This article was downloaded by: [46.64.248.158] On: 22 February 2024, At: 08:10
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA

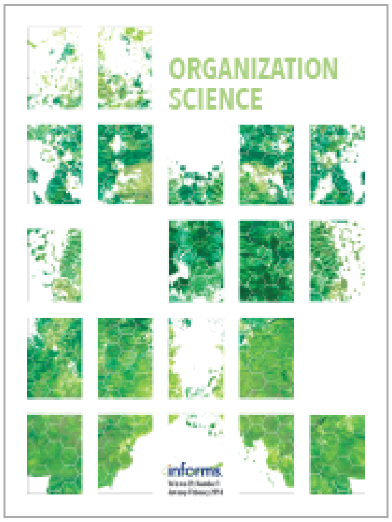

Organization Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Cognitive Economy and Product Categorization

Michele Piazzai, Min Liu, Martina Montauti

To cite this article:
Michele Piazzai, Min Liu, Martina Montauti (2024) Cognitive Economy and Product Categorization. Organization Science

Published online in Articles in Advance 22 Feb 2024

. https://doi.org/10.1287/orsc.2021.15751

Full terms and conditions of use: https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2024 The Author(s)

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

https://pubsonline.informs.org/journal/orsc

ORGANIZATION SCIENCE
Articles in Advance, pp. 1–24
ISSN 1047-7039 (print), ISSN 1526-5455 (online)

Cognitive Economy and Product Categorization

Michele Piazzai,[a,*] Min Liu,[b] Martina Montauti[c]

[a] Universidad Carlos III de Madrid, 28903 Getafe, Spain; [b] Durham University, Durham DH1 3LB, United Kingdom; [c] University of St. Gallen, CH-9000 St. Gallen, Switzerland

*Corresponding author

Contact: michele.piazzai@uc3m.es, https://orcid.org/0000-0001-9676-7422 (MP); min.liu@durham.ac.uk, https://orcid.org/0000-0001-8860-0374 (ML); martina.montauti@unisg.ch, https://orcid.org/0000-0003-4903-7798 (MM)

Received: October 1, 2021
Revised: January 4, 2023; September 30, 2023
Accepted: January 2, 2024
Published Online in Articles in Advance: February 22, 2024

https://doi.org/10.1287/orsc.2021.15751

Copyright: © 2024 The Author(s)

Abstract. In mediated markets, the categorization of products by mediators is critical to efficient interaction between producers and consumers. As organizational research tends to focus on the consequences of categorization rather than its antecedents, however, we know relatively little about why mediators assign one category label or another to a product. In this study, we argue that two informational properties of labels, *specificity* and *distinctiveness*, determine the outcomes of mediators' categorization decisions. Our analysis of product categorization decisions made by members of an online music community, 2000–2020, supports this argument. We find that a label's odds of being assigned to a product increase (a) if this label encodes information that is neither too similar nor too different from that which is encoded by a superordinate label, that is, it has moderate specificity; and (b) if it encodes information that differs as much as possible from that which is encoded by horizontally related labels, that is, it has maximal distinctiveness. These findings persist after controlling for other possible determinants of mediators' categorization decisions, including producers' claims to labels, products' typicality, and mediators' expertise.

Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "*Organization Science*. Copyright © 2024 The Author(s). https://doi.org/10.1287/orsc.2021.15751, used under a Creative Commons Attribution License: https://creativecommons.org/licenses/by/4.0/."

Funding: M. Piazzai received funding from the Community of Madrid [EPUC3M12, VPRICIT] and the Spanish Ministry of Science and Innovation [RYC2021-032325-I]. M. Liu received funding from Durham University Business School.

Supplemental Material: The e-companion is available at https://doi.org/10.1287/orsc.2021.15751.

Keywords: categorization • category labels • mediated markets • conditional logit • recorded music

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

1. Introduction

Organizational research characterizes markets as interfaces between producers and audiences, where producers vie for resources that members of the audience control, and audience members grant producers access to resources in exchange for products they like (Zuckerman 1999, 2017). According to this literature, an audience consists primarily of consumers, but it can also include actors like critics (Durand et al. 2007; Negro et al. 2010, 2011, 2015), analysts (Zuckerman 1999, Ruef and Patterson 2009), or enthusiasts (Koçak et al. 2014, Formilan and Boari 2021), who play the role of mediators by providing information about products in the form of category labels. In such "mediated markets" (Zuckerman 1999, p. 1400), consumers rely on category labels assigned to products by mediators to screen what is on offer. Despite their centrality to producer-audience interaction, the categorization of products by mediators continues to be framed in organizational research almost exclusively as an explanatory device. It is something that helps researchers make sense of variance downstream, in product evaluations or organizational performance, but apparently not something that deserves a more thorough understanding in itself. In empirical models, the category labels assigned to products by mediators are regularly used to compute independent variables but hardly ever figure as the dependent (see Younkin and Kashkooli 2020, for an exception). As a result, scholars concur that mediators' categorization decisions are critical to organizational outcomes but know relatively little about why a mediator assigns one category label or another.

Our objective in this study is to answer this question: given a product with some observable features and a set of labels that could be used to categorize it, what makes a mediator more likely to assign a particular label? This question is important to organization theory for two reasons. First, as explained by Zuckerman (1999), the

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

categorization of products by mediators is instrumental to the selection process by which producers gain access to resources necessary to their survival, such as money or attention. In this sense, being assigned one label or another can be "a matter of life and death" for organizations (Ozcan and Gurses 2018, p. 1789). In fact, many organizations will find that their chances to acquire these resources are curtailed already during mediators' categorization decisions, as their products end up on the wrong side of some relevant category boundary. Second, mediators' categorization decisions facilitate the social construction of economic reality by allowing products and organizations to get counted as members of the same market segments (Kennedy 2008). Through the consistent assignment of labels by mediators and their subsequent use by consumers, categories tend to achieve "ontological status" (Kennedy and Fiss 2013, p. 1144), that is, they grow embedded in the collective cognition of market actors, by virtue of which they become valuable tools for communication, coordination, and competition (Rosa et al. 1999, Cattani et al. 2017). Whether a category manages to retain this "ontological status" over time depends in no small measure on how frequently, or with what probability, the category label is assigned by mediators.

To make sense of this phenomenon, we account for potential determinants of label assignment that were suggested by previous literature. In particular, we consider recent studies in strategy, which suggest that producers can influence the categorization of their products through deliberate claims to advantageous labels (Pontikes 2018, Barlow et al. 2019), and in economic sociology, which suggest that typicality is a primary driver of categorization decisions (Goldberg et al. 2016, Smith and Chae 2017, Hannan et al. 2019). Our theoretical propositions, however, introduce another possible explanation. Building on a growing body of research that underscores the importance of hierarchical relations between category labels (Younkin and Kashkooli 2020, Cudennec and Durand 2023) and draws attention to informational properties stemming from these relations (Hannan et al. 2019, Lo et al. 2020), we propose that two such properties, *specificity* and *distinctiveness*, determine the outcomes of mediators' categorization decisions. We hypothesize that moderate levels of specificity and either moderate or maximal levels of distinctiveness allow mediators to economize on cognitive costs, making category labels better candidates for assignment. To test these hypotheses, we analyze over two decades' worth of product categorization decisions made by members of an online music community. Estimating choice models, we find that moderate levels of specificity and maximal levels of distinctiveness make a label more likely to be assigned to music products by community members. The effects of specificity and distinctiveness on these mediators' categorization decisions complement the effects of producers' claims, of products' typicality, and of mediators' own expertise. These effects persist in robustness models where we alter the criteria for inclusion in our sample, the composition of choice sets, and the operationalization of controls. They also seem robust to social influence and to the level of abstraction at which categorization occurs.

Our study makes theoretical, empirical, and practical contributions to organizational literature. On the theoretical front, it contributes to research on strategic categorization (Pontikes and Kim 2017, Barlow et al. 2019, Verhaal and Pontikes 2022) by showing how category labels' specificity and distinctiveness could constrain producers' strategic agency. It also contributes to research on category viability (Rhee et al. 2017, Lo et al. 2020, Soublière et al. 2023) by showing that moderate specificity and maximal distinctiveness make categories more likely to be assigned by mediators, and thus more likely to endure in the collective cognition of market actors. On the empirical front, our study highlights a potential bias in research that aims to track firms' positioning in a product market based on archival records of categorization decisions made by mediators (e.g., Montauti and Wezel 2016, Piazzai and Wijnberg 2019). Inasmuch as these decisions are shaped by labels' specificity and distinctiveness, they may suggest changes in a firm's positioning even if none occurred, undermining the internal validity of empirical measures meant to capture these changes. Finally, on the practical front, our study has implications for managers as it helps them understand why particular labels are preferentially assigned to their products by mediators. Managers should ultimately strive to influence mediators' categorization decisions, but to that end, knowing what drives these decisions in the absence of strategic intervention is necessary. By addressing this question, our study can help managers form realistic expectations and formulate their "category strategy" (Pontikes 2018).

2. Theory and Hypotheses

Organizational research views categorization as a probabilistic process in which a boundedly rational agent or decision maker, presented with an object and a set of labels, must assign one or more labels for the purpose of describing the object (Hannan et al. 2019). This process is ubiquitous to human mental life and constitutes a stepping-stone toward most everyday decisions. It turns out to be a process of great interest to organization theory because it is preparatory to, and inseparable from, selection mechanisms through which members of a market audience allocate resources to products and producers (Zuckerman 1999, 2017). Mediators, in particular, play a fundamental role in this selection by assigning category labels to products, which allows consumers to efficiently screen producers' offer. Mediators' categorization decisions are hence precursory to a two-

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

stage process whereby a consumer first uses mediator-assigned labels to determine which products are worthy of consideration, and only then looks at these products more closely to determine which of them is worthy of selection. Thus, a supermarket customer aiming to buy healthy groceries may not give equal consideration to all foodstuffs in a store but reserve most of it for those labeled "organic" by a certification agency. Similarly, someone looking to buy a car that comfortably sits a family of five may not look up technical details for all models on the market but consider only those categorized as "minivans" by specialized magazines or enthusiasts. It remains possible for a product fitting one's requirements to be summarily denied consideration because it fails to be assigned some crucial category label. As a result, mediators' categorization decisions come to affect important organizational outcomes, and the reasons why these key members of the audience tend to assign particular labels to products become relevant to organizational analysis.

Studies that showcase the consequence of categorization for organization theory can be broadly classified in two streams. The first one focuses on individual products or producers as the units of analysis, aiming to explain how their competitive success depends on category labels assigned to them by audiences (Cattani et al. 2017, Zhao et al. 2017). Labels are assigned by audience members to products or producers on the basis of observable features, such as product attributes (Younkin and Kashkooli 2020) and production techniques (Verhaal et al. 2015), but also on the basis of signals received from producers by way of advertisements (Hsu and Grodal 2015), narratives (Barlow et al. 2019, Taeuscher et al. 2022), and press releases (Pontikes 2012, Granqvist et al. 2013, Zunino et al. 2019). The effects of these categorizations are examined on a variety of outcomes, including sales (Zhao et al. 2013), growth (Smith and Chae 2016), access to venture capital (Wry and Lounsbury 2013), and valuation by investors (Zuckerman 1999, Cudennec and Durand 2023), critics (Durand et al. 2007; Negro et al. 2010, 2015), or consumers (Hsu 2006, Barlow et al. 2018). Although research in this first stream has been instrumental in bringing categorization to the forefront of organizational literature, it remains mostly preoccupied with consequences for phenomena that are traditionally relevant to studies in this field. This distinguishes it from research in the second stream, which focuses on categories themselves as the units of analysis and addresses outcomes like category creation or emergence (Durand and Khaire 2017), legitimation (Navis and Glynn 2010, Slavich et al. 2020), change (Lounsbury and Rao 2004, Negro et al. 2011, Jones et al. 2012), and demise (Kuilman and Van Driel 2013). Studies in this second stream suggest that categorization decisions made by audience members, and particularly mediators (Kennedy 2008, Khaire 2017), dictate the timing of a category's life cycle. For emergent categories, more frequent assignment of the category label helps legitimation; for established ones, it makes growth and survival more likely; and for declining ones, it can delay or prevent dissolution. Shifting focus from products and organizations to categories thereof, these studies noticeably expanded the scope of organizational literature and consolidated the study of categories as one of its most feverish domains (Vergne and Wry 2014).

Their respective merits notwithstanding, these two streams of research share a limitation in that they tend to consider categorization decisions as a black box. Lacking a model of how these decisions are made, they leave us with a partial picture of how categorization shapes economically relevant outcomes at the level of products, organizations, or categories (Hannan et al. 2019). Opening this black box would be not only interesting, but even necessary to ensure that theoretical arguments about categorization already circulating in organizational literature actually hold. Consider for instance Lo et al. (2020) proposal that a category should not be too distinctive, or dissimilar to related categories in a hierarchical classification system, in order to be viable for organizational activity. These authors argue that an overly distinctive category "is often less likely to be recognized as part of the system" (Lo et al. 2020, p. 91). This is ultimately because there is tension between a category's need to be similar enough to other categories to obtain recognition, and its need to be dissimilar enough from other categories to make useful distinctions. There is an intuitive analogy here with the tension described by optimal distinctiveness theory between the needs for conformity and differentiation at the level of products and producers (Deephouse 1999, Zuckerman 2016, Zhao et al. 2017). The validity of this analogy, however, depends on whether audience members, including mediators, are less likely to apply a highly distinctive category label as opposed to one that is just moderately distinctive, like they are less likely to select a highly distinctive product or producer relative to a moderately distinctive one. There is no evidence that this is the case.

We aim to fill this gap in organizational literature by proposing a model of how categorization decisions occur. We focus on mediators as key decision makers because their categorizations provide the "social screen" (Zuckerman 1999, p. 1404) on which consumers rely to select products for purchase. We build on the relational perspective on categories proposed by Lo et al. (2020, but see also Wry and Lounsbury 2013). Like these authors, we view categories as elements of a hierarchical system bound together by vertical and horizontal relations. Vertical relations arise between categories located at different levels of abstraction, such as "computers" and "workstations," or "jazz" and "free jazz." In these cases, one category is superordinate to the other in the

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

system's hierarchical ordering (Rosch et al. 1976). Horizontal relations arise between categories located at the same level of abstraction, such as "workstations" and "servers," or "free jazz" and "swing." In these cases, both categories are connected to the same superordinate. The labels of categories bound by vertical or horizontal relations can be more or less differentiated in terms of information they encode about product features: this variance is captured by two informational properties that, following previous studies (Hannan et al. 2019, Lo et al. 2020, Cudennec and Durand 2023), we term *specificity* and *distinctiveness*. In what follows, we develop hypotheses about how these properties affect a category label's probability of assignment. But first, we introduce two theoretical assumptions.

