American Journal of Cultural Sociology (2024) 12:575–619
https://doi.org/10.1057/s41290-024-00210-2

ORIGINAL ARTICLE

Making Obama iconic: iconic potential and the affordances of images

Terence E. McDonnell[1] · **Marshall A. Taylor**[2] · **Michael Lee Wood**[3]

Accepted: 29 December 2023 / Published online: 20 April 2024

Abstract

"Iconic" images possess cultural power. This paper asks: What makes an image "iconic"? To become publicly recognized as "iconic," images undergo a process of consecration. We argue that this process of consecration is facilitated by the capacity to perceive the degree to which an image "looks iconic." From this perspective, iconicity is an affordance that emerges from the relation between an image's visual qualities and the cognitive associations of a perceiving subject. The paper presents a novel method for measuring subjective perceptions of iconicity using pairwise image comparison tasks. The analysis finds that people systematically evaluate some images as more iconic than others, suggesting that images vary in their "iconic potential." The paper assesses whether image qualities are associated with an image's iconic potential and whether iconic potential varies across groups. There is minimal variation in people's evaluations of iconicity, suggesting a high degree of stability. These results indicate that the process by which images become publicly recognized as "iconic" is not entirely arbitrary. Some images are perceived by individuals as more iconic than others, which may influence which images get promoted and circulated.

Keywords Iconicity · Cultural power · Affordances · Iconic potential · Visual sociology

✉ Terence E. McDonnell
tmcdonn2@nd.edu

[1] Department of Sociology, University of Notre Dame, 4060 Jenkins Nanovic Hall, Notre Dame, IN 46556, USA

[2] Department of Sociology, New Mexico State University, Las Cruces, New Mexico, USA

[3] Department of Sociology, Brigham Young University, Provo, Utah, USA

Introduction

Images are an essential part of public culture and collective memory. Like other cultural objects, images can possess cultural power, defined as the capacity to shape belief and behavior (McDonnell 2016, 2023; Schudson 1989), remain in the public consciousness, and enter the canon (Griswold 1987). Through images, the visages of individuals can persist and play active roles in public long after their deaths. Those of Che Guevara and Marilyn Monroe do cultural work symbolizing and constraining notions of revolution and beauty. Some images, such as Dorothea Lange's "Migrant Woman" and *National Geographic's* "Afghan Girl," endure in the public's imagination.

"Iconic" images possess cultural power. Various scholars have theorized what cultural work already-iconic images do (Alexander 2008; Zubrzycki 2016). Icons have been associated with critical social functions, such as creating shared meanings and a sense of "shared-ness" (Sturken 2015, p. 3) and becoming sites of contestation and "aesthetic revolt" in the building of national identity (Zubrzycki 2016). Similarly, scholars have argued that iconic images, or images with "iconic power" (Alexander et al. 2012), can communicate the meaning of a subject, often instantaneously, and "symbolically condense" stories, narratives, and imaginaries (Sonnevend 2020). Iconic images are said to communicate not only symbolic meaning but produce affective and sensual effects as well (Alexander et al. 2012, p. 4; Bartmanski 2015)—a necessary capacity to enact "cultural work" successfully (Schudson 1989).

Previous research primarily focuses on the dynamics of already-iconic images (Alexander 2008; Alexander et al. 2012), but in this paper, we ask the causally prior question: What makes an image "iconic"? There are innumerable images in the world, but only specific images possess the capacity to effect social change, capture a moment, and distill meaning to an essence. Furthermore, images vary in their degree of iconicity (Alexander 2012, p. 27). How do images come to possess this extraordinary cultural power? Can any image become iconic, or only images with certain visual qualities? Following Sonnevend's (2020) brief mention of the concept, we call this capacity of images to become iconic "iconic potential."

