. 2006. "Associative and Propositional Processes in Evaluation: An Integrative Review of Implicit and Explicit Attitude Change." *Psychological Bulletin* 132(5):692–731.

Ghosh, Vanessa E., and Asaf Gilboa. 2014. "What Is a Memory Schema? A Historical Perspective on Current Neuroscience Literature." *Neuropsychologia* 53:104–14.

Gill, Simone V., Karen E. Adolph, and Beatrix Vereijken. 2009. "Change in Action: How Infants Learn to Walk Down Slopes." *Developmental Science* 12(6):888–902.

Goldberg, Amir. 2011. "Mapping Shared Understandings Using Relational Class Analysis: The Case of the Cultural Omnivore Reexamined." *American Journal of Sociology* 116(5):1397–1436.

Goldstone, Robert L., Alan Kersten, and Paulo F. Carvalho. 2018. "Categorization and Concepts." Pp. 1–43 in *Stevens' Handbook of Experimental Psychology and Cognitive Neuroscience*. American Cancer Society.

Gorman, Elizabeth H. 2005. "Gender Stereotypes, Same-Gender Preferences, and Organizational Variation in the Hiring of Women: Evidence from Law Firms." *American Sociological Review* 70(4):702–28.

Graves, Alex, Greg Wayne, and Ivo Danihelka. 2014. "Neural Turing Machines." *ArXiv:1410.5401 [Cs]*.

Haidt, Jonathan, Fredrik Björklund, and Scott Murphy. 2000. "Moral Dumbfounding: When Intuition Finds No Reason." *Lund Psychological Reports* Vol 1 no 2.

Hampe, Beate. 2005. *From Perception to Meaning: Image Schemas in Cognitive Linguistics*. Vol. 29. edited by Beate Hampe and J. E. Grady. Berlin, New York: Mouton de Gruyter.

Hodnik Čadež, Tatjana, and Vida Manfreda Kolar. 2015. "Comparison of Types of Generalizations and Problem-Solving Schemas Used to Solve a Mathematical Problem." *Educational Studies in Mathematics* 89(2):283–306.

Hopfield, John J. 1982. "Neural Networks and Physical Systems with Emergent Collective Computational Abilities." *Proceedings of the National Academy of Sciences* 79(8):2554–58.

Hsu, Greta, and Stine Grodal. 2015. "Category Taken-for-Grantedness as a Strategic Opportunity: The Case of Light Cigarettes, 1964 to 1993." *American Sociological Review* 80(1):28–62.

Hunzaker, M. B. Fallin. 2016. "Cultural Sentiments and Schema-Consistency Bias in Information Transmission." *American Sociological Review* 81(6):1223–50.

Hunzaker, M. B. Fallin, and Lauren Valentino. 2019. "Mapping Cultural Schemas: From Theory to Method." *American Sociological Review* 0003122419875638.

Kahan, Dan M. 2015. "Laws of Cognition and the Cognition of Law." *Cognition* 135:56–60.

van Kesteren, Marlieke T. R., Dirk J. Ruiter, Guillén Fernández, and Richard N. Henson. 2012. "How Schema and Novelty Augment Memory Formation." *Trends in Neurosciences* 35(4):211–19.

Kovács, Balázs, and Michael T. Hannan. 2010. "The Consequences of Category Spanning Depend on Contrast." *Research in the Sociology of Organizations* 31:175–201.

Kruglanski, Arie W., and Gerd Gigerenzer. 2011. "Intuitive and Deliberate Judgments Are Based on Common Principles." *Psychological Review* 118(1):97–109.

Lakoff, George. 1990. *Women, Fire, and Dangerous Things*. 1997th ed. University Of Chicago Press.

Lamont, Michèle, Laura Adler, Bo Yun Park, and Xin Xiang. 2017. "Bridging Cultural Sociology and Cognitive Psychology in Three Contemporary Research Programmes." *Nature Human Behaviour* 1(12):866.

Lewis, Amanda E. 2003. "Everyday Race-Making: Navigating Racial Boundaries in Schools." *American Behavioral Scientist* 47(3):283–305.