2.1. Assumptions

We assume mediators to behave as boundedly rational decision makers (Simon 1955). They are ostensibly interested in making the best possible decisions, but they must cope with a number of constraints that render this impractical: for example, they have limited memory to store information and limited computational power to process it. Because of these constraints, they are averse to the rigor that perfect rationality demands. Their "entire behavior […] is at all times motivated by the urge to minimize effort" (Zipf 1949, p. 3). They thus give up on finding the best possible solution to a decision problem and settle for solutions that are merely good enough. Such a "satisficing" attitude is well documented in the context of purchase decisions by consumers (Caplin et al. 2011), and it appears logical to expect it in categorization decisions by mediators as well. Indeed, psychological research suggests there are cognitive costs inherent to any categorization, which humans were conditioned by evolution to avoid (Anderson 1991). Remembering meanings for a set of labels inevitably come at a price, as information must be retrieved from memory and this requires expense of finite cognitive resources. Mediators strive to preserve these resources by favoring labels that make efficient use of them (cf. Pothos and Chater 2002), encoding information about products that is relevant to them and to consumers without imposing unnecessary costs.

This boundedly rational approach to categorization is at the core of recent studies in cognitive science (Lieder and Griffiths 2020). However, it is not a recent trend in cognitive psychology: in embryonic form, it was proposed by Rosch (1978, p. 28) in her principle of *cognitive economy*, which states that "as an organism, what one wishes to gain from one's categories is a great deal of information about the environment while conserving finite resources as much as possible." Experimental evidence shows that this principle affects virtually all categorizations, including those about products in markets (Johnson and Fornell 1987). Our first assumption is that it affects the categorization of products by mediators. To preserve finite resources, mediators do not strive to assign labels that encode the greatest amount of information about a product, but rather seek a trade-off between the amount of information encoded by a label and the cognitive costs or effort needed to remember it. This is not only because they, as decision makers, need to remember the meanings of labels in order to assign them—something they would like to accomplish with minimal effort (Zipf 1949)—but also because consumers, as beneficiaries of their decisions, need to remember the meanings of labels in order to interpret them. Thus, mediators assign labels that encode more information only if this is worth the cognitive costs, and among labels that encode comparable amounts of relevant information, they tend to assign those that encode the least amount of information overall.

Assumption 1. *Mediators are more likely to assign category labels that impose smaller cognitive costs for the same amount of relevant information.*

Assumption 1 will be pivotal to our arguments about the effects of specificity and distinctiveness. However, we do not see it as particularly restrictive because the urge to minimize cognitive costs is hardwired into human rationality and not even expert decision makers can be considered exempt. At the same time, we note that tolerance to cognitive effort can vary across individuals and correlates with prior knowledge of a domain. A category label that is cognitively costly for a novice could be cognitively economical for an expert (Chi et al. 1981). Therefore, we find it necessary to introduce another assumption to clarify just how much we expect mediators to know about the product domain.

Within the domain, we assume mediators to be sufficiently familiar with the meaning of each category label at their disposal to be able to evaluate it as a candidate for assignment to a product. If a label is not assigned, this must be because another one was preferable and not because its meaning was unknown. In addition, we assume mediators to know the vertical and horizontal relations binding category labels together into a hierarchical system. Following organizational literature, we define this system as "a cognitive representation of the structural relationships between categories that has achieved some consensus among […] audiences [and that] describes the various levels at which categorization takes place and defines how these levels are embedded or nested within one another" (Vergne and Wry 2014, p. 68). Given any two labels, our decision makers are able to determine whether these labels are related, and if so, whether their relationship is vertical or horizontal. If two labels are vertically related, then they know which one is the superordinate. They also know that the subordinate label inherits part of its informational content from the superordinate, if only in probability (Hannan

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

et al. 2019), because this is a consequence of nesting (see also Cudennec and Durand 2023). If two labels are horizontally related, then our decision makers know that these labels share part of their informational content by way of inheritance from the same superordinate. For example, given the three labels "computer," "workstation," and "server," they know how a "workstation" differs from a "server," but also that both "workstations" and "servers" are "computers," so some generic features of "computers" are expected of both. In effect, these subordinate labels expand a common ground set by the superordinate with additional information of their own.

Assumption 2. *Mediators know the meanings of category labels available for assignment and the hierarchical relationships among them.*

We find this assumption to be considerably more restrictive than the former. To begin with, there is the nontrivial provision built into it that hierarchical relationships among categories exist, and as per Vergne and Wry's (2014) definition, they are agreed-upon by audience members. This limits the scope of our study to settings where (a) there is an established hierarchical classification system, and (b) mediators know the meanings and relative positions of labels in this system. Condition (a) seems reasonable in settings where categorization is contingent on a highly institutionalized set of labels, like the Standard Industrial Classification (Zuckerman 1999), the United States Patent Classification (Ferguson and Carnabuci 2017), and the system of art genres (DiMaggio 1987). It is unlikely to hold if a classification system is "emergent or in flux" (Ruef and Patterson 2009, p. 486), which is often the case in nascent industries. Condition (b), instead, seems reasonable in settings where the average mediator has general knowledge of the market, possibly in addition to specialized knowledge about parts of it. This is easily true of professionals, like industry analysts, patent examiners, and art critics, but it is also true of highly engaged or "vanguard" consumers (Koçak et al. 2014). For example, music and beer enthusiasts who categorize products online have sufficient expertise, at least on average, to know what category labels mean and how they relate to one another (Barlow et al. 2018, Montauti 2019, Formilan and Boari 2021).

Joint consideration of Assumptions 1 and 2 gives rise to the following question: if mediators already know the meanings of category labels at their disposal, and need not learn these meanings as they approach categorization decisions, why do they incur cognitive costs when assigning labels to products? There are two reasons why cognitive costs arise even if label meanings have already been learned. One is that, after being learned, information needs to be activated or transferred from long-term to working memory in order to be used for decision making. Because working memory has limited capacity, however, its allocation is subject to opportunity costs (Kurzban et al. 2013). Any information committed to it precludes the commitment of other information, some of which could be relevant to the decision at hand. The severity of this problem grows in proportion to the amount of information a label encodes, because with more information encoded by a label, less space remains available for possible alternatives, potentially leading to a comparison among inferior options and a worse decision outcome. The foregone benefit of alternative uses of memory space is part of the cognitive costs a mediator pays when remembering the meanings of labels.

Another reason why cognitive costs arise is that transferring information to working memory comes with the risk of memory interference (Oberauer and Kliegl 2006, Jonides et al. 2008). This occurs when information is poorly recalled, partly omitted, or mixed up with other information: for example, a mediator may consider the assignment of two labels and mistakenly believe that one of them requires product features that are actually required by the other. This too can lead to a worse decision outcome. Like before, this is a problem that grows in proportion to the amount of information a label encodes, because "with increased memory demand, each individual item suffers from interference from more other items" (Oberauer and Kliegl 2006, p. 622). Together, constraints on memory space and the risk of memory interference give rise to a cognitive price. Mediators are required to pay this price even in the absence of learning because they incur it in the course of remembering label meanings. With this in mind, we can develop testable hypotheses about how category labels' specificity and distinctiveness affect mediators' decision process.

2.2. Specificity

The specificity of a category label is defined as the extent to which the information encoded by this label differs from the information encoded by its superordinate (Cudennec and Durand 2023). In organizational literature, this is also referred to as *informativeness* (Hannan et al. 2019). Because the information encoded by a superordinate tends to be inherited by its subordinate, any difference between the two labels is primarily due to extra information that the subordinate adds. Consider the case of a music listener browsing the catalog of a streaming service and finding a product categorized as "math rock." From this label, the consumer may readily infer that the product has features consistent with the superordinate "rock," but specificity requires "math rock" to say something more about the product's features. Provided one knows its meaning, this label happens to be quite specific, because it offers strong indications as to the product's musicological features: for example, it suggests odd time signatures, irregular rhythmic structure, extended chords, angular melodies,

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

counterpoint, and possibly other attributes about which "rock" remains relatively vague. Other labels sharing the same superordinate, like "classic rock" and "hard rock," are not nearly as specific. They add less to their common baseline. A greater level of specificity means that the subordinate label encodes and has potential to convey a greater amount of information relative to its superordinate.

We remark that by amount of information here we refer to a continuous quantity. In principle, it is possible for a label to encode a single detail about a product, for example, that it has a particular value for one feature while saying nothing about other features, and still be highly specific. This is because that single detail can be very different from what one would expect on the basis of the superordinate. In this regard, our approach is consistent with information theory, where the amount of information encoded in a message depends on how different that information is from a contextual expectation, not on the length of the message (see, for example, Ferrer-i-Cancho and Del Prado Martín 2011).

Having labels encode more information about product features is generally helpful to mediators because it allows them to more accurately determine if a label fits a product to be categorized. It is also helpful to consumers, as by inferring product features more accurately from a label previously assigned by mediators they are better able to reduce uncertainty. Therefore, we may expect a more specific label to be a more attractive candidate for assignment. Because of cognitive limitations to which humans are subject, however, greater specificity is not always desirable: to begin with, information yields diminishing returns as only so much of it can be relevant to either mediators or consumers. After a critical threshold is reached, having more information no longer makes a difference because there is already enough to "satisfice" (Caplin et al. 2011). Furthermore, information does not come for free: both evaluating labels for potential assignment and interpreting labels previously assigned by others to infer the features of products entails cognitive costs, which increase with the amount of information encoded by a label on top of its superordinate, and continue to do so even after the threshold is reached after which more information is no longer useful.

Empirical evidence of the connection between information and cognitive costs can be found in experimental research by Shepard et al. (1961), who examined human subjects' performance in categorization tasks and observed that category labels loaded with more information about object features cause subjects to spend more time thinking and increase the probability of mistakes. The authors observed that these negative consequences ensue even if subjects already know the meanings of labels, because they learned them weeks before, so dealing with more information seems to entail greater costs regardless of learning. Indeed, opportunity costs generated by constraints on working memory and the risk of memory interference occur when information previously learned needs to be activated or remembered. By increasing the amount of information a label encodes compared with is superordinate, specificity expands the total bundle of information to be activated or remembered: therefore, opportunity costs increase, interference is more probable, and the label imposes a greater cognitive price. This is true not only for mediators, who must remember the meaning of a label to evaluate it for assignment, but also for consumers who rely on mediators' assignments to decide which products are worth considering. To the extent that mediators take consumers' perspective, they are aware that consumers are subject to the same cognitive limitations as themselves (Epley et al. 2004). Thus, they are likely to realize that excessive specificity makes labels cumbersome not only for them, but also for the beneficiaries of their categorization decisions.

Given this combination of benefits and costs, the relationship between a category label's degree of specificity and its probability of being assigned to a product reflects an interplay of positive and negative consequences. On the positive side, greater specificity allows mediators to specialize the meaning of a superordinate label and better describe a product. On the negative side, it increases the cognitive costs involved in the label's assignment and subsequent interpretation. A mediator should find the drawbacks of greater specificity worth incurring if the additional information encoded on top of the superordinate is relevant, for example, because it is necessary for them to sort products effectively or it is necessary for consumers to decide if a product is worth considering. But relevant information is eventually exhausted, and at that point specificity makes the label more costly without concurrently making it more useful. We hence expect the benefits of specificity to initially outweigh its costs, but while costs continue to build up as specificity increases, benefits tend to subside. The combination of costs increasing at a constant rate and benefits increasing at a decreasing rate makes the functional relationship between specificity and the probability of label assignment nonmonotonic, and more concretely, inverse U-shaped (see Haans et al. 2016).

Hypothesis 1. *The specificity of a category label has an inverse U-shaped relationship with the label's probability of assignment to a product in a mediator's categorization decision.*

2.3. Distinctiveness

The distinctiveness of a category label is defined as the extent to which information encoded by this label differs from information encoded by other labels connected to

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

the same superordinate (Hannan et al. 2019). As labels connected to the same superordinate tend to inherit the same informational content, any difference between them is primarily due to extra information that each of them adds. Greater distinctiveness requires this to be more differentiated, and in the extreme unique. Unlike specificity, however, differentiation between labels connected to the same superordinate depends on the kind of information these labels encode, rather than the amount. Labels sharing the same superordinate may add the same amount of information to their common baseline, and thus be equivalent in terms of specificity, but because the information they add could be different for each, they can still be more or less distinctive. Consider "punk rock" and "prog rock," which share the same superordinate, as do "hard rock," "glam rock," and other varieties of "rock." Any of these labels is apt to categorize products with the generic features of rock, and all of them add something to this superordinate, but what they add is different for each. "Prog rock," for instance, is associated with compositional complexity, elaborate instrumentation, and meticulous sound editing, whereas "punk rock" is associated with compositional simplicity, stripped-down instrumentation, and minimal editing that aims to recreate the roughness and imperfection of live takes. Due to their diametric features, these labels differ in the kind of information they encode and not necessarily in the amount. "Punk rock" is relatively indistinctive insofar as products are expected to feature loud aggressive vocals and distorted guitars, but this is also true for other forms of "rock." By comparison, "prog rock" is highly distinctive because products are expected to have features uncommon within "rock," such as long solos that showcase technical virtuosity, and these set the label apart.