Any investigation of what makes images iconic requires a conceptualization of "iconicity." Most generally, we conceptualize iconicity as a perceived quality of images.[1] This perceived quality has two types. The first type is the recognition of an image's privileged public status. Images with this quality undergo a consecration process similar to Durkheim's description of sacralization (Alexander 2012, pp. 34–35). As Alexander observes, in contemporary societies there is a "series of mediations" between a cultural object and the audience, shaping the audience's perceptions. This includes processes of production, distribution, and public evaluation (e.g., art criticism). These mediations may heighten or diminish an image's iconic status (Salganik et al. 2006). Most research on iconicity focuses on this quality of public status by selecting and analyzing images with this quality.

[1] This conceptualization is consistent with contemporary research on iconicity focusing on historically contingent cultural meanings, but it is distinct from Peircean approaches that define icons in formal terms as signs that share properties with the objects they represent (Peirce 1955).

Fig. 1 From iconic potential to public recognition as iconic

The second type of perceived quality of iconicity is the recognition that a given image "looks iconic." This is a subtle but essential distinction from the first form because it implies a capacity to intuit the iconicity of images, even *before* the images undergo any consecration process. It is the difference between, say, recognizing a famous picture of Marilyn Monroe as iconic and looking at one's personal photo album and feeling, "This is an iconic photograph of Mom." This perceived quality is more substantive insofar as it requires that images possess specific attributes (rather than being iconic because they have been publicly consecrated as such).

We argue that before images achieve the public quality of iconicity (the first type discussed above), they are first perceived and chosen by producers and distributors among a set of possible images. This selection process is not arbitrary or random. Not all images are equal; some have more "iconic potential" than others. We hypothesize that this selection process is facilitated by an encultured capacity to subjectively perceive and evaluate the relative iconicity of images, *even when seeing them for the first time* (See Fig. 1).

Our focus on iconic potential and perceived qualities of iconicity differs from previous research. Some research has tried to determine what makes images iconic by looking for commonalities across recognized iconic images. However, these efforts more often frustrate than illuminate (Kemp 2011).[2] We argue that adequately answering what makes images iconic requires, first, differentiating iconicity as a publicly ascribed quality and a substantive, subjectively perceived quality and, second, disaggregating public processes of consecration and distribution from individual-level perceptions, as shown in Fig. 1. If achieving the publicly recognized status as iconic involves individual and collective processes, then different measures are needed for different processes. There is too much noise among images widely recognized as iconic to make inferences about individual-level perceptions. What is initially perceived as having iconic potential may vary across persons and domains. Any set of publicly consecrated iconic images may, thus, vary widely and have no necessary relations. A different and more direct approach is needed.

This article takes a first step in understanding iconicity as a result of both individual perception and collective consecration by developing a way of measuring and analyzing "iconic potential." Our goals are threefold: (1) To determine whether people actually possess an encultured capacity to "see" images as more or less iconic, (2) to identify image-level features associated with iconic potential, and (3)

[2] Art historian Kemp said as much of his attempt to identify universal qualities of iconic images. As he notes, "Do iconic images have anything in common?...we can refine the premiss to read: iconic images necessarily share a certain set of key or essential features if they are to achieve the highest status…I think this premiss *[sic]*…is untenable" (339–340).

to determine the degree to which these evaluations vary across persons and groups. The method we develop involves a direct comparison of the relative iconicity of images that are not widely circulated. We choose lesser-known images to avoid social influence effects (Salganik et al. 2006) because popular images could color people's perceptions by evoking the public status meaning of iconicity rather than the substantive "iconic-like" perception we aim to capture. We also choose images that represent the same image subject and are created by the same person to avoid noise introduced when using images across domains or by different creators. We then collect Americans' assessments of iconicity around a random sample of less circulated images of President Obama, all taken by White House Photographer Pete Souza.

The paper proceeds as follows. In the next section, we develop the concept of iconic potential, discuss relevant research on iconicity and image qualities, and introduce our novel approach to measuring and analyzing iconicity. We then discuss our data and measures. In the paper's central section, we present our analysis and interpret our findings. We find strong evidence that images vary in their ability to afford iconicity at the level of individual judgment. We discuss visual qualities that influence evaluations of iconicity in our sample. Additionally, we find that iconicity as an affordance is quite stable across groups of people, meaning that despite some variation, most people's evaluations of iconicity appear to be based on similar visual qualities. We conclude by discussing the studies' limitations and the implications for the study of iconicity in cultural sociology.