Lizardo, Omar. 2014. "Beyond the Comtean Schema: The Sociology of Culture and Cognition Versus Cognitive Social Science." *Sociological Forum* 29(4):983–89.

Lizardo, Omar. 2017. "Improving Cultural Analysis: Considering Personal Culture in Its Declarative and Nondeclarative Modes." *American Sociological Review* 82(1):88–115.

Lizardo, Omar, Robert Mowry, Brandon Sepulvado, Dustin S. Stoltz, Marshall A. Taylor, Justin Van Ness, and Michael Wood. 2016. "What Are Dual Process Models? Implications for Cultural Analysis in Sociology." *Sociological Theory* 34(4):287–310.

Lizardo, Omar, and Michael Strand. 2010. "Skills, Toolkits, Contexts and Institutions: Clarifying the Relationship between Different Approaches to Cognition in Cultural Sociology." *Poetics* 38(2):205–28.

Mandler, Jean M., and Cristóbal Pagán Cánovas. 2014. "On Defining Image Schemas." *Language and Cognition* 6(4):510–32.

Mandler, Jean M., and Gary H. Ritchey. 1977. "Long-Term Memory for Pictures." *Journal of Experimental Psychology: Human Learning and Memory* 3(4):386–96.

Mandler, Jean Matter. 1984. *Stories, Scripts, and Scenes: Aspects of Schema Theory.* Lawrence Erlbaum Associates, Inc., Publishers, 365 Broadway, Hillsdale, NJ 07642 ($19.95).

Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. Cambridge, Mass: The MIT Press.

Martin, John Levi. 2010. "Life's a Beach but You're an Ant, and Other Unwelcome News for the Sociology of Culture." *Poetics* 38(2):229–44.

Martin, John Levi, and Matthew Desmond. 2010. "Political Position and Social Knowledge." *Sociological Forum* 25(1):1–26.

Murphy, Gregory, and Aaron Hoffman. 2012. "Concepts." Pp. 275–91 in *Cambridge Handbook of Cognitive Science*.

Payne, B. Keith, Alan J. Lambert, and Larry L. Jacoby. 2002. "Best Laid Plans: Effects of Goals on Accessibility Bias and Cognitive Control in Race-Based Misperceptions of Weapons." *Journal of Experimental Social Psychology* 38(4):384–96.

Pitt, David. 2020. "Mental Representation." in *The Stanford Encyclopedia of Philosophy*, edited by E. N. Zalta. Metaphysics Research Lab, Stanford University.

Quinn, Naomi. 2011. "The History of the Cultural Models School Reconsidered: A Paradigm Shift in Cognitive Anthropology." Pp. 30–46 in *A Companion to Cognitive Anthropology*, edited by D. B. Kronenfeld, G. Bennardo, V. C. de Munck, and M. D. Fischer. Wiley-Blackwell.

Railton, Peter. 2006. "Normative Guidance." in *Oxford Studies in Metaethics*, edited by R. Shafer-Landau. Clarendon Press.

Railton, Peter. 2017. "Moral Learning: Conceptual Foundations and Normative Relevance." *Cognition* 167:172–90.

Ray, Victor. 2019. "A Theory of Racialized Organizations." *American Sociological Review* 84(1):26–53.

Reyna, Valerie F., and Charles J. Brainerd. 2011. "Dual Processes in Decision Making and Developmental Neuroscience: A Fuzzy-Trace Model." *Developmental Review* 31(2):180–206.

Ridgeway, Cecilia L., and Shelley J. Correll. 2004. "Unpacking the Gender System: A Theoretical Perspective on Gender Beliefs and Social Relations." *Gender & Society* 18(4):510–31.

Roberts, Maxwell J., and Elizabeth J. Newton. 2001. "Inspection Times, the Change Task, and the Rapid-Response Selection Task." *The Quarterly Journal of Experimental Psychology Section A* 54(4):1031–48.

Roeder, Jessica L., and F. Gregory Ashby. 2016. "What Is Automatized during Perceptual Categorization?" *Cognition* 154:22–33.

Rosch, Eleanor. 1978. "Principles of Categorization." in *Concepts: Core Readings*, edited by E. Margolis and S. Laurence. MIT Press.