Because distinctiveness does not require a label to encode more information with respect to its superordinate, using more distinctive labels does not necessarily impose greater cognitive costs on mediators. As a result, greater distinctiveness does not necessarily hinder their pursuit of cognitive economy. Indeed, more distinctive features do not occupy more space in working memory, nor is remembering such features somehow more liable to interference. On the contrary, holding distinctive information in working memory means interference is less likely to occur (Oberauer and Kliegl 2006). In addition, distinctive information tends to be remembered more easily because it stimulates deeper cognitive processing when it is encountered (Hunt et al. 1992). Therefore, and provided that a label's meaning is known, distinctiveness can even make the label less cognitively demanding for mediators. Meanwhile, for consumers, products to which a more distinctive label was assigned may be easier to interpret unambiguously. Being associated with a combination of features that rarely co-occur elsewhere, a highly distinctive label allows consumers "to have as many properties as possible predictable from knowing any one property" of a product (Rosch 1978, pp. 28–29). For example, knowing that a piece of "rock" features long and complex guitar solos enables someone familiar with the "prog rock" label to infer the product is also likely to feature elaborate instrumentation and extensive sound engineering. Someone who likes this combination of features could conclude that the product is worth listening, whereas someone who dislikes it might find it better to avoid. Thus, distinctiveness facilitates the production and activation of behavioral patterns that make it easier to deal with an uncertain world. To be maximally effective at reducing uncertainty, category labels should be "clearly demarcated bins, into which any object […] will neatly and uniquely fit" (Bowker and Star 2000, p. 10). In other words, they should be maximally distinctive.

Although from a purely cognitive standpoint we could expect greater distinctiveness to be always desirable for a category label, recent research in organization theory offers a different perspective. Examining the possible consequences of category distinctiveness, Lo et al. (2020, p. 91) pointed out that horizontal differentiation between categories can have important drawbacks:

> [T]here exists a fundamental tension when it comes to the optimal level of distinctiveness for a focal category. On the one hand, a category that shares too little overlap with or is perceived as being too distant from other categories in a classification system has a lower likelihood of being viable. The relational approach to categories suggests that a category's meaning is derived from its relationship to other categories in the system; a category that occupies a very peripheral position is often less likely to be recognized as part of the system. On the other hand, if a category overlaps too much with adjacent categories in the system, its classificatory utility declines, rendering it less useful as a stand-alone category.

This perspective is familiar to organization scholars because it resonates with a prolific stream of literature based on the social-psychological theory of optimal distinctiveness (Brewer 1991). Research in this stream has long considered distinctiveness as a property of products or producers, analyzing how differentiation from products or producers within the same category affects product appeal and organizational performance (Zuckerman 2016, Zhao et al. 2017). Early studies in this stream focused on a trade-off between differentiation and conformity, claiming that producers—and by extension, their products—must be perceived as sufficiently similar to their competitors to be legitimate, yet sufficiently dissimilar to be valuable or difficult to substitute. They should be "as different as legitimately possible" (Deephouse 1999, p. 147) because of an inverse U-shaped relationship between distinctiveness and competitive outcomes like performance or appeal.

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

More recent research on optimal distinctiveness revisited this fundamental tension and examined situations where the aforementioned relationship deviates from an inverse U-shape. For example, Zhao et al. (2018) considered the consequences of product distinctiveness during the emergent phase of a category's life cycle, and found that the relationship with product performance is initially positive but consolidates into an inverse U-shape as the category becomes mature and institutionalized. Barlow et al. (2019) studied distinctiveness with respect to multiple reference points within a product category, including highly typical and highly successful competitors, and found that depending on the reference point the relationship can be positive or negative. Moving from products to producers, Taeuscher et al. (2022) analyzed the optimal value of distinctiveness for organizations in more or less distinctive categories and found that, when organizational distinctiveness is matched with greater category distinctiveness, its relationship with organizational performance switches from an inverse U-shape to a U-shape. Haans (2019) reported a similar reversal when comparing organizations in heterogeneous vs. homogeneous categories. Despite these recent studies, organizational literature on optimal distinctiveness did not abandon Deephouse's (1999) prediction of an inverse U-shape: on the contrary, this prediction continues to be represented in contemporary studies (Goldenstein et al. 2019, Taeuscher and Rothe 2021). It is commonly argued, for instance, that in highly institutionalized contexts such as cultural or financial markets an inverse U-shaped relationship arises (Zuckerman 2016, 2017; Zhao et al. 2018). The optimal distinctiveness argument that Lo et al. (2020, p. 91) "extend [...] to also apply to inter-category relations" presumes an inverse U-shape.

Although optimal distinctiveness research originally concerned a different unit of analysis, that is, products or producers, and primarily focuses on the consequences of categorization decisions made by producers themselves rather than by mediators, the category-level approach outlined by Lo et al. (2020) is directly relevant to our theoretical arguments. It suggests that distinctiveness should be moderate, not maximal, for a category label to make a good candidate for assignment. Therefore, psychological and organizational research point to different expectations about the ideal level of distinctiveness a category label should possess. Based on psychological literature, we would expect the ideal level of distinctiveness to be as high as possible, because more distinctive labels impose smaller cognitive costs and enable less ambiguous inferences. Based on organizational literature, instead, we would expect the ideal level of distinctiveness to be only moderate, because indistinctive labels are easier to substitute but overly distinctive labels are unlikely to be recognized as legitimate. Rather than siding with one perspective or the other, we formulate alternative hypotheses.

Hypothesis 2a. *The distinctiveness of a category label has a positive relationship with the label's probability of assignment to a product in a mediator's categorization decision.*

Hypothesis 2b. *The distinctiveness of a category label has an inverse U-shaped relationship with the label's probability of assignment to a product in a mediator's categorization decision.*

3. Data and Methods

Testing our hypotheses empirically requires us to observe a large number of categorization decisions made by mediators in a product market. We aim to model the choice an individual mediator makes when, presented with a product having certain features, she selects one or more category labels to describe this product. To specify this model, we must know all candidate labels the mediator can consider for assignment, that is, all the elements of her choice set. For this reason, it is preferable for us to analyze a context where mediators can only assign labels from a predetermined list, fixed and known to us as modelers, rather than freely coming up with labels in an open-ended process. Data fitting these requirements is readily available from websites that allow their users to categorize products in a given market using predetermined lists of labels. These websites are especially common in markets for cultural goods, where large communities of enthusiasts gather online to build searchable databases of products on offer. Examples familiar to organizational literature include BeerAdvocate (Verhaal et al. 2015, Barlow et al. 2018), Discogs (Montauti and Wezel 2016, Montauti 2019, Formilan and Boari 2021), Goodreads (Kovács and Sharkey 2015), and the Internet Movie Database (Hsu 2006, Hsu et al. 2009). Members of these communities are motivated by passion for the product domain, and invest considerable time and effort in collaboration with likeminded peers (Faraj et al. 2011). Their categorizations are intended to build a public repository of knowledge on which consumers can rely to find products that fit their tastes. Beside considerations of data availability, an advantage of studying an online community is that the relevant classification system tends to be institutionalized (DiMaggio 1987), and being part of a "vanguard audience" (Koçak et al. 2014), community members have sufficient domain knowledge to satisfy Assumption 2.

We collect data from Discogs, a comprehensive music database and marketplace where users can categorize products using genre and subgenre (*style*) labels. Discogs was born with a mission "to build the biggest and most comprehensive interactive public music database

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

in the world […] Because music is what makes us human, and keeping a well-organized, public archive of all the recorded music in the world helps preserve a full picture of who we are" (Discogs 2018). Since its launch in 2000, over 600,000 users have contributed toward this goal, making Discogs a major informational resource for music listeners and one of the largest music websites by search traffic (Alexa 2022). Genres in the database comprise relatively broad types of music, such as "rock," "electronic," and "jazz." Styles are nested within genres so as to form a two-level system, and comprise more specific types of music like "rockabilly," "techno," and "free jazz." Sociological research suggests that both genre and style labels on Discogs map to meaningful categories in audience members' perception (Van Venrooij 2015). Consistent with this, labels assigned to products by Discogs users are used by computer scientists as a ground truth corpus to train music recommendation algorithms (Bogdanov and Herrera 2012). They are also used in strategy and organization research to analyze competition among record companies (Montauti 2019, Piazzai and Wijnberg 2019, Zanella et al. 2021).

We consider Discogs users as mediators in the market for recorded music. This is justified by the visibility that Discogs enjoys in the market. Users are aware that the information they provide through category labels is relied upon by ordinary consumers, as Discogs pages rank highly in search results whenever music products are looked up online. Moreover, third-party applications connect with the Discogs API to relay information from the Discogs database to their own users. On Discogs itself, category labels assigned to products by Discogs users serve to organize a thriving marketplace for second-hand records, which at the time of writing includes over 68 million listings, with prices ranging from 0.01 to 1.2 million USD. The match of buyers and sellers in this marketplace depends on products being listed in the genre and style categories where people expect to find them. For this reason, members of the Discogs community have an incentive to make accurate categorization decisions, and are encouraged by community guidelines to provide correct and complete information. They are also encouraged to correct existing database entries if they think some of the information provided is inaccurate. We must note that, although Discogs users act as mediators, they are also consumers of music themselves, so this setting does not allow a neat distinction between mediator and consumer roles as in Zuckerman's (1999) study of the market for securities. However, this is not a problem for our analysis as we are only interested in the decisions Discogs users make in the capacity of mediators, that is, categorization decisions. We do not analyze their consumption choices, nor those other consumers make on the basis of their categorizations.

Our intention is to look at Discogs users' assignment of category labels to products in the Discogs database as the outcome to be explained. To that end, we could focus on predicting either users' assignment of genres or their assignment of styles. We replicated our analysis at both levels of the hierarchy and found similar results. The main analysis below concerns styles, because we expect this to be the most relevant level of abstraction for highly engaged consumers (see also Formilan and Boari 2021). Moreover, at this level of abstraction we observe more labels as well as greater variance in their characteristics, including specificity and distinctiveness. Estimates explaining users' assignment of genres as opposed to styles will be presented as robustness tests.

Any user on Discogs can categorize any product in the database at any time by picking genres from a checklist and then choosing one or more styles connected to these genres from a dropdown menu, as illustrated in Appendix A. The complete list of labels, curated by Discogs developers and updated over time at the community's request, currently includes 15 genres and 580 styles. After being assigned to a product, genre and style labels become visible on the product's web page and can later be revised by other users. These revisions could also be considered categorization decisions, but they are more complicated because they are based on knowledge of the labels previously assigned by someone else. This creates the possibility of social influence (Bodoff and Vaknin 2016), which is a confounding factor for our study. To minimize social influence and ensure a cleaner test of our hypotheses, we only analyze decisions where products were being categorized on Discogs for the first time—that is, in the absence of labels previously assigned by others. This means that we exclude revisions from our sample. Though it is still possible that users influence each other outside the website, we think this approach reduces the confound as much as reasonably possible.

As in other cultural markets, music genre and style labels encode information about the aesthetic attributes of goods. In recorded music, they primarily encode acoustic or musicological features, which are the axes or dimensions of a feature space wherein product differentiation is perceived by music listeners (Askin and Mauskapf 2017). To represent the meanings of genres and styles in this space, we collect feature data from AcousticBrainz, an open-source service ran by computer scientists in collaboration with music consumers for the purpose of developing better music information software (Porter et al. 2015). Contributors to this service can download a suite of algorithms to create numerical descriptions of music on their computers. These algorithms analyze audio files and automatically submit their features to AcousticBrainz, where they are further processed and made public. At the time of writing, numerical descriptions are available for about 7.5

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

million songs or *tracks*. Each description encompasses thousands of features, sorted into low-level and high-level. Low-level features describe a track's basic acoustic profile, for example, "overall loudness, dynamics and spectral shape of a signal, rhythm descriptors […] and tonal information" (AcousticBrainz 2022). High-level features, instead, describe more sophisticated aspects like mood, timbre, and gender of vocals, estimated from low-level data after its submission to the server. High-level features are considered unstable by AcousticBrainz developers, in the sense that they are extracted by experimental algorithms and can change as a result of technological improvements. Low-level features, instead, are considered stable and insensitive to technological change.

Our analysis is based on a subset of low-level features that AcousticBrainz showcases on every track's web page for the purpose of summarizing its acoustic profile to human readers (see Appendix A). This includes (1) the track's main key, (2) the scale and (3) relative frequency of this main key, (4) the most frequent key of the chord progression, (5) the scale of this most frequent key, (6) an estimate of the track's danceability, (7) the average number of beats per minute, and (8) the total count of beats. We also consider (9) track length as an additional feature. These nine attributes were used in previous research to model the feature space of music (Piazzai and Wijnberg 2019). Considering no more than nine is appropriate because, although music can be a complex good, its feature space should be low-dimensional enough to remain cognitively manageable for humans (Verheyen et al. 2007). Features (1) and (4) are nominal variables taking values from the circle of fifths, that is, A to G, with ♯ denoting sharp notes; (2) and (5) are dichotomous variables corresponding to major or minor scale; (3) is a percentage; (6) and (7) are continuous; and (8) and (9) are integer. For statistical purposes, (1) and (4) are recoded into dummy variables with the key of A as reference, while (2) and (5) are recoded into binary variables with the major scale as reference.

3.1. Sampling Procedure

In merging data from our sources, we preempt possible issues that arise from differences in granularity. These arise because AcousticBrainz provides information at the level of individual songs or tracks, whereas Discogs users categorize records, which normally consist of multiple tracks. To avoid problems of aggregation, we restrict our analysis to singles, which Discogs defines as records that include one main track and possibly a few additional tracks or alternative versions. Because singles have a short duration and heavily revolve around one track, usually to the point of inheriting its title, there is less variance in acoustic features within a single's tracklist compared with longer records like extended-play and full-length albums. As a result, we can create numerical descriptions of singles by taking the mean features of songs in their tracklist without destroying nearly as much information. We can empirically confirm that this approach is sensible because descriptive statistics and pairwise correlations of track-level features (Appendix B, Tables B3 and B4) are nearly identical to those of single-level features we obtain from averaging (Appendix B, Tables B5 and B6).