Theoretical background

Iconicity and iconic potential

Scholars have argued that iconic images have qualities that afford emotional encounters through the senses. Alexander (2008) describes interacting with an icon as an "experience" that "sucks the feeling viewer into meaning" (p. 783). He describes this as understanding "by feeling, by contact, by the 'evidence of the senses' rather than the mind" (p. 782). Sonnevend (2020) describes the "image encounter" as "a diffuse sensory experience… our regular meeting with images, in which we come in contact with the image and negotiate our relationship with it" (p. 220). She observes that sometimes, images rise to the level of iconic "when we strongly feel toward an image and develop a deep emotional bond." Roland Barthes (1981) was similarly attentive to the qualities of images and their capacity to elicit different effects by distinguishing between the "studium" and the "punctum." The studium describes the more mundane interest that a photograph evokes—an image that draws attention and observation but not necessarily emotion or deep connection. Alternatively, the punctum is the quality of an image that "pierces" or "wounds" the viewer. It is often a surprising and arresting quality of an image—an object that triggers a memory or the aesthetic feel of the image that attracts people to the object. Notably, the punctum could be different for different people. This is to say the qualities of images that

lead one person to see it as iconic may vary from person to person, even as people have a shared implicit understanding of iconicity.

If images vary in their ability to afford iconicity, what sets apart "iconic" images? Scholars have primarily approached this question by focusing on the semiotic functions that icons perform. For example, Peirce (1955) defines icons as signs that "represent its object mainly by its similarity." Smith (2012, p. 172) echoes and extends Peirce's definition, suggesting that iconic images both *resemble* their intended subject and *represent* cultural narratives and myths. Similarly, Alexander views icons as "symbolic condensations" rooting "social meanings in a specific and 'material' form" (2008, p. 782)—the aesthetic surface of an image collapses an underlying depth of meaning that comes from the icon's social embeddedness. Kemp (2011) says that iconic images "are endowed with a special presence, as if some quality of the original [subject] is embedded in them" (p. 342), much like Benjamin's (1968) "aura."

We define iconic potential as an image-level relational property constituted by people's subjective perceptions of the image. Images perceived as more iconic by more individuals have more "iconic potential," meaning more potential to be selected, consecrated, and recognized as a public icon. The question of iconic potential—whether any image can become iconic or only images with certain visual qualities—is a specific instance of a more general question about the dynamics of materiality and perception in meaning-making. Extending research on materiality to questions of iconicity, we might ask if iconic images are "passive surface[s] upon which social forces act and impart meaning" or are they "active and constitutive elements in the production of social forms, relations, and meanings?" (Domínguez Rubio 2014, p, 618).

If images are passive, then iconicity may best be explained via a process of applying the label of "iconic" to an image (Becker 2008). From this perspective, the qualities of the image are less relevant, and any image could rise to the level of an icon because "all hits are flukes" (Bielby and Bielby 1994). The most robust version of this argument is that any image could achieve iconic status—from a *Time Magazine* photojournalist's portrait to the blurry snapshot taken on your uncle's flip phone—if it goes through the requisite socio-historical process that ritually consecrates it as iconic. At the level of individual judgment, the capacity to judge the "iconicity" of an image would be understood as the ability to *recognize* the socially imposed iconic status (McDonnell et al. 2017). For example, a person could identify an image of Che Guevara or Marilyn Monroe as "iconic" only insofar as they have a sense of its iconic status in the public sphere. Such an account is bolstered by work explaining what makes a cultural object rise in status, from mechanisms of recognition to experimental tests of self-fulfilling prophecies and social influence (Salganik et al. 2006; Salganik and Watts 2008).