Rumelhart, David E., James L. McClelland, and PDP Research Group. 1986. *Parallel Distributed Processing: Explorations in the Microstructure of Cognition: Foundations*. 1st edition. Cambridge, Mass: A Bradford Book.

Rumelhart, David E., and Andrew Ortony. 1977. "The Representation of Knowledge in Memory." in *Schooling and the Acquisition of Knowledge*.

Rumelhart, David E., Paul Smolensky, and James L. McClelland. 1986. "Schemata and Sequential Thought Processes in PDP Models." Pp. 3–57 in *Parallel Distributed Processing: Explorations in the Microstructures of Cognition*. Vol. 2, edited by J. L. McClelland and D. E. Rumelhart.

Ryder, Dan. 2009. "Problems of Representation I: Nature and Role." P. 233 in *The Routledge Companion to Philosophy of Psychology*, edited by J. S. P. Calvo. Routledge.

Schröder, Tobias, and Paul Thagard. 2013. "The Affective Meanings of Automatic Social Behaviors: Three Mechanisms That Explain Priming." *Psychological Review* 120(1):255–80.

Sewell, William H. 1992. "A Theory of Structure: Duality, Agency, and Transformation." *American Journal of Sociology* 98(1):1–29.

Shaw, Lynette. 2015. "Mechanics and Dynamics of Social Construction: Modeling the Emergence of Culture from Individual Mental Representation." *Poetics* 52:75–90.

Shepherd, Hana, and Emily A. Marshall. 2018. "The Implicit Activation Mechanism of Culture: A Survey Experiment on Associations with Childbearing." *Poetics* 69:1–14.

Shore, Bradd. 1998. *Culture in Mind: Cognition, Culture, and the Problem of Meaning*. Oxford University Press.

Smith, Eliot R., and Jamie DeCoster. 2000. "Dual-Process Models in Social and Cognitive Psychology: Conceptual Integration and Links to Underlying Memory Systems." *Personality and Social Psychology Review* 4(2):108–31.

Smolensky, Paul. 1988. "On the Proper Treatment of Connectionism." *Behavioral and Brain Sciences* 11(1):1–23.

Sperber, Dan. 1996. "Explaining Culture: A Naturalistic Approach." *Mind* 110(439):845–854.

Sperber, Dan. 2006. "Why a Deep Understanding of Cultural Evolution Is Incompatible with Shallow Psychology." Pp. 431–449 in *Roots of human sociality*.

Strauss, Claudia, and Naomi Quinn. 1997. *A Cognitive Theory of Cultural Meaning*. Cambridge University Press.

Swidler, Ann. 1986. "Culture in Action: Symbols and Strategies." *American Sociological Review* 51(2):273–86.

Tenenbaum, Joshua B., Thomas L. Griffiths, and Charles Kemp. 2006. "Theory-Based Bayesian Models of Inductive Learning and Reasoning." *Trends in Cognitive Sciences* 10(7):309–18.

Tenenbaum, Joshua B., Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. 2011. "How to Grow a Mind: Statistics, Structure, and Abstraction." *Science* 331(6022):1279–85.

Thagard, Paul. 2005. *Mind: Introduction to Cognitive Science, , 2nd Edition*. 2nd edition. Cambridge, Mass: A Bradford Book.

Thagard, Paul. 2012. "Cognitive Architectures." Pp. 275–91 in *Cambridge Handbook of Cognitive Science*.

Tsoukalas, Ioannis. 2006. "A Method for Studying Social Representations." *Quality and Quantity* 40(6):959–81.

Vaisey, Stephen. 2009. "Motivation and Justification: Toward a Dual-Process Theory of Culture in Action." *American Journal of Sociology* 114(6):1675–1715.

Vaisey, Stephen, and Margaret Frye. 2019. "The Old One-Two." in *The Oxford Handbook of Cognitive Sociology*. Oxford University Press.

Vaisey, Stephen, and Omar Lizardo. 2010. "Can Cultural Worldviews Influence Network Composition?" *Social Forces* 88(4):1595–1618.

Vanpaemel, Wolf, and Gert Storms. 2008. "In Search of Abstraction: The Varying Abstraction Model of Categorization." *Psychonomic Bulletin & Review* 15(4):732–49.