We start building our sample by retrieving 22,010 singles on Discogs for which a cross-reference exists to one or more tracks on AcousticBrainz. This is not the case for all singles on Discogs: there are many for which no such link exists, and because these are impossible for us to represent in terms of features they are necessarily excluded from analysis. We checked whether this attrition biases our sample by the following procedure. We began by randomly drawing 100 tracks from singles in our preliminary sample whose features were not on AcousticBrainz at the time of data collection. For 62 of these, we were able to retrieve audio files legally available online, and after analyzing them through AcousticBrainz's suite of algorithms, we submitted their features to the AcousticBrainz database. For each track we submitted, we then randomly drew 10 tracks from singles in our preliminary sample whose features were already on AcousticBrainz at the time of data collection. We hence arrived to a random subsample of 620 tracks whose features were already on AcousticBrainz, plus 62 tracks whose features we submitted ourselves. With these data in hand, we performed two graphical comparisons: first, we compared the features of tracks that were already in the database to the features of those we submitted ourselves; second, among tracks that were already in the database, we compared the features of those that were submitted to AcousticBrainz in the same year as their release on the market to the features of those submitted one to five years later. We may consider tracks submitted many years after release to be relatively similar to those that were never submitted, just like late respondents to a survey are considered similar to nonrespondents (Armstrong and Overton 1977). As shown in Figure 1, neither of these comparisons points to distributional differences among acoustic features. We thus concluded that tracks not available on AcousticBrainz are acoustically similar to those that are available, for which reason attrition does not seem to bias our sample.

For each of the 22,010 singles thus retrieved, we browsed the edit history on Discogs to find the first time the product was categorized by a member of the community. To reduce the influence of critical discourse generated around products over time, we required this decision to have been made in the same year as the single's release on the market. This narrowed down our sample to 2,872 decisions made about singles released between the launch of Discogs and data collection (2000–2020). We performed robustness tests using even

Figure 1. Multidimensional Scaling Solution for a Random Sample of Tracks

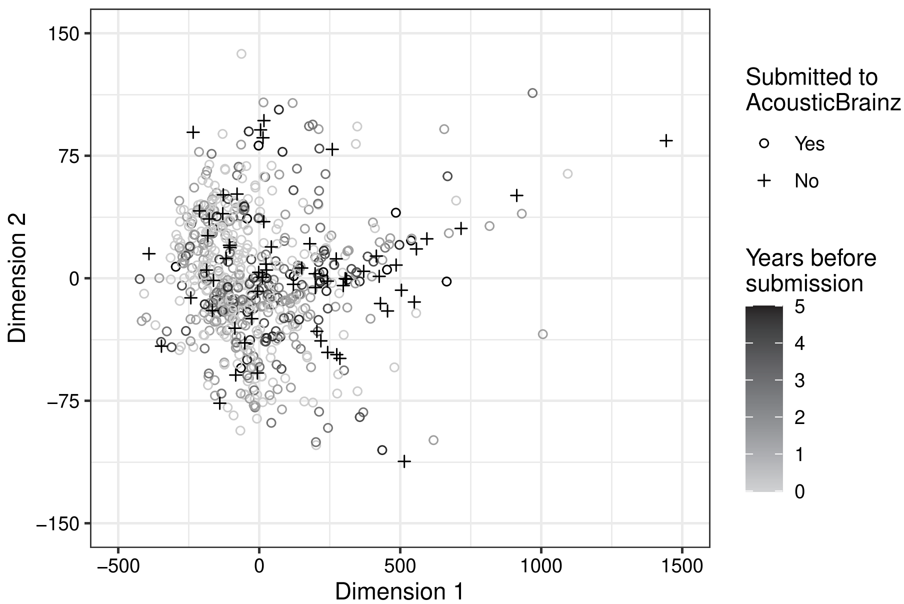

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

more restrictive requirements, limiting our sample to decisions made within six months, one month, or one week of a single's release. As will be shown in our robustness tests, these results are fully consistent with our main analysis. Singles in our sample were categorized into 13 genres and 200 styles in total. Their distribution over genres is relatively uneven, with 70.6-percent having been labeled "electronic," 30.4-percent "rock," and 14.1-percent "pop." This reflects the greater popularity of electronic music during our study period, the low barriers to entry for producers of this genre, and the fact that Discogs was originally born as a community of electronic music enthusiasts. The distribution of singles over styles is comparatively even: the top three styles, all subordinate to "electronic," are "house" (13.6-percent), "electro" (13.4-percent), and "drum'n'bass" (12.4-percent); the fourth and fifth are "indie rock" (11.2-percent) and "alternative rock" (7.8-percent). The singles were categorized by 1,494 users in total, with 72.6-percent of users having made only one decision each, 13-percent having made two, 4.5-percent having made three, 3.5-percent having made four, and so on in smaller numbers. The top three users, who individually made 43, 46, and 63 decisions, jointly account for 5.3-percent of all decisions in the sample. To ensure that these relatively active users do not drive our results, we replicated our analysis after excluding their decisions. The resulting estimates are consistent with our main analysis and reported as robustness tests.

3.2. Statistical Model

Our outcome of interest is a Discogs user's decision to assign a given style label to a single. This is a qualitative choice, which can be appropriately modeled by conditional logit regression (McFadden 1974). The conditional logit model, often used in organizational research to predict strategic decisions (e.g., Carnabuci et al. 2015), is analogous to the logit model except that its estimates are conditional on a choice set. In our case, the choice set includes every style label in the dropdown menu the user sees on Discogs. In the resulting model, every product-label pair becomes an observation, with the dependent variable (*Assignment*) being equal to one if the label was assigned to the product by the user, and zero otherwise. It is possible for the dependent variable to equal one for multiple product-label pairs within the same choice set, representing situations in which multiple style labels were assigned to the same product. Pairs are not formed for labels that were not yet available on Discogs at the time of the user's decision. This leaves us with a total of 228,509 product-label observations. Because we cannot assume a user to pay equal attention to every label in the choice set, we cluster regression errors by user-label pairs. Along with control variables, this clustering contains the impact of users' personal proclivities in favor or against particular labels, which can arise from familiarity (Zunino et al. 2019) and cognitive fluency (Unkelbach 2006).

We remark that, even with clustering, our approach assumes every style label available on Discogs at the time of a user's decision to be considered a candidate for assignment. It could be that users do not actually consider all labels as candidates: for example, they could make a preliminary selection so that, for some elements of the choice set, the probability of assignment approaches zero. Though this might be the case, all the labels remain available for assignment and so outcomes in which they are assigned could be improbable, but

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

not impossible. We believe that retaining every label in the choice set keeps our model faithful to the data-generating process. Nonetheless, the concern may arise that our model includes a large number of product-label pairs for which the probability of assignment is virtually null. To address this concern, we replicated our analysis after including in a user's choice set only those labels that present a minimum fit with the product on the basis of acoustic features, where fit is calculated according to the procedure outlined in the next paragraph. In these models, users' choice sets are limited to the 75-percent, 50-percent, or 25-percent best-fitting labels, so that users are effectively assumed to consider fewer candidates in their decisions. Model estimates, very similar to our main analysis, are reported as robustness tests.

Calculating the fit between a label and a product on the basis of acoustic features requires us to introduce some notion of distance in feature space. This is a kind of informational distance, and in order to compute it, we need to represent both the product and the label in terms of features. For any single, the means of track-level features provide a vector of numerical coordinates, that is, a location in the feature space. Therefore, for each single, we automatically obtain a point-like representation. We assume singles sharing the same genre or style label to be points sampled from the same normal distribution over the feature space, which can be regarded as a function expressing the label's plausibility or fit at any given location. This distribution represents the label's meaning (Hannan et al. 2019). We compute the extent to which a label fits a single by way of Mahalanobis distance, which is commonly used in machine learning to assess goodness-of-fit between a point and a multivariate normal distribution. In essence, the Mahalanobis distance expresses how likely it is for the point to be sampled from the distribution, and is hence analogous to a z-score. More details on the Mahalanobis distance and its statistical properties are included in Appendix C.

To compute the Mahalanobis distance from a genre or style label to a single, we first take the set of singles in our sample to which this label was assigned by Discogs users during the year before the focal categorization decision; then, we compute the means and covariance matrix of these singles' features. We only consider singles categorized during one year before because the meanings of category labels change over time, as producers engage in behaviors that affect the features of products released on the market or their perception by the audience (Durand et al. 2007, Negro et al. 2011). To account for this time-dependence, we do not infer a label's meaning from *all* singles to which the label was ever assigned, but only from those to which it was assigned in the recent past. If a single x and a label Y are considered at year t, the label's meaning is inferred from singles $y_1 \ldots y_N$ to which Y was assigned at $t-1$, and the distance from Y to x at t is:

$$d_t(x, Y) = \sqrt{(x - \overline{y})^T \mathbf{K}_{yy}^{-1} (x - \overline{y})}, \tag{1}$$

where $\overline{y}$ is the mean feature vector of $y_1 \ldots y_N$, and $\mathbf{K}_{yy}$ is the covariance matrix of these features. Intuitively, $\overline{y}$ represents the average member or prototype of the category denoted by Y, and $\mathbf{K}_{yy}$ represents the category's internal correlational structure (Rosch 1978). This contains information about the co-occurrence of features among category members: for example, the fact that "speedcore" music tends to have a relatively short duration and a high number of beats per minute, whereas "ambient" music has a relatively long duration and few beats. Note that the inclusion of $\mathbf{K}_{yy}$ makes Equation (1) fundamentally different from a Euclidean distance. By allowing categories to have not only different prototypes but also different correlational structures, our measure takes into account that any variation in acoustic features that seems trivial for products in a category could be meaningful for products in another, or even for products in the same category at a different time. In this sense, our measure reflects the adaptive nature of categorical reasoning (Anderson 1991).

3.2.1. Specificity and Distinctiveness. Our predictors of theoretical interest relate to distance from one label to another rather than distance from a label to a single. In the case of specificity, this is the distance from a genre to a style nested within this genre, for example, from "electronic" to "house." In the case of distinctiveness, it is the mean distance from a style to other styles nested within the same genre, for example, from "house" to "dubstep," "techno," etc. In their analytical treatment of categories and labels, Hannan et al. (2019) proposed to measure the distance between label meanings on the basis of text data using Kullback-Leibler divergence. This approach is convenient if one happens to know the theoretical distributions representing the meanings of labels, or if these distributions can be easily estimated from data, for example, via topic modeling and other techniques in natural language processing. But this is not our case, because unlike Hannan et al. (2019), we are not dealing with a corpus of text and cannot estimate a topic model. We prefer to use a label-to-label distance that directly builds on the single-to-label distance defined in Equation (1). If two labels Y and Z are considered at year t, we take all singles $y_1 \ldots y_N$ and $z_1 \ldots z_M$ to which Y and Z were respectively assigned at $t-1$, and the distance from Z to Y at t takes the form:

$$D_t(Y,Z) = \frac{1}{N}\sum_{i=1}^{N} d_t(y_i, Z) = \frac{1}{N}\sum_{i=1}^{N} \sqrt{(y_i - \overline{z})^T \mathbf{K}_{zz}^{-1} (y_i - \overline{z})}. \tag{2}$$

In short, this is the mean Mahalanobis distance from Z to $y_1 \ldots y_N$. This measure shares many desirable properties of Kullback-Leibler divergence. First, it is a "statistical

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

distance" (Criado et al. 2011), because it is sensitive to the distribution of points $z_1 \ldots z_M$ thanks to the inclusion of $\mathbf{K}_{zz}$. Second, it is a directed distance, that is, $D_t(Y,Z) \neq D_t(Z,Y)$ in general, because $D_t(Y,Z)$ depends on $\mathbf{K}_{zz}$ while $D_t(Z,Y)$ depends on $\mathbf{K}_{yy}$. This is desirable because similarity judgments between category labels can be asymmetric (Gleitman et al. 1996). In any event, the measure positively correlates with Kullback-Leibler divergence because if the divergence from Z to Y increases, then the Mahalanobis distance from Z to any point $\boldsymbol{y}$ sampled from Y is expected to increase as well. We provide computational evidence of this correlation in Appendix C.

If Y is a style nested within genre Z, Equation (2) can be used directly to measure the specificity of Y at t. This corresponds to our variable *Specificity*. Because membership in Y implies membership in Z, any distance from Z to Y represents extra information that Y adds to what is already encoded by Z. If, instead, Y and Z are styles nested within the same genre, which may also encompass other styles $Z', Z'', \ldots$, then Equation (2) can be used to compute pairwise distances from $Z, Z', Z'', \ldots$ to Y, and the mean of these distances taken to measure the distinctiveness of Y at t. This corresponds to our variable *Distinctiveness*. Recall that greater distinctiveness does not imply a higher amount of information encoded with respect to the superordinate: this depends on specificity, and we expect little correlation between specificity and distinctiveness. In fact, only in the case that all the styles within a genre have low specificity do they also then necessarily have low distinctiveness, as being close to their common genre will also make them close to one another. But as long as their degrees of specificity vary, distinctiveness can vary too. Within our data set, the most specific styles by median value throughout our observation period are "ethereal," "drone," and "avantgarde." Discogs considers "drone" to be a style of "electronic" music, while the other two are listed as styles of "rock." In terms of distinctiveness, the top styles by median value are "jump blues," "soundtrack," and "easy listening." These are considered styles of "blues," "stage and screen" music, and "jazz," respectively.

3.2.2. Control Variables. We aim to take into account, as much as reasonably possible given the limits of our data, any variance in the outcomes of Discogs users' categorization decisions that can be attributed to competing explanations. By requiring decisions to have been made by Discogs users who do not see category labels previously assigned by peers, we have already taken steps to reduce the possibility of social influence. By requiring decisions to have been made in the same year as products' release on the market, we have also limited the impact of critical discourse accumulating around products over time. As explained above, we also cluster regression errors at the level of user-label pairs so as to address cognitive factors that can make specific users more or less inclined toward particular labels. Through control variables we account for three other determinants of users' decisions, all considered important by organizational literature, including products' typicality in candidate categories, producers' claims to category labels, and mediators' expertise.

Typicality is perhaps the most obvious predictor of whether a user assigns a given category label to a single. Cognitive scientists define the typicality of an object in a category as the extent to which the object is a representative member of that category, and measure it by proximity to the category prototype (Hampton 2007). This approach is also common in organizational research (e.g., Smith and Chae 2016). Greater typicality of a product in a category should make the corresponding label a stronger candidate for assignment, because people want to apply labels descriptive of objects' actual features (Rosch 1978). For example, if a piece of music averages 300 beats per minute, it is unlikely to be labeled "ambient" and quite likely to be labeled "speedcore" or some other style of music for which fast tempo is expected. Typicality is commonly understood to be a decreasing function of distance in feature space (Hannan et al. 2019), so it can be measured by a transformation of Equation (1). For every product-label pair, we compute the value of this equation and then take its multiplicative inverse as the value of our control variable *Typicality*.