Alternatively, if images are active, then adequately explaining iconicity requires attending to how the qualities of images interact with people's perceptual and cultural capacities. From this perspective, iconicity is not a quality that can be ascribed via a label but a quality that specific images *afford*. Guided by affordance theory (Davis 2020; Gibson 1979; Norman 1988), this hypothesis treats iconicity as an *affordance*—the potential uses and meanings of objects made manifest in

interactions with people. Here, we draw on recent work in the affordance tradition that blends work on material and symbolic affordances, linking perception and semiotics (McDonnell 2016, 2023; DeNora 2000; Childress et al. 2021). Affordances are inherently *relational*, constituted by the qualities of objects and the experience and capacities of the people interacting with them. From this perspective, only certain images can gain iconic status—images with the right qualities to be read as iconic by audiences. From this perspective, then, iconic potential emerges at the intersection of image qualities and the meanings and uses they afford to enculturated and attuned people. If iconicity is an affordance, then images with certain qualities (e.g., material, formal, conventionally symbolic qualities) should be more likely to afford iconicity, in the same way that only objects with certain qualities evoke the perception of beauty.[3] This is not to say that there is an objective, universal set of qualities that afford beauty; instead, given people's perceptual and cultural capacities, only objects with certain qualities will evoke the perception.

An affordance-based explanation of iconicity complements labeling-based explanations. Images cannot achieve a culturally powerful iconic status without the individuals responsible for circulating them and boosting their public status. In this light, affordance theory proposes that consequential individual judgments that lead to images gaining cultural power are not arbitrary or random but partially anchored by the fact that images differ in their capacity to afford subjective perceptions of iconicity. Said differently, an affordance-based explanation of iconicity posits that individuals can perceive the "iconicity" of images, regardless of whether these images have been widely circulated or consecrated as "iconic" in the public sphere. From this perspective, an image has more or less "iconic potential" (i.e., the potential to become a culturally powerful image) depending on its ability to afford "iconicity" at the level of individual judgment. More than just a matter of popularity and consecration, individuals judge images as iconic and images achieve iconicity based on "interdependent ensembles of their properties" experienced as a gestalt (Seguin 2023; Taylor et al. 2019).

There is a danger in reading affordances as determinative. Affordances are relative; the affordances of images depend on their visual qualities *and* the cultural experiences and capacities of the individuals interacting with them (McDonnell 2023). Therefore, variation in either image qualities or individual capacities or experiences may change the available affordances—in this case, how much an image evokes iconicity for any given person. One question this raises for cultural analysis is how stable an affordance is across people. If iconicity is stable across people, we would expect that evaluations of iconicity would vary more by image qualities and less by individual or group differences.

[3] Semiotically, we draw from Peirce's distinction between icon, index, and symbol. These photographs must first and foremost be an "iconic" sign in that they physically depict Obama. Still, the *iconic potential* afforded by the photograph is also possibly shaped by the presence of indexical and conventionally symbolic signs (e.g., a handshake and smile indexes friendliness, or a statue of a bald eagle conventionally symbolizes America.) that might influence its iconicity. Material arrangements of symbols, or these symbols or qualities of the image itself (e.g., blurriness) may enhance or detract from the iconicity as well.

Addressing limitations of previous research on iconicity

Previous research on iconicity has proceeded by selecting cases of images that have achieved iconic status—images that have undergone some form of social consecration. While often revealing, this approach is limited for studying people's perceptions of iconic potential. One limitation is that any public canon of iconic images will result from many individual evaluations of iconic potential. We have the outcome (iconic image), but scholars tend to attribute the reasons for its iconic status to idiosyncratic historical explanations. This approach blackboxes how qualities of images might afford iconic potential. Without systematic data, analysts are forced to rely on their own interpretation of what image qualities make it iconic, which may not map onto how individuals in the world actually assess whether images are iconic.

Attention to perception at the individual level reveals patterns in the act of meaning-making. Because individuals do not necessarily evaluate the iconic potential of images in the same way, a public canon of images may result from many different substantive judgments of iconic potential. Any attempt to identify what canonical iconic images have in common in such cases may fail if different images achieve iconic status via different communities with distinct ways of evaluating images. This could explain (in part) why Kemp's (2011) attempt to identify key features of iconic images ended in frustration. Stated differently, analyzing images that have achieved iconic status tells us relatively little about iconic potential because they are far removed from initial, individual-level acts of evaluation.