Wegner, Daniel M., Ralph Erber, and Sophia Zanakos. 1993. "Ironic Processes in the Mental Control of Mood and Mood-Related Thought." *Journal of Personality and Social Psychology* 65(6):1093–1104.

Widrow, B., and M. A. Lehr. 1990. "30 Years of Adaptive Neural Networks: Perceptron, Madaline, and Backpropagation." *Proceedings of the IEEE* 78(9):1415–42.

Wilde, Melissa J. 2004. "How Culture Mattered at Vatican II: Collegiality Trumps Authority in the Council's Social Movement Organizations." *American Sociological Review* 69(4):576–602.

Wilson, Margaret. 2002. "Six Views of Embodied Cognition." *Psychonomic Bulletin & Review* 9(4):625–36.

Wood, Michael Lee, Dustin S. Stoltz, Justin Van Ness, and Marshall A. Taylor. 2018. "Schemas and Frames." *Sociological Theory*.

Zerubavel, Eviatar. 1996. "Lumping and Splitting: Notes on Social Classification." *Sociological Forum* 11(3):421–33.

Figure 1. Diagram of our multilevel conceptualization of schemas

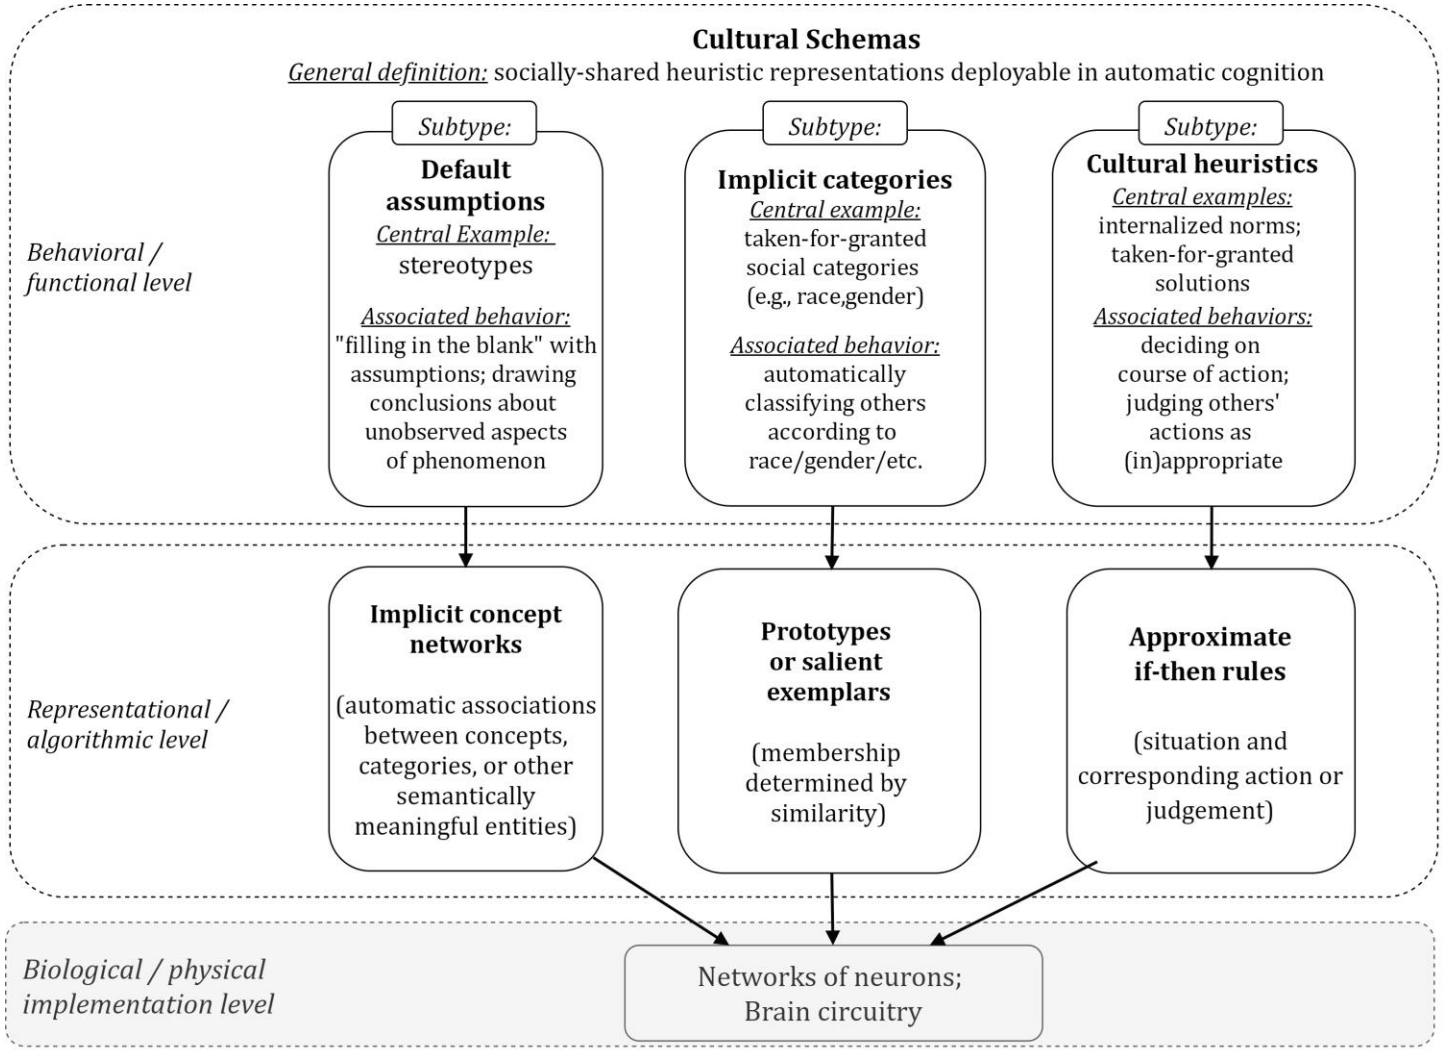

Figure 2. Example of an implicit concept network (cf. Rumelhart, Smolensky, et al. 1986).