A more sociologically interesting and perhaps more managerially relevant predictor of categorization decisions is the strength of producer-side signals about which category labels should be assigned to their products. Producers are interested in having mediators make categorization decisions that cast their products in a favorable light: for this reason, they can lay claims to particular labels as part of a category strategy (Verhaal and Pontikes 2022). In cultural markets, these claims are made routinely—though not exclusively (see, e.g., Sgourev et al. 2023)—by way of product names (Verhaal et al. 2015, Khessina and Reis 2016). These provide a direct channel of communication from producers to audiences that is low-cost, fully under producers' control, and practically impossible for mediators to ignore. In our setting, it is possible for the title of a single to incorporate a category label. This seems to occur with some regularity in our sample, where 156 singles include one of the style labels without modification in their titles. In some cases, titles exactly correspond to one of the style labels (e.g., *Black Metal* by Ascii Disko); in others, labels appear in titles with small modifications (e.g., *Elektro* by Outwork), as part of a compound word (e.g., *Gutterpunk* by Noisia), or disguised as wordplay (e.g., *Saxphunk* by Criss Source).

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

We measure the extent to which producers claim a label in a single's title through a string similarity index. Various algorithms exist that can be used to quantify the similarity between two character strings: we use one of the most common in text analysis, termed Optimal String Alignment (OSA), which counts the number of insertions or deletions of any character and transpositions of two adjacent characters required to turn one string into the other. To obtain our control variable *Producer claim*, we compute the OSA index for every product-label pair and then reverse-code its value so that it increases with similarity. We also estimate models where the variable *Producer claim* is computed via different algorithms, including Levenshtein distance, bigram cosine similarity, and Jaro-Winkler distance with 10-percent penalty factor. Compared with OSA, these algorithms involve different operations on strings and do not necessarily yield similar values. Their Pearson's correlation coefficients with the OSA index are 0.999, −0.077, and −0.061, respectively ($p = 0.000$). Results from these models are consistent with our main analysis and reported as robustness tests.

The final possible determinant that we consider is the expertise of decision makers (Cudennec and Durand 2023). Compatibly with our theoretical assumptions, we expect Discogs users to possess sufficient knowledge of the product domain to know the meanings of category labels and hierarchical relations between them. However, we do not expect their knowledge to be evenly distributed over categories. There could be styles of music in which their expertise is greater, for example, because they have a preference for product with those features, consume them more often, and categorize them more frequently. This can influence their categorizations through the availability heuristic (Tversky and Kahneman 1973). The clustering of regression errors by user-label pairs already adjusts the model to account for this heterogeneity, but to provide a more convincing empirical test, we introduce two control variables, *Style expertise* and *Genre expertise*, which respectively count the number of decisions previously taken by the focal user that resulted in the assignment of the focal style, and the number of decisions previously taken by this user that resulted in the assignment of any style connected to the same genre as the focal style. If a user's expertise is localized in certain genre or style categories, this will be captured by variance in these counts.

4. Results

The descriptive statistics and pairwise correlations of variables in our main analysis are summarized in Table 1. Pearson's coefficients do not suggest harmful collinearity among independent variables, and the condition number of the data matrix (25.241) is below the upper bound of 30 that statistical literature considers indicative of collinearity problems (Belsley et al. 1980). All independent variables are mean-centered and standardized before regression, which further decreases the condition number to 1.622. We remark that this transformation has no substantive impact on our estimates because it rescales the variables without altering their distributions. In addition to minimizing collinearity, standardization is useful because it facilitates comparison of effect sizes.

As in ordinary logit regression, coefficients in a conditional logit model correspond to natural logarithms of odds ratios. By exponentiation, they can be converted to odds ratios, which in our case represent changes in the odds of a label's assignment associated with one-standard deviation increases in the value of an independent variable, all else being equal. Because odds and probabilities are directly related, these coefficients allow us to evaluate support for our hypotheses. We estimate five conditional logit models, each nested into the previous, to assess changes in model fit resulting from the inclusion of each additional independent variable. Model estimates are included in Table 2 in their original logarithmic form, along with robust standard errors and conventional notation for thresholds of statistical significance. When reporting these estimates below, we automatically convert them to odds ratios, appending 95-percent confidence intervals (CI) and p-values.

We begin by estimating a model that includes only control variables: *Typicality*, *Producer claim*, *Style expertise*, and *Genre expertise* (Model 1). We then extend the list of regressors with variables of theoretical interest, that is, *Specificity* (Model 2), *Distinctiveness* (Model 3),

Table 1. Descriptive Statistics and Correlations of Variables in Models of Style Assignment

	Variable	Mean	SD	Min.	Max.	(1)	(2)	(3)	(4)	(5)	(6)
(1)	Assignment	0.018	0.133	0	1						
(2)	Typicality	0.032	0.045	0	0.486	0.20***					
(3)	Producer claim	65.239	8.221	0	80	0.01***	0.01***				
(4)	Style expertise	0.923	7.610	0	586	0.22***	0.14***	0.01***			
(5)	Genre expertise	21.449	59.326	0	678	0.04***	0.01***	0.03***	0.37***		
(6)	Specificity	4.938	0.831	1.500	8.554	0.01***	0.04***	0.01**	0.01***	0.04***	
(7)	Distinctiveness	540.617	588.724	1.155	11,275.612	0.00	−0.11***	0.01***	0.03***	0.05***	−0.05***

*$p < 0.05$; **$p < 0.01$; ***$p < 0.001$.

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

Table 2. Conditional Logit Estimates of Style Assignment

	Model 1		Model 2		Model 3		Model 4		Model 5	
Variable	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE	Coef.	SE
Typicality	0.711***	(0.013)	0.712***	(0.013)	0.713***	(0.013)	0.679***	(0.014)	0.679***	(0.014)
Producer claim	0.251*	(0.100)	0.243*	(0.100)	0.245*	(0.100)	0.243*	(0.099)	0.245*	(0.099)
Style expertise	0.261***	(0.022)	0.261***	(0.022)	0.261***	(0.022)	0.257***	(0.022)	0.256***	(0.022)
Genre expertise	0.702***	(0.152)	0.694***	(0.151)	0.695***	(0.152)	0.681***	(0.153)	0.666***	(0.156)
Specificity			0.105***	(0.021)	0.100***	(0.020)	103.551***	(17.038)	86.099***	(17.585)
Specificity^2							−192.664***	(22.217)	−170.297***	(23.351)
Distinctiveness					0.049*	(0.012)	0.147***	(0.021)	−266.689	(221.145)
Distinctiveness^2									−467.131	(248.434)
No. product-label pairs	228,509		228,509		228,509		228,509		228,509	
No. user-label clusters	120,229		120,229		120,229		120,229		120,229	
Log-likelihood	−14,329.6		−14,318.89		−14,316.10		−14,257.55		14,223.87	
Likelihood ratio χ^2			21.423***		5.579*		117.09***		67.365***	

$^*p < 0.05$; $^{**}p < 0.01$; $^{***}p < 0.001$.

Specificity squared (Model 4), and *Distinctiveness* squared (Model 5). We include the squared terms of *Specificity* and *Distinctiveness* only at the end to better assess changes in model fit resulting from the use of quadratic polynomials. To track model fit, we perform likelihood-ratio tests between every pair of consecutive models. The resulting χ^2 values are reported at the bottom of Table 2, along with their levels of statistical significance. These results point to significant increases in log-likelihood following the addition of each independent variable ($p \leqslant 0.018$). Note that the increase is especially substantial in the test of Model 4 vs. 3, which suggests that, although Model 3 provides evidence of a positive relationship between *Specificity* and *Assignment*, the inverse U-shaped relationship specified in Model 4 better explains the data. This will be corroborated in the course of our analysis by Lind and Mehlum's (2010) three-step test for an inverse U-shape. A relatively high increase in log-likelihood is also evidenced in the test of Model 5 vs. 4, but here it rather seems a consequence of overfitting, as will become apparent by Lind and Mehlum's (2010) test. Estimates are stable in terms of sign and magnitude throughout Models 1 to 4, so we look directly at Model 4.

Both *Typicality* and *Producer claim* appear to have positive and significant relationships with the dependent variable. In particular, if *Typicality* increases by one standard deviation, the odds of assignment for the focal label change by a factor $e^{0.679} = 1.972$, which implies a 97.2-percent increase (CI = [92.1, 102.6], $p = 0.000$). This is consistent with our expectation that products located closer to the average for a particular music style are more likely to be categorized into that style. A smaller but nonetheless significant increase occurs in the case of *Producer claim*: if this variable increases by one standard deviation, the odds of label assignment increase by 27.5-percent (CI = [4.9, 54.9], $p = 0.015$). This suggests that Discogs users respond to signals they receive from artists or record companies about the styles in which their music should be categorized. In fact, they are more likely to assign whatever label is cued by a single's title. Because this effect persists with *Typicality* in the model, it seems to occur regardless of how accurately cued labels describe a product's features. We find this indicative of leeway that producers enjoy to strategically influence the categorization of their products. With regard to expertise variables, we also find positive and significant relationships. In particular, if *Style expertise* increases by one standard deviation, the odds of label assignment increase by 29.4-percent (CI = [23.9, 35.0], $p = 0.000$), and if *Genre expertise* increases by one standard deviation, the odds increase by 97.6-percent (CI = [46.3, 166.9], $p = 0.000$). This could be a consequence of self-selection, in the sense that Discogs users tend to categorize products that make good candidates for categories they know especially well.

It remains to be ascertained whether *Specificity* and *Distinctiveness* have the hypothesized relationships with label assignment. In Model 4, both variables have highly significant coefficients, but while interpretation is straightforward for *Distinctiveness*—a one-standard deviation increase in the value of this variable is linked to a 15.9-percent increase in the odds of assignment (CI = [11.3, 20.7], $p = 0.000$)—for *Specificity* it is rendered more difficult by the variable's polynomial specification. In this case, it is advisable to implement a formal test of the inverse U-shape. Statistical significance for the coefficients of *Specificity* and *Specificity* squared in Model 4 is not a sufficient condition, "as the estimated [turning] point may be too close, given the uncertainty, to an end point of the data range" (Lind and Mehlum 2010, p. 115), in which case the relationship between *Specificity* and *Assignment* could be better modeled by a logarithmic curve. A formal test of an inverse U-shape against simpler functional forms can be performed by Lind and Mehlum's (2010) three-step procedure (see also Haans et al. 2016). According to this procedure, three conditions must hold for an inverse U-shape to be

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

empirically supported: first, the coefficient of the variable's squared term should be negative and significant; second, the slope should be significant at both ends of some meaningful interval $[L, H]$ within the variable's range, more concretely positive at L and negative at H; third, the 95-percent confidence interval of the turning point should be included in $[L, H]$. The choice of values L, H is important: these can be set to the variable's minimum and maximum, so that $[L, H]$ is the variable's range, but "if we want to make sure that the inverse U-shape is not only a marginal phenomenon the interval could also be in the interior of the [range]" (Lind and Mehlum 2010, p. 111).

We can check if these conditions hold for *Specificity* in Model 4. The first condition is satisfied because the variable's squared term has a negative and significant coefficient ($z = -8.672$, $p = 0.000$). To check the second, we set L, H to the first and 99th percentiles of the variable's distribution on the mean-centered and standardized scale, which correspond to $L = -3.216$ and $H = 2.586$. As a result, the interval $[L, H]$ contains the central 98-percent of observations. We test the slopes of the curve at L and H using two-tailed t-tests. The resulting statistics t_L and t_H are both significant and of the appropriate sign ($t_L = 8.860$, $p = 0.000$, and $t_H = -8.263$, $p = 0.000$), so the second condition is also satisfied. To check the third, we estimate the 95-percent confidence interval of the turning point using Fieller's (1954) method, as recommended by Lind and Mehlum (2010). We obtain the interval $[0.192, 0.355]$, which is a subset of $[L, H]$, so the third condition is satisfied as well. Because all three conditions hold, there is evidence of an inverse U-shaped relationship between *Specificity* and *Assignment*. This evidence persists in Model 5 as the coefficient of the variable's squared term remains negative and significant ($z = -7.293$, $p = 0.000$), the slope is positive at the first percentile ($t_L = 7.400$, $p = 0.000$) and negative at the 99th ($t_H = -7.027$, $p = 0.000$), and the turning point's interval $[0.165, 0.351]$ is still included in $[L, H]$. Therefore, Hypothesis 1 is supported.

The purpose of Model 5 to assess whether a similar inverse U-shaped relationship exists for *Distinctiveness*, as predicted in Hypothesis 2b, or if we should stick to the positive relationship estimated in Model 4, which is predicted in Hypothesis 2a. It is immediately clear from Table 2 that an inverse U-shape does not exist according to Lind and Mehlum's (2010) three-step test, despite the significant increase in log-likelihood, because the coefficient of *Distinctiveness* squared is negative but not significant ($z = -1.880$, $p = 0.060$). Therefore, the first condition required by Lind and Mehlum (2010) is violated. This would be already enough to reject Hypothesis 2b, but we can go on to check the second and third conditions. To check the second, we set $L = -0.873$ and $H = 3.726$. As before, these correspond to the first and 99th percentiles of the variable's distribution on the mean-centered and standardized scale. We find a positive slope at L but no significant slope at H ($t_L = 2.561$, $p = 0.010$, and $t_H = -1.809$, $p = 0.070$), which means that the relationship does not turn negative throughout the central 98-percent of the variable's distribution. Hence, the second condition is violated as well. Finally, to check the third, we estimate the turning point's interval and obtain $[-\infty, +\infty]$, which obviously extends outside $[L, H]$, so the third condition is violated too. These test results do not change if we set L, H to the variable's minimum and maximum, that is, $L = -0.916$ and $H = 18.234$ on the mean-centered and standardized scale, because the slope at H is still not significant ($t_L = 2.501$, $p = 0.012$, and $t_H = -1.864$, $p = 0.062$) and the turning point's interval remains $[-\infty, +\infty]$. Therefore, an effect reversal does not even occur as a marginal phenomenon. We conclude that our data supports Hypothesis 2a but not Hypothesis 2b.

4.1. Robustness Tests and Additional Analyses

We estimate a number of additional models to check the sensitivity of our findings to some of our methodological choices. We start by evaluating their robustness to alternative model specifications and sampling procedures. Tables D7 to D10 in Appendix D summarize the results of our robustness tests. In particular, Table D7 reports estimates from models in which the sample is limited to categorization decisions that occurred within six months (Model 6), one month (Model 7), or one week (Model 8) of the focal product's release. Table D8 reports estimates from a subsample where categorization decisions made by the three most active users are excluded (Model 9). Table D9 reports estimates from subsamples where users' choice sets are restricted to category labels for which the focal product's typicality is in the upper 75th (Model 10), 50th (Model 11), or 25th percentile (Model 12). Finally, Table D10 reports estimates from models where producers' claims to category labels are measured by different string similarity indices, including Levenshtein distance (Model 13), bigram cosine similarity (Model 14), and Jaro-Winkler distance (Model 15). For each of these models, we replicate Lind and Mehlum's (2010) three-step test for an inverse U-shaped relationship between *Specificity* and *Assignment*, and append the results of this test at the bottom of the table. Support for Hypotheses 1 and 2a holds everywhere except in Model 12, where the 75-percent reduction in sample size severely truncates the range of *Specificity* and causes Lind and Mehlum's (2010) test to fail.

Next, we check the robustness of our estimates to an alternative explanation for the inverse U-shaped effect of *Specificity*. We predict an inverse U-shape because specificity generates cognitive costs, and people tend to avoid cognitive costs if they already have enough information to "satisfice." As a result, we expect mediators to

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

assign a highly specific label with lower probability than a moderately specific one. An alternative reason to expect the very same pattern, however, is that a highly specific label is applicable to a small set of products, and this makes it less useful for the purpose of communication between mediators and consumers. Indeed, in order to be useful, a label must allow consumers to infer a product's features, but consumers cannot do this if they do not know the label's meaning. They are unlikely to know the meanings of labels that include very few products, simply because they are unlikely to have come across them. In this study, we assume mediators to know the meanings of category labels (Assumption 2) but we do not extend the same assumption to consumers. Therefore, mediators could be concerned that consumers are unable to interpret a highly specific label, and for this reason be disinclined to assign it.

To verify if our main results hold after taking this alternative explanation into account, we estimate a model where the probability that consumers know the meaning of a candidate label is approximated by the count of products to which the label was assigned before to any product by any Discogs user (*Past assignments*). The greater this count, the greater the probability that consumers have come across some products with this label. If Discogs users refrain from assigning a label because they are concerned that ordinary consumers do not know it, then we should find a positive relationship between *Past assignments* and our dependent variable. This is exactly what we find, as shown in Table D11 (Model 16). However, the inverse U-shaped relationship between *Specificity* and *Assignment* continues to hold, as confirmed by Lind and Mehlum's (2010) test. It thus seems that, while mediators' concern about the intelligibility of labels for consumers affects their categorization decisions, it cannot fully explain the reversal of *Specificity*'s effect.

Beyond robustness to alternative model specifications, we also aim to check whether our findings depend on two fundamental analytical choices made in our study, namely our choice to model the categorization of products into styles and not genres, and our choice to measure the representativeness of products in each style by way of spatial proximity to category averages, or prototypes. With regard to the former choice, concerns may arise because, although styles tend to be more relevant to music enthusiasts (Montauti and Wezel 2016, Formilan and Boari 2021), categorization also occurs at the level of genres. Appendix E provides descriptive statistics and pairwise correlations for variables computed at the level of genres (Table E12), as well as conditional logit estimates from a model that replicates the specification of Model 4 at this level of analysis (Table E13, Model 17). The control variable *Style expertise* does not figure in this model because user-label pairs are created for genres, not styles. Aside from this distinction, estimates are similar to those from Model 4. The inverse U-shaped relationship between *Specificity* and *Assignment* is confirmed by Lind and Mehlum's (2010) test, as shown at the bottom of Table E13. This reassures us that support for Hypotheses 1 and 2a extends to categorization decisions made at a higher level of abstraction.

With regard to the other choice mentioned above, concerns may arise because cognitive theories of categorization do not necessarily consider typicality, as measured by proximity to a prototype (Hampton 2007), to be the sole or even the foremost indicator of an object's representativeness as a potential category member. Some studies argue that proximity to the most salient members of the category, termed exemplars, provides a more accurate measure of representativeness (Medin and Schaffer 1978). Thus, exemplar-based models of categorization are sometimes considered alternative to models based on prototypes (Smith and Minda 2002). In recent years, exemplar-based models also garnered attention in organizational analysis (Zhao et al. 2018, Barlow et al. 2019). Appendix F (Table F14) presents results from models where we include *Exemplarity*, or proximity to exemplars, as an additional control variable that either substitutes (Model 18) or complements (Model 19) the variable *Typicality*. In this appendix we also describe supplementary data collected to reliably identify style exemplars and detail our approach to calculating the *Exemplarity* variable. Model estimates continue to support Hypotheses 1 and 2a.

5. Discussion

What makes a mediator more likely to assign a particular category label to a product? Existing research in organization theory tends to sidestep this question, keeping mediators' categorization decisions confined to the explanatory side of conceptual and empirical models that explain organizational or product-level outcomes. In this study, however, we framed mediators' categorization decisions as the outcome to be explained. In doing so, we arrived at multiple and concurrent answers to the question above. First, mediators are more likely to assign labels previously used for products that possess features similar to the product at hand. Second, mediators are more likely to assign labels already claimed by producers. Third, mediators are more likely to assign labels whose meaning they know more closely. These three answers resonate with existing literature on the cognitive and strategic foundations of categorization in markets (Barlow et al. 2019, Hannan et al. 2019, Cudennec and Durand 2023). Yet, unlike existing literature, our analysis points to a new and complementary explanation: mediators are more likely to assign labels that are moderately specific and maximally distinctive. We find that labels are assigned with greater probability if they encode neither too little nor too much information with respect to their superordinate, and if they

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

encode different information with respect to other labels under the same superordinate. This suggests that mediators' categorization decisions are affected by informational properties that have nothing to do with the product under consideration. Beyond product features, producer signals, and personal expertise, mediators tend to assign labels that encode information in a cognitively economical fashion.

Our finding that assignment is more likely for labels that are maximally distinctive, as opposed to moderately so, may initially appear at odds with Lo et al. (2020) argument that categories should have moderate distinctiveness to be attractive tools for categorization. In this study, we do not find evidence of distinctiveness having to be moderate. There are, in our view, two possible reasons for this apparent misalignment with Lo et al. (2020) proposition and with optimal distinctiveness literature more generally. One is that optimal distinctiveness theory may not extend to our setting because of key differences in the units and outcomes analyzed. Indeed, research on optimal distinctiveness normally concerns the distinctiveness of products or producers, and examines consequences for competitive performance (Zuckerman 2016, Zhao et al. 2017). In contrast, our study concerns the distinctiveness of category labels and examines consequences for label assignment. On the one hand, arguments developed for products and producers may not generalize to labels; on the other hand, these arguments may not explain label assignment as well as they explain more conventional competitive outcomes. Categorization is certainly a competitive process, as labels vie for finite resources in the mind of decision makers, but competition takes a far more abstract meaning in this case and comparisons with products or producers are not straightforward.

Another possible reason why our results appear at odds with optimal distinctiveness literature relates to how we empirically measured distance in a feature space. As explained above, we applied a measure based on Mahalanobis distance, which is commonly used in machine learning to quantify the distance between a point and a multivariate distribution (see also Appendix C). However, literature on optimal distinctiveness mostly opts for a simpler approach based on the Euclidean distance, which equals the sum of squared differences between two points' coordinates (e.g., Zhao et al. 2018). Although some studies on optimal distinctiveness deviate from this default approach (e.g., Haans 2019, Taeuscher and Rothe 2021, Taeuscher et al. 2022), they nonetheless use measures that, like Euclidean distance, involve a sum of coordinate-wise differences. The Mahalanobis distance used in our study differs radically from Euclidean distance and similar functions because it involves multiplication by a variance-covariance matrix. This serves to ensure that coordinate-wise differences are weighed so as to be discounted when they are expected, as in the case of two positively correlated features that change in the same direction, and magnified when they are unexpected, as in the case of two positively correlated features that change in opposite directions. We think this is a more refined approach to measuring distances in feature space, but it also represents a methodological variation that could make empirical results more difficult to benchmark.

Despite these important differences, we believe that interesting points of contact exist between our theoretical framework and optimal distinctiveness research. We think there is, at least conceptually, a parallel to be drawn between the research agenda of optimal distinctiveness theorists and the general objective of our study. Barlow et al. (2019, p. 1220) argued that one of the top priorities for optimal distinctiveness research is to clarify "whom or [what] organizations conform to and differentiate from to gain a competitive advantage." Our analysis rests on a similar question: what should category labels be similar to or different from to be more frequently chosen for assignment by mediators? Along the vertical axis of a classification system, where the reference is a superordinate label, a balance between similarity and differentiation is conducive to assignment. Along the horizontal axis, where the reference is the set of other labels connected to the same superordinate, maximal differentiation is instead conducive to assignment. Interestingly, Barlow et al. (2019) also detailed conditions in which products benefit from being maximally distinctive with respect to other products in their category: therefore, a strictly positive relationship between distinctiveness and competitive success finds precedent in optimal distinctiveness literature. In light of this, the results of our analysis could be more compatible with this literature than is initially apparent.

Our finding that maximal distinctiveness positively relates to label assignment is intuitively consistent with ecological research on idea diffusion. This research argued that, among established ideas, low similarity with other ideas is helpful to sustain popularity (Denrell and Kovács 2020). A parallel reasoning seems to apply to institutionalized categories, such as the ones we analyze in this study, because we find that low similarity to related categories at the same level of abstraction, i.e., high distinctiveness, is helpful to increase the probability of label assignment. At the same time, we find that moderate similarity to superordinate categories, that is, moderate specificity, increases the probability of label assignment, so the effect of similarity is nonmonotonic. This goes to show that it is not just similarity that matters but also, and crucially, similarity to what. The same question may be relevant to studies on idea diffusion: the popularity of ideas, like technological inventions (Kovács et al. 2021) or academic concepts (Denrell and Kovács 2020), could depend on similarity in ways that are contingent on a hierarchical structure in which ideas

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

are embedded. It could be that between ideas connected to the same superordinate, similarity should be low, but between ideas such that one is subordinate to the other, some degree of similarity benefits the subordinate.

5.1. Theoretical Implications

Our study shows that category labels' specificity and distinctiveness affect categorization decisions made by mediators in a product market. This bears implications for various streams of literature in organization theory, including literature on strategic categorization and literature on category viability. With regard to strategic categorization, our study speaks to a growing body of research that examines how producers can best position their offerings within product categories (Pontikes 2018, Barlow et al. 2019, Verhaal and Pontikes 2022). The results reported here have potential to change how this research conceives the limits of producers' strategic agency. Indeed, previous studies suggested that producers can strategically influence audience members' categorization decisions by claiming labels in product names or descriptions (e.g., Barlow et al. 2019) and implementing product features that are typical of those labels (e.g., Verhaal et al. 2015). This can be effective: however, specificity and distinctiveness also play a role, because a label's baseline probability of being used for categorization tends to be low if that label is unspecific, too specific, or indistinctive. Inferences about the determinants of audience members' categorization decisions can be misleading if researchers fail to account for specificity and distinctiveness. For example, it could be that audience members do not use the category labels intended by producers, despite explicit claims and carefully chosen product features; and this is not because producers' claims and feature choices are ignored but because the assignment of some label is, and remains, a low-probability event due to specificity and distinctiveness. In this case, the effects of claims and typicality can be underestimated. Alternatively, it could be that decision makers are particularly inclined to use category labels intended by producers due to specificity and distinctiveness, but their inclination is misattributed to the strength of producers' claims or the typicality of product features. Thus, these effects can be overestimated.

Our core message to researchers interested in strategic categorization is that an organization's capacity to strategically manipulate the categorization of its products is constrained by the values of the two informational properties we analyze here. This represents a contribution for this stream of literature because, when considering constraints on strategic categorization, previous studies focused on economic factors like resource endowments and mobility frictions (Cattani et al. 2017), or social factors like audience engagement (Verhaal and Pontikes 2022), as opposed to informational factors like category labels' specificity and distinctiveness. The influence of these factors leads us to believe that the scope of producers' agency could be more limited than existing literature suggests.

Next, in connection to category viability, our study caters to organization theorists' longstanding interest in the persistence or durability of categories (Lounsbury and Rao 2004, Pontikes and Barnett 2015, Rhee et al. 2017). Recent studies in this stream developed the notion of viability as a latent construct that dictates the timing of a category's life cycle, arguing that categories are more likely to endure if they stay viable for new products or firms (Lo et al. 2020, Soublière et al. 2023). In order to stay viable, however, a category needs to be used by the audience with some probability as a sorting device; and above all, it needs to be used by mediators, because their judgment has a disproportionate influence on consumer behavior. The more probable a category label's assignment to a product by mediators, the more often this label enters comparisons with other labels, reinforcing its embedding in a shared cognitive network that legitimates its existence (Rosa et al. 1999). Such legitimacy makes the category attractive to new entrants (Kennedy 2008). Our study connects with the nascent literature on category viability by asking what makes category labels more likely to be assigned to products by mediators. We can view the probability of a category label's assignment as a positive correlate of viability: from this perspective, our results suggest that categories stay viable for new product entry if they retain moderate specificity and maximal distinctiveness.

A more general implication for literature on category viability is that legitimating processes sustaining a category's existence occur not only at a macro level of analysis, as in the case of producers banding together to stimulate change in institutionalized classification systems (Rao et al. 2003, Ozcan and Gurses 2018, Slavich et al. 2020), but also at a micro level, as part of everyday decisions made by individuals about products or firms. The range of arguments that can be mobilized by organization theorists to explain why market categories endure can be consequently expanded: beyond conventional sociological argumentation, researchers can build on a wealth of psychological evidence that details aspects of human cognition, including perception, memory, heuristics, and their impact on decision making. This micro-level approach is novel to the literature because, so far, the focus on macro-level sociological explanations has been overwhelming. As organization scholars now consider the notion of category viability, we think there is an unprecedented opportunity to bridge sociological and cognitive perspectives. Our study makes an attempt by leveraging cognitive arguments to clarify the determinants of product categorization by mediators.

5.2. Empirical Implications

The effects of specificity and distinctiveness on mediators' categorization decisions also have consequences for

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

empirical analysis. This is because research in organization, as well as cognate fields like strategic management, regularly uses category labels assigned to products by mediators to compute variables intended to capture aspects of an organization's product strategy or positioning. For example, Montauti and Wezel (2016) used genre and style labels on Discogs to track the entry of new record companies in particular segments of the market for recorded music. Similarly, Piazzai and Wijnberg (2019) used genre labels assigned to products on Discogs to track record companies' movements within or across market segments. In these cases, labels assigned to products by mediators were considered indicative of product features and used to measure firm-level constructs. Because of specificity and distinctiveness, however, this approach can lead to biased inferences.

Suppose we are interested in measuring how firms change the characteristics of products in their portfolio over time. To that end, we turn to archival sources that provide historical information about firms' products, including the category labels assigned to them by mediators. We reasonably assume that different labels map to different configurations of product features: therefore, if labels assigned to a firm's products change over time, the firm must be changing some of its products' features. Now suppose we observe that a particular firm, for example, a record company whose products were mostly categorized as "synth-pop" at a given time, tends to release new products at a later time that are labeled "Europop." We may be tempted to infer that the firm started to release a different kind of music, but is our reading of the situation correct? Not necessarily. It is possible that the firm changed absolutely nothing about the music it releases, and instead, it is the specificity or distinctiveness of "synth-pop" that shifted, causing this label to lose ground to "Europop" in mediators' categorization decisions. This can induce bias in our data sources, as some of the labels will be more likely to appear in them, as a result of which any inference we might draw on the basis of our measurements could be biased as well. If we meant to analyze how changes in a firm's product portfolio relate to firm survival, for instance, we could grossly misestimate the relationship between change and survival because we simply do not have a valid measure of change.

We believe, in summary, that if informational properties of category labels such as specificity and distinctiveness lead mediators to prefer particular labels, irrespective of product features, then treating label assignments as indicative of product or firm-level characteristics can be inaccurate. In this sense, our study points to a problem inherent to the decision process that generates the data organization scholars collect from archival records. This is a difficult problem to solve, as it ultimately stems from selective availability of historical information (Denrell and Kovács 2008). Solving it may require access to a database of decisions made by actors who are unconstrained by cognitive costs, which is hard to find. Alternatively, corrections may be implemented that account for labels' differential probabilities of being represented in a biased database. At any rate, given the number of studies in our field that use archival records of categorization decisions to measure characteristics of products and firms, being at least aware of the problem is important.

5.3. Practical Implications

In addition to theoretical and empirical aspects, our study has practical implications for managers of organizations. There are many situations where managers can benefit from knowing how mediators make categorization decisions: for example, they can use this knowledge to make better choices when it comes to marketing and commercializing their products. In fact, while managers cannot entirely control which labels get assigned to their products, they have power to decide how to position these products in the eyes of mediators. Suppose a manager can choose between two equally suitable labels for a new product's marketing campaign, but while one of the labels has moderate specificity and high distinctiveness, the other is unspecific, too specific, or indistinctive. In this case, the manager should choose the label with moderate specificity and high distinctiveness, because mediators' resistance to the other means that the marketing campaign will need to be stronger and more costly in order to get the product's label accepted. For similar reasons, an R&D manager choosing between two equally suitable classes for the purpose of filing a new patent application should opt for the class with moderate specificity and high distinctiveness. Patent examiners will be more inclined to assign the same class and deem the submission appropriate.

There is something else managers could do in light of our findings. Given the role that specificity and distinctiveness play in mediators' categorization decisions, learning about the specificity and distinctiveness of labels available to mediators should be an important step in the formulation of a "category strategy" (Pontikes 2018). This involves deciding where the firm will position itself and its products in the market, trying to convert a favorable position into long-term value. Crucially, the success of any such strategy depends on audiences, and especially mediators, assigning exactly the category labels the firm's managers foresee. As an aid to estimating the chance of success of a candidate strategy, managers can use simple experiments to gauge category labels' specificity and distinctiveness. Highly engaged consumers could be called to a laboratory and asked to make inferences about products, answering questions about what features these products are likely to possess given particular labels. The extent to which these inferences change as a result of providing one label or

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

another can be used to measure specificity and distinctiveness (Hannan et al. 2019). With these measures in hand, managers could rank candidate labels for a product according to their values of specificity and distinctiveness, knowing that moderate specificity and maximal distinctiveness increase a label's probability of being assigned. It then becomes possible to formulate a category strategy that takes these odds into account. Candidate strategies could be evaluated depending on the categorization decisions they necessitate mediators to make. If a strategy hinges on mediators assigning labels that appear improbable, for example, because they are unspecific, too specific, or indistinctive, it may be prudent to set this strategy aside. Conversely, if a strategy hinges on mediators assigning labels that are moderately specific and maximally distinctive, this strategy could be recommended, because mediators are more likely to prefer labels with such properties anyway.

5.4. Limitations and Further Research

We conclude with some reflections on the limits and scope conditions of our study. Our theoretical arguments concern situations where mediators assign category labels to products from a pool of possible options, that is, a choice set. Our hypotheses rest on the assumption that this choice set is fixed in advance, and moreover, that mediators know the meanings and hierarchical relations of category labels comprising it (Assumption 2). We believe this assumption holds in our empirical setting because the choice set is determined by an online system, and users of this system are expert enough on average to know the meanings of labels and their relative positions in a classification hierarchy. We believe this assumption also holds in many settings familiar to organizational research, such as the markets for beer (Verhaal et al. 2015), feature films (Hsu 2006, Hsu et al. 2009), and mobile apps (Barlow et al. 2019), but it may not hold in markets where the classification system is emergent (Ruef and Patterson 2009) or contested (Negro et al. 2011, Jones et al. 2012). In these cases, the choice set is not fixed in advance and category meanings are not necessarily known to mediators, even if they are expert, which limits the generalizability of our findings.

In addition to markets where a classification system is not yet established, Assumption 2 could fail in markets where the classification system is established, but mediators do not know enough about it. This can occur if mediators are novices, and their knowledge of label meanings and hierarchical relations is insufficient for labels' specificity and distinctiveness to show their influence. In these cases, it is possible that specificity and distinctiveness have a muted effect or no effect at all. Alternatively, it could be that the effects of specificity and distinctiveness become even more pronounced, because expert decision makers have an advantage over novices when it comes to bearing the cognitive costs of information encoded by category labels. From this perspective, we might expect a decrease in the optimal level of specificity. Moreover, novices could be less forgiving toward indistinctive labels, making the relationship between distinctiveness and label assignment even steeper. Both of these alternatives seem plausible to us. Studying how mediators' level of expertise moderates the effects of specificity and distinctiveness is an interesting direction for follow-up research.

Another limitation of our analysis is that, while we proposed cognitive economy as the mechanism underlying the effects of specificity and distinctiveness, we cannot conclusively show this is the reason for the observed patterns. We can only make predictions consistent with this cognitive argument and offer evidence in support of these predictions. But unfortunately, we cannot entirely rule out alternative mechanisms that lead to the same predictions. One such mechanism was addressed in our robustness tests, but to lend stronger empirical support to our argument that cognitive economy drives the effects of specificity and distinctiveness, future research should consider an experimental design. Incidentally, this would also allow one to examine more direct repercussions of cognitive economy on consumer behavior, including product discovery and evaluation.

Finally, our analysis could not consider mediators' personal agendas or motivations for making categorization decisions. We characterized their assignments as a function of specificity, distinctiveness, producer claims, product typicality, and personal expertise, but if additional factors are at play that prompt them to reject or deliberately misapply particular labels, such as social identity concerns or allegiance to certain categories, this is not reflected in our analysis. Furthermore, our analysis did not closely consider the temporality of categorization decisions. We limited our sample to decisions taken within some time after a product's release, but variance in the delay with which products are categorized would be interesting to explain as well. We also restricted our sample to the very first decisions taken about products on Discogs: however, subsequent decisions could be worth studying as telltale signs of disagreement or shifting consensus about the meaning of labels. These aspects of the categorization process represent interesting questions for future research to address.

Acknowledgments

The authors thank their Senior Editor, Jerker Denrell, and three anonymous reviewers for excellent feedback. All remaining omissions and mistakes are the authors' own. The authors also thank Alex van Venrooij, Filippo Carlo Wezel, and all participants of the 2018 Nagymaros Workshop for helpful comments.

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

Data Availability

The digitally-shareable data necessary to reproduce the reported results is available at https://github.com/piazzai/os-ms-21-15751.

References

AcousticBrainz (2022) AcousticBrainz data. Accessed May 30, 2022, https://acousticbrainz.org/data.

Alexa (2022) Discogs.com. Accessed April 15, 2022, https://www.alexa.com/siteinfo/discogs.com.

Anderson JR (1991) The adaptive nature of human categorization. *Psych. Rev.* 98(3):409–429.

Armstrong SJ, Overton TS (1977) Estimating non-response bias in mail surveys. *J. Marketing Res.* 14(3):396–402.

Askin N, Mauskapf M (2017) What makes popular music popular? Product features and optimal differentiation in music. *Amer. Sociol. Rev.* 82(5):910–944.

Barlow MA, Verhaal JC, Angus RW (2019) Optimal distinctiveness, strategic categorization, and product market entry on the Google Play app platform. *Strategic Management J.* 40(8):1219–1242.

Barlow MA, Verhaal JC, Hoskins JD (2018) Guilty by association: Product-level category stigma and audience expectations in the U.S. craft beer industry. *J. Management* 44(7):2934–2960.

Belsley DA, Kuh E, Welsch RE (1980) *Regression Diagnostics: Identifying Influential Data and Sources of Collinearity* (Wiley, New York).

Bodoff D, Vaknin E (2016) Priming effects and strategic influences in social tagging. *Human Comput. Interact.* 31(2):133–171.

Bogdanov D, Herrera P (2012) Taking advantage of editorial metadata to recommend music. Aramaki M, Barthet M, Kronland-Martinet R, Ystad S, eds. vol. 7900, *Lecture Notes in Computer Science* (Springer, Berlin), 618–632.

Bowker GC, Star SL (2000) *Sorting Things Out: Classification and Its Consequences* (MIT Press, Cambridge, MA).

Brewer MB (1991) The social self: On being the same and different at the same time. *Pers. Soc. Psych. Bull.* 17(5):475–482.

Caplin A, Dean M, Martin D (2011) Search and satisficing. *Amer. Econom. Rev.* 101(7):2899–2922.

Carnabuci G, Operti E, Kovács B (2015) The categorical imperative and structural reproduction: Dynamics of technological entry in the semiconductor industry. *Organ. Sci.* 26(6):1734–1751.

Cattani G, Porac JF, Thomas H (2017) Categories and competition. *Strategic Management J.* 38(1):64–92.

Chi MT, Feltovich PJ, Glaser R (1981) Categorization and representation of physics problems by experts and novices. *Cogn. Sci.* 5(2):121–152.

Criado C, Alamo N, Rabal H, Cap N (2011) Decision and classification problems using Mahalanobis statistical distance. Ding Y, Li Y, Fan Z, Li S, Wang L, eds. *Proc. 8th Internat. Conf. Fuzzy Systems Knowledge Discovery* (IEEE, Piscataway, NJ), 1012–1016.

Cudennec A, Durand R (2023) Valuing spanners: Why category nesting and expertise matter. *Acad. Management J.* 66(1):335–365.

Deephouse DL (1999) To be different, or to be the same? It's a question (and theory) of strategic balance. *Strategic Management J.* 20(2):147–166.

Denrell J, Kovács B (2008) Selective sampling of empirical settings in organizational studies. *Admin. Sci. Quart.* 53(1):109–144.

Denrell J, Kovács B (2020) The ecology of management concepts. *Strategy Sci.* 5(4):293–310.

DiMaggio P (1987) Classification in art. *Amer. Sociol. Rev.* 52(4): 440–455.

Discogs (2018) Overview of how Discogs is built. Accessed May 29, 2022, https://support.discogs.com/hc/en-us/articles/360008545114-Overview-Of-How-Discogs-Is-Built.

Durand R, Khaire M (2017) Where do market categories come from and how? Distinguishing category creation from category emergence. *J. Management* 43(1):87–110.

Durand R, Rao H, Monin P (2007) Code and conduct in French cuisine: Impact of code changes on external evaluations. *Strategic Management J.* 28(5):455–472.

Epley N, Keysar B, Van Boven L, Gilovich T (2004) Perspective taking as egocentric anchoring and adjustment. *J. Pers. Soc. Psych.* 87(3):327–339.

Faraj S, Jarvenpaa SL, Majchrzak A (2011) Knowledge collaboration in online communities. *Organ. Sci.* 22(5):1224–1239.

Ferguson J-P, Carnabuci G (2017) Risky recombinations: Institutional gatekeeping in the innovation process. *Organ. Sci.* 28(1):133–151.

Ferrer-i-Cancho R, Del Prado Martín FM (2011) Information content vs. word length in random typing. *J. Stat. Mech. Theory Exp.* 2011(12):L12002.

Fieller EC (1954) Some problems in interval estimation. *J. R. Stat. Soc. B.* 16(2):175–185.

Formilan G, Boari C (2021) The reluctant preference: Communities of enthusiasts and the diffusion of atypical innovation. *Indust. Corporate Change* 30(3):823–843.

Gleitman LR, Gleitman H, Miller C, Ostrin R (1996) Similar, and similar concepts. *Cognition* 58(3):321–376.

Goldberg A, Hannan MT, Kovács B (2016) What does it mean to span cultural boundaries? Variety and atypicality in cultural consumption. *Amer. Sociol. Rev.* 81(2):215–241.

Goldenstein J, Hunoldt M, Oertel S (2019) How optimal distinctiveness affects new ventures' failure risk: A contingency perspective. *J. Bus. Venturing* 34(3):477–495.

Granqvist N, Grodal S, Woolley JL (2013) Hedging your bets: Explaining executives' market labeling strategies in nanotechnology. *Organ. Sci.* 24(2):395–413.

Haans RF (2019) What's the value of being different when everyone is? The effect of distinctiveness on performance in homogeneous vs. heterogeneous categories. *Strategic Management J.* 40(1):3–27.

Haans RF, Pieters C, He ZL (2016) Thinking about U: Theorizing and testing U- and inverted U-shaped relationships in strategy research. *Strategic Management J.* 37(7):1177–1195.

Hampton JA (2007) Typicality, graded membership, and vagueness. *Cogn. Sci.* 31(3):355–384.

Hannan MT, Le Mens G, Hsu G, Kovács B, Negro G, Pólos L, Pontikes EG, Sharkey AJ (2019) *Concepts and Categories: Foundations for Sociological and Cultural Analysis* (Columbia University Press, New York).

Hsu G (2006) Jacks of all trades and masters of none: Audiences' reactions to spanning genres in feature film production. *Admin. Sci. Quart.* 51(3):420–450.

Hsu G, Grodal S (2015) Category taken-for-grantedness as a strategic opportunity: The case of light cigarettes, 1964 to 1993. *Amer. Sociol. Rev.* 80(1):28–62.

Hsu G, Hannan MT, Koçak Ö (2009) Multiple category memberships in markets: An integrative theory and two empirical tests. *Amer. Sociol. Rev.* 74(1):150–169.

Hunt JM, Kernan JB, Bonfield EH (1992) Memory structure in the processing of advertising messages: How is unusual information represented? *J. Psych.* 126(4):343–356.

Johnson MD, Fornell C (1987) The nature and methodological implications of the cognitive representation of products. *J. Consumer Res.* 14(2):214–228.

Jones C, Maoret M, Massa FG, Svejenova S (2012) Rebels with a cause: Formation, contestation, and expansion of the de novo category "modern architecture," 1870–1975. *Organ. Sci.* 23(6):1523–1545.

Jonides J, Lewis RL, Nee DE, Lustig CA, Berman MG, Moore KS (2008) The mind and brain of short-term memory. *Annual Rev. Psych.* 59:193–224.

Kennedy MT (2008) Getting counted: Markets, media, and reality. *Amer. Sociol. Rev.* 73(2):270–295.

Kennedy MT, Fiss PC (2013) An ontological turn in categories research: From standards of legitimacy to evidence of actuality. *J. Management Stud.* 50(6):1138–1154.

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

Khaire M (2017) The importance of being independent: The role of intermediaries in creating market categories. Durand R, Granqvist N, Tyllström A, eds. *Research in the Sociology of Organizations*, vol. 51 (Emerald, Bingley, UK), 259–293.

Khessina OM, Reis S (2016) The limits of reflected glory: The beneficial and harmful effects of product name similarity in the U.S. network TV program industry, 1944–2003. *Organ. Sci.* 27(2): 411–427.

Koçak Ö, Hannan MT, Hsu G (2014) Emergence of market orders: Audience interaction and vanguard influence. *Organ. Stud.* 35(5):765–790.

Kovács B, Sharkey AJ (2015) The paradox of publicity: How awards can negatively affect the evaluation of quality. *Admin. Sci. Quart.* 59(1):1–33.

Kovács B, Carnabuci G, Wezel FC (2021) Categories, attention, and the impact of inventions. *Strategic Management J.* 42(5):992–1023.

Kuilman JG, Van Driel H (2013) You too, Brutus? Category demise in Rotterdam warehousing, 1871–2011. *Indust. Corporate Change* 22(2):511–548.

Kurzban R, Duckworth A, Kable JW, Myers J (2013) An opportunity cost model of subjective effort and task performance. *Behav. Brain Sci.* 36(6):661–679.

Lieder F, Griffiths TL (2020) Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources. *Behav. Brain Sci.* 43:E1.

Lind JT, Mehlum H (2010) With or without U? The appropriate test for a U-shaped relationship. *Oxf. Bull. Econom. Stat.* 72(1): 109–118.

Lo JY-C, Fiss PC, Rhee EY, Kennedy MT (2020) Category viability: Balanced levels of coherence and distinctiveness. *Acad. Management Rev.* 45(1):85–108.

Lounsbury M, Rao H (2004) Sources of durability and change in market classifications: A study of the reconstitution of product categories in the American mutual fund industry, 1944–1985. *Soc. Forces* 82(3):969–999.

McFadden D (1974) Conditional logit analysis of qualitative choice behavior. Zarembka P, ed. *Frontiers in Econometrics* (Academic Press, New York), 105–142.

Medin DL, Schaffer MM (1978) Context theory of classification learning. *Psych. Rev.* 85(3):207–238.

Montauti M (2019) Anchors in rough seas: Understanding category spanning as a source of market coordination. *J. Management Stud.* 56(4):823–853.

Montauti M, Wezel FC (2016) Charting the territory: Recombination as a source of uncertainty for potential entrants. *Organ. Sci.* 27(4):954–971.

Navis C, Glynn MA (2010) How new market categories emerge: Temporal dynamics of legitimacy, identity, and entrepreneurship in satellite radio, 1990–2005. *Admin. Sci. Quart.* 55(3): 439–471.

Negro G, Hannan MT, Fassiotto M (2015) Category signaling and reputation. *Organ. Sci.* 26(2):584–600.

Negro G, Hannan MT, Rao H (2010) Categorical contrast and audience appeal: Niche width and critical success in winemaking. *Indust. Corporate Change* 19(5):1397–1425.

Negro G, Hannan MT, Rao H (2011) Category reinterpretation and defection: Modernism and tradition in Italian winemaking. *Organ. Sci.* 22(6):1449–1463.

Oberauer K, Kliegl R (2006) A formal model of capacity limits in working memory. *J. Mem. Lang.* 55(4):601–606.

Ozcan P, Gurses K (2018) Playing cat and mouse: Contests over regulatory categorization of dietary supplements in the United States. *Acad. Management J.* 61(5):1789–1820.

Piazzai M, Wijnberg NM (2019) Product proliferation, complexity, and deterrence to imitation in differentiated-product oligopolies. *Strategic Management J.* 40(6):945–958.

Pontikes EG (2012) Two sides of the same coin: How ambiguous classification affects multiple audiences' evaluations. *Admin. Sci. Quart.* 57(1):81–118.

Pontikes EG (2018) Category strategy for firm advantage. *Strategy Sci.* 3(4):620–631.

Pontikes EG, Barnett WP (2015) The persistence of lenient market categories. *Organ. Sci.* 26(5):1415–1431.

Pontikes EG, Kim R (2017) Strategic categorization. Durand R, Granqvist N, Tyllström A, eds. *Research in the Sociology of Organizations*, vol. 51 (Emerald, Bingley, UK), 71–111.

Porter A, Bogdanov D, Kaye R, Tsukanov R, Serra X (2015) AcousticBrainz: A community platform for gathering music information obtained from audio. Müller M, Wiering F, eds. *Proc. 16th Internat. Soc. Music Inform. Retrieval Conf.* (ISMIR, Canada), 786–792.

Pothos EM, Chater N (2002) A simplicity principle in unsupervised human categorization. *Cogn. Sci.* 26(3):303–343.

Rao H, Monin P, Durand R (2003) Institutional change in Toque Ville: Nouvelle cuisine as an identity movement in French gastronomy. *Amer. J. Sociol.* 108(4):795–843.

Rhee EY, Lo JY, Kennedy MY, Fiss PC (2017) Things that last? Category creation, imprinting, and durability. Durand R, Granqvist N, Tyllström A, eds. *Research in the Sociology of Organizations*, vol. 51 (Emerald, Bingley, UK), 295–395.

Rosa JA, Porac JF, Runser-Spanjol J, Saxon MS (1999) Sociocognitive dynamics in a product market. *J. Marketing* 63(4):64–77.

Rosch EH (1978) Principles of categorization. Rosch EH, Lloyd BB, eds. *Cognition and Categorization* (Lawrence Erlbaum, Hillsdale, NJ), 27–48.

Rosch EH, Mervis CB, Gray WD, Johnson DM, Boyes-Braem P (1976) Basic objects in natural categories. *Cognit. Psych.* 8(3): 382–439.

Ruef M, Patterson K (2009) Credit and classification: The impact of industry boundaries in nineteenth-century America. *Admin. Sci. Quart.* 54(3):486–520.

Sgourev SV, Aadland E, Formilan G (2023) Relations in aesthetic space: How color enables market positioning. *Admin. Sci. Quart.* 68(1):146–185.

Shepard RG, Hovland CI, Jenkins HM (1961) Learning and memorization of classifications. *Psych. Monogr. Gen. Appl.* 75(13):1–42.

Simon HA (1955) A behavioral model of rational choice. *Quart. J. Econom.* 69(1):99–118.

Slavich B, Svejenova S, Opazo MP, Patriotta G (2020) Politics of meaning in categorizing innovation: How chefs advanced molecular gastronomy by resisting the label. *Organ. Stud.* 41(2): 267–290.

Smith EB, Chae H (2016) "We do what we must, and call it by the best names": Can deliberate names offset the consequences of organizational atypicality? *Strategic Management J.* 37(6): 1021–1033.

Smith EB, Chae H (2017) The effect of organizational atypicality on reference group selection and performance evaluation. *Organ. Sci.* 28(6):1134–1149.

Smith JD, Minda JP (2002) Distinguishing prototype-based and exemplar-based processes in dot-pattern category learning. *J. Exp. Psych. Learn. Mem. Cogn.* 28(4):800–811.

Soublière J-F, Lo JY, Rhee EY (2023) Coherence within and across categories: The dynamic viability of product categories on Kickstarter. *Acad. Management J.* Forthcoming.

Taeuscher K, Rothe H (2021) Optimal distinctiveness in platform markets: Leveraging complementors as legitimacy buffers. *Strategic Management J.* 42(2):435–461.

Taeuscher K, Zhao EY, Lounsbury M (2022) Categories and narratives as sources of distinctiveness: Cultural entrepreneurship within and across categories. *Strategic Management J.* 43(10): 2101–2134.

Downloaded from informs.org by [46.64.248.158] on 22 February 2024, at 08:10 . For personal use only, all rights reserved.

Tversky A, Kahneman D (1973) Availability: A heuristic for judging frequency and probability. *Cognit. Psych.* 5(2):207–232.

Unkelbach C (2006) The learned interpretation of cognitive fluency. *Psych. Sci.* 17(4):339–345.

Van Venrooij A (2015) A community ecology of genres: Explaining the emergence of new genres in the U.K. field of electronic/dance music, 1985–1999. *Poetics* 52:104–123.

Vergne JP, Wry T (2014) Categorizing categorization research: Review, integration, and future directions. *J. Management Stud.* 51(1):56–94.

Verhaal JC, Pontikes EG (2022) Shaping cultural meanings in markets with category strategy and optimal distinctiveness: An agency-based perspective. Lockwood C, Soublière J-F, eds. *Research in the Sociology of Organizations*, vol. 80 (Emerald, Bingley, UK), 179–191.

Verhaal JC, Khessina OM, Dobrev SD (2015) Oppositional product names, organizational identities, and product appeal. *Organ. Sci.* 26(5):1466–1484.

Verheyen S, Ameel E, Storms G (2007) Determining the dimensionality in spatial representations of semantic concepts. *Behav. Res. Methods* 39(3):427–438.

Wry T, Lounsbury M (2013) Contextualizing the categorical imperative: Category linkages, technology focus, and resource acquisition in nanotechnology entrepreneurship. *J. Bus. Venturing* 28(1):117–133.

Younkin P, Kashkooli K (2020) Stay true to your roots? Category distance, hierarchy, and the performance of new entrants in the music industry. *Organ. Sci.* 31(3):604–627.

Zanella P, Cillo P, Verona G (2021) Whatever you want, whatever you like: How incumbents respond to changes in market information regimes. *Strategic Management J.* 43(7):1258–1286.

Zhao EY, Ishihara M, Lounsbury M (2013) Overcoming the illegitimacy discount: Cultural entrepreneurship in the US feature film industry. *Organ. Stud.* 34(12):1747–1776.

Zhao EY, Fisher G, Lounsbury M, Miller D (2017) Optimal distinctiveness: Broadening the interface between institutional theory and strategic management. *Strategic Management J.* 38(1):93–113.

Zhao EY, Ishihara M, Jennings PD, Lounsbury M (2018) Optimal distinctiveness in the console video game industry: An exemplar-based model of proto-category evolution. *Organ. Sci.* 29(4): 588–611.

Zipf GK (1949) *Human Behavior and the Principle of Least Effort* (Addison-Wesley, Cambridge, MA).

Zuckerman EW (1999) The categorical imperative: Securities analysts and the illegitimacy discount. *Amer. J. Sociol.* 104(5):1398–1438.

Zuckerman EW (2016) Optimal distinctiveness revisited. Pratt MG, Schultz M, Ashforth BE, Ravasi D, eds. *The Oxford Handbook of Organizational Identity* (Oxford University Press, New York), 183–199.

Zuckerman EW (2017) The categorical imperative revisited: Implications of categorization as a theoretical tool. Durand R, Granqvist N, Tyllström A, eds. *Research in the Sociology of Organizations*, vol. 51 (Emerald, Bingley, UK), 31–68.

Zunino D, Suarez FF, Grodal S (2019) Familiarity, creativity, and the adoption of category labels in technology industries. *Organ. Sci.* 30(1):169–190.

Michele Piazzai is Assistant Professor of Management and Ramón y Cajal fellow at Universidad Carlos III de Madrid. His current research interests include formal models of categorization and decision making, with particular regard to their implications for the behavior of firms and their audiences. He received his PhD from Delft University of Technology (The Netherlands).

Min Liu is Associate Professor in Management at Durham University Business School. Her main research interests lie in organizational ecology, categorization in market contexts, as well as industry and market evolution. She received her PhD from Università della Svizzera italiana (Lugano, Switzerland).

Martina Montauti is Assistant Professor of Strategic and International Management at the University of St. Gallen. She received her PhD in Management from Università della Svizzera italiana (Lugano, Switzerland). Her research focuses on the influence of social cognition on decision making and, ultimately, on outcomes of strategic relevance such as market competition, entry, and performance returns.