Suppose the analyst aims to understand why certain images gain public recognition as iconic. In that case, analyzing already canonized images is limiting because one has effectively sampled on the dependent variable, risking biased observations. Ignoring the "losers" in an iconic competition may lead the researcher to miss essential factors influencing evaluations of iconicity. Additionally, without comparison to "not iconic" images, we lack the evidence to reject the possibility that *any* image, independent of visual qualities, could achieve iconic status.

In this paper, we advance previous research with a novel methodological approach with three key features: First, we elicit individual judgments of iconicity amid a set of images with relatively controlled producer and subject status signals to understand how people evaluate the iconicity of image qualities. Second, we elicit these judgments via pairwise comparison tasks using a set of relatively obscure images of a single subject (former U.S. President Obama). Third, we seek to maximize differences in iconic taste by creating three different subgroup analyses (micro taste communities, demographic groups, and groups who differ by how they define the core function of the presidency) as a conservative test of the stability of an image's iconic potential across groups. This unique approach is well-suited to identify (1) whether images do vary in their ability to afford iconicity, (2) which visual qualities are salient for evaluating iconicity by examining both winners and losers in iconic contests, and (3) whether these salient visual qualities are stable or vary across groups.

Data and measures

Our study has three main research aims. First, we seek to establish whether people's evaluation of "iconic potential" is grounded in image qualities—or, said differently, whether certain images "afford" evaluations of iconicity without previously being labeled by cultural gatekeepers as "iconic." Iconic potential is measured in the aggregate—the more people see an image as iconic, the more likely they will experience it as iconic and be "nominated" for a public consecration process. Assuming people possess this capacity to identify images as more or less iconic, our second goal is to assess which image qualities are salient for evaluating iconic potential and, if so, whether and the degree to which this varies across persons. Said differently, if iconic potential is an affordance, what visual qualities afford it, and how stable is it?

This section discusses the data we collected and measures we constructed to address these three research questions. We begin by discussing our image-level "relative iconicity" measure, our primary outcome variable. Next, we discuss our measures of image qualities used to identify which image qualities matter for evaluating iconic potential and then discuss how we subset the data to measure the stability of this affordance across people.

Measuring relative iconicity

Studying iconic potential requires developing adequate measures of people's subjective perceptions. As an individual judgment, iconicity can be measured in two main ways: As an absolute concept (i.e., "iconic" or "not iconic") or a relative concept ("more iconic" or "less iconic"). We chose to measure iconicity in terms of relative judgment for two reasons. First, relative iconicity more closely aligns with the concept of iconic potential. Consider: when presented with a pair of images, a person may perceive neither as iconic. However, that does not mean the images have the same *iconic potential*. One image might have all the material qualities of iconic images but fails to evoke the "iconic" categorization because it has not undergone a public consecration ritual. Measuring iconicity in terms of relative judgment can better determine whether iconic potential of image qualities exists. The relative iconicity of images will either be random (because the judgment of iconicity is wholly arbitrary and dependent on labels, such that there is no "grounding"), or it will vary systematically if it is indeed based on the presence or absence of certain image qualities.

Relative iconicity is a relational measure derived from individual judgments, meaning images are more or less iconic based on aggregations of individual decisions. Our measurement of relative iconicity was derived in two steps: collecting individual-level judgments using pairwise image comparison tasks and aggregating the results of these comparisons to create an image-level measure. We discuss each step below.

Measuring relative iconicity: pairwise image comparison survey

To construct our measure of relative iconicity, we collected data using a novel survey of U.S. citizens ages 18 and older recruited using Amazon Mechanical Turk.[4] The data were collected from October 23–25, 2020. After eliminating duplicate responses and non-US respondents and removing respondents who either skipped or incorrectly answered a "top-three" ranking question (which we discuss in more detail in the micro-community analysis outlined later), we had 919 cases.

Participants were shown a series of pairs of images and instructed to choose each pair's most "iconic" image. This productive task aimed to inductively draw out people's understandings based on their pairwise evaluations. Had we given a definition of iconic, it may have clashed with people's perceptions and offered a biased test of whether iconicity is a stable perception in the world. Because we wanted to draw out people's understandings of "iconicity," we did not define "iconicity" for the participants.

The images used for the pairwise comparisons consist of a random sample of 20 images (without replacement) of President Obama taken by Pete Souza, the Presidential photographer during the Obama presidency (see Figs. 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21).[5] Several factors guided our image selection. We chose Obama as our subject because we wanted to focus on a single person who was almost universally known, thereby controlling for variation due to familiarity with or status of the subject. Additionally, as a political figure, Obama provides a straightforward way to control for potentially influential individual qualities (i.e., measuring people's voting patterns and political identity/beliefs). Using Obama as our subject also allowed us to limit the sample to images from the same photographer, offering some control over the variation in the aesthetic sensibilities across images.

Given that this is the first study we know of to measure individual judgments of relative iconicity, we sampled images randomly because our approach is necessarily inductive. In this case, a random sample of images is more appropriate than a more focused sampling based on specific image qualities because it opens more possibilities for inductive hypothesis generation. We sampled 20 images to allow for variability in image qualities while limiting the demands on research participants and ensuring enough quality observations. After sampling 20 random images, we performed a reverse Google image search for each of them to ensure that none of the images had been widely circulated, minimizing the potential of the social status of

[4] Amazon Mechanical Turk is a popular online platform for crowd-sourcing tasks, including survey-taking. The platform is a popular respondent pool for survey research—and, although not as representative as probability samples (Goodman and Paolacci 2017), MTurk samples outperform college-based samples in a variety of ways. For a review of how MTurk samples compare to nationally-representative probability samples (namely, from the Current Population Survey and the American National Elections Study 2008 panel), see McDonnell, Stoltz, and Taylor (2022, pp. 892-893).

[5] https://www.flickr.com/photos/obamawhitehouse/ Sampled from 6668 images managed by the National Archives and Records Administration (NARA).

Fig. 2 1st ranked image. *Note*: Photography by Pete Souza, The White House

images influencing people's judgments.[6] The images show Obama in various contexts (e.g., alone, with a crowd) and positions (e.g., sitting, standing).

The 20 images yielded 190 possible image pairs, making it impossible to create a direct individual-level image ranking with pairwise comparison tasks because we could not reasonably expect each participant to make 190 evaluations. However, we created aggregate measures of relative iconicity by randomly assigning a limited number of pairwise tasks to participants and ensuring that each of the 190 possible pairs was seen many times (Salganik and Levy 2015; Salganik and Watts 2008). To this end, we gave

[6] We performed a reverse Google image search of the 20 images and found that none were widely circulated, measured in the number of results relative to popular images of Obama. The most circulated image of the set had 1/10 the number of results as the iconic Obama "HOPE" image.

Fig. 3 2nd ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 4 3rd ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 5 4th ranked image. *Note*: Photography by Pete Souza, The White House

each participant 25 pairwise image comparison tasks. Of these 25 tasks, the first 10—comprising a set of 10 pairs that included all 20 images—were identical for everyone, ensuring participants saw each image exactly once before seeing repeated images. Participants then completed an additional 15 pairwise image comparison tasks randomly selected from the remaining 180 possible pairs. These pairs were randomly sampled evenly, meaning that each of these 180 pairs was seen by 158 or 159 different participants. To avoid bias based on image position, we randomized the top/bottom placement of images every time a pair was shown. In total, we have data on 20,250 pairwise image comparison tasks.

Fig. 6 5th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 7 6th ranked image. *Note*: Photography by Pete Souza, The White House

Measuring iconic potential: aggregating results of the pairwise image comparison tasks

We used the results of the 20,250 pairwise image comparison tasks to create the image-level measure of aggregated relative iconicity (iconic potential). The process involved

Fig. 8 7th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 9 8th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 10 9th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 11 10th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 12 11th ranked image. *Note*: Photography by Pete Souza, The White House

creating a "win percentage" for each image: the percentage of times an image was evaluated as more iconic than the images it was paired with. Formally speaking, let N be the set of all pairwise image comparison tasks, N_i be the subset of N that includes image i (i.e., all comparison tasks that feature image i), and $|N_i|$ be that subset's cardinality. Also, let x be a binary outcome variable from those comparisons between any two images i and k, where $i, k \in S$ and S is the set of all 20 images. Then, with indicator function $\mathbf{1}_{i>k}(x)$ being a 1 when i is selected over k and a 0 otherwise, the relative iconicity r for image i is defined as:

$$r_i = \frac{\sum \mathbf{1}_{i>k}(x)}{|N_i|} \times 100 \tag{1}$$

Fig. 13 12th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 14 13th ranked image. *Note*: Photography by Pete Souza, The White House

After calculating the win percentage (r_i) for each image, we ranked the images from 1 (highest win percentage) to 20 (lowest win percentage). When we speak of an image being "the most iconic" or "more iconic," we mean that the image had a higher win percentage or ranking than other images.

Coding for image qualities

To assess which image qualities mattered, we quantitatively coded each of our sample of 20 images for their qualities. Our selection of quantitative image quality variables was loosely guided by previous research, but because our study is

Fig. 15 14th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 16 15th ranked image. *Note*: Photography by Pete Souza, The White House

inductive, we included a wide range of plausible image-level variables. We took two approaches to constructing image quality variables. First, we used the "imagefluency" R package (Mayer 2020) to measure low-level image features, such as contrast and symmetry. We then measured high-level image qualities by hand-coding the Obama images based on substantive features (e.g., the presence or absence of presidential symbols or being outdoors). Most of these qualities were measured dichotomously as the feature's presence (1) or absence (0). For a list of the qualities we measured, see Table 1.

Fig. 17 16th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 18 17th ranked image. *Note*: Photography by Pete Souza, The White House

Using subsets to assess varying salience of qualities and the stability of iconic potential

To assess the role of audience, we sought to examine whether different groups valued different qualities and whether the distribution of iconic potential changed from group to group. Iconic potential can be measured across different groups, depending on how individual judgments of relative iconicity are aggregated. Our full-sample measure of iconic potential can reveal whether there are consistent patterns in people's perceptions across a population. Still, it may conceal hidden heterogeneity that can only be revealed with more fine-grained measures. Thus, we also constructed measures of iconic potential from grouped subsets of the data.

Fig. 19 18th ranked image. *Note*: Photography by Pete Souza, The White House

Fig. 20 19th ranked image. *Note*: Photography by Pete Souza, The White House

In sum, we created three types of measures of iconic potential by subsetting our sample and reaggregating the results of the pairwise comparison tasks along three dimensions: individual-level demographic variables, individual-level perceptions of the Presidency, and taste-based subsets based on micro-community cluster assignments created from a "top-three" ranking task. From there, we compared these subsets to each other and the full sample to assess the stability of each image's iconic potential.

We performed a series of *t* tests for each subset to identify which image qualities were salient for evaluating iconic potential for the group in question. To assess the stability of an image's iconic potential, we compare the similarity in images' iconic

Fig. 21 20th ranked image. *Note*: Photography by Pete Souza, The White House

Table 1 Image-level variables–bolded variables were included in the analysis

Kind	Variables
Symbols	**U.S. Flag,** Presidential Seal, Representations of other Presidents (e.g., portraits, busts)
Co-present	**Kids;** military; **advisors/staff;** other world leaders; **citizens**
Obama qualities	**Formality of dress; giving speech; smiling; frowning; talking**
Body position	**Facing forward; in profile;** back turned to camera; neutral body position; head position; hand activity; **standing; sitting;** squatting
Context	**Indoors**; outdoors; at White House; Obama alone; **Obama with a crowd;** number of people; **political backdrop; legibility of event;** formality of event; racial make-up of those co-present
Camera	Camera angle; **out of focus; close up**
Image Qualities Measured with imagefluency	**Symmetry (vertical); symmetry (horizontal); complexity; typicality; contrast; self-similarity**

potential (via win percentage and ranked position) between the full sample and the subsets.

Subsets