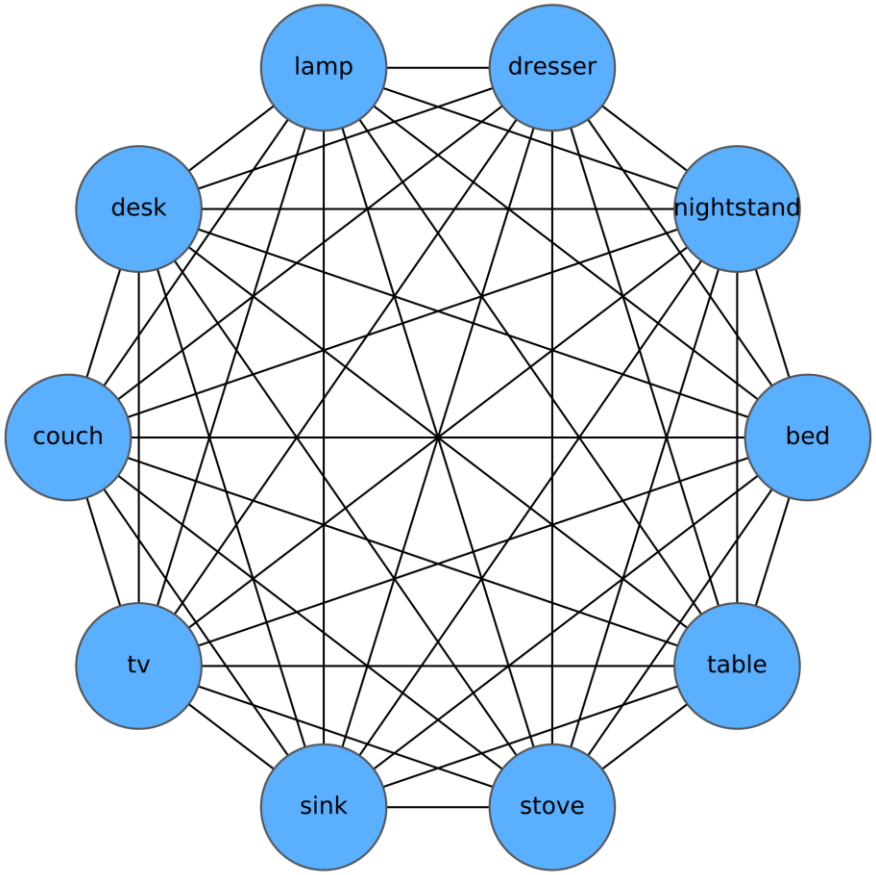

Table 1. Two major conceptions of schemas in sociology, with corresponding definitions.

Type	Source	Definition / explanation
Default assumptions		
	DiMaggio (1997)	First major sociological treatment to use this conception: "knowledge structures that represent objects or events and provide default assumptions about their characteristics, relationships, and entailments under conditions of incomplete information" (DiMaggio 1997).
	Cerulo (2010)	"knowledge structures such as stereotypes, scripts, etc. that, with broad strokes, represent the characteristics of people, places, objects or events and allow us to infer what these entities do, where they fit, and what to expect of them." (Cerulo 2010:117)
	Rumelhart and Ortony (1977)	Influential early usage from cognitive science: "data structures for representing the generic concepts stored in memory" (Rumelhart and Ortony 1977: 101).
Implicit associations		
	Vaisey (2009)	First major soc. use of this conception: "cultural-cognitive structures [...] built up out of experience [that] allow a person to respond to stimuli in ways that are automatically generated by the weighted connections between the elements of the inputs at hand. Proponents of this view do not talk of schemas as things that are "deployed" like tools but rather as *deep, largely unconscious networks of neural associations* that facilitate perception, interpretation, and action" (Vaisey 2009:1685–8, emphasis added).
	Wood et al. (2018)	"a form of personal culture—literally, entrenched multimodal (i.e., visual, aural, tactile, olfactory, kinesthetic, etc.) *neural associations* developed via repeated embodied experience (i.e., perceptual, sensorimotor, interactional) and stored in long term memory" (Wood et al. 2018, emphasis added).
	Frye (2017)	"Schemas influence cognition primarily through nonconscious channels; at their most basic level, they *associate related concepts* in our minds, such as coffee with warm or snake with danger" (Frye 2017:947, emphasis added)
	Rumelhart, Smolensky, and McClelland (1986)	Influential early example from cognitive science: "[S]chemata emerge at the moment they are needed from the interaction of *large numbers of much simpler elements all working in concert* with one another. [...]Input comes into the system, activating a set of units. These *units are interconnected with one another, forming a sort of constraint satisfaction network*. The inputs determine the starting state of the system and the exact shape of the goodness-of-fit landscape. The system then moves toward one of the goodness maxima." (Rumelhart, Smolensky, et al. 1986:20).

Table 2. What it means if a piece of cognitive contents either meets or does not meet each of our three criteria: automaticity, representational character, and sharedness. Cultural schemas meet all three criteria (Representational=Yes, Shared=Yes, and Automatic=Yes). Five other entries refer to a different kind of mental contents. The two bottom rows are left blank because declarative cognitive contents must be representational, making the corresponding combinations impossible.

Automatic	Represen-tational	Shared	Interpretation	Example
Yes	Yes	Yes	Cultural schema	Common stereotype (e.g., "Girls are studious")
		No	Non-cultural schema	Idiosyncratic stereotype (e.g., "Men love tomatoes")
	No	Yes	Non-schematic non-declarative culture	Mere association (e.g., association of "cat" with "bat", which is due to rhyme alone)
		No	Non-declarative non-culture	Most motor schemas (e.g., ability to rapidly stabilize yourself when you sense you are beginning to slip on ice)
No	Yes	Yes	Declarative culture	Many products of classroom learning (e.g., steps for integration by parts; a Chaucer poem memorized for class.)
		No	Declarative non-culture	Specific personal beliefs (e.g., knowledge of contents of my bag; my plans for next weekend)
	No	Yes	—	—
		No	—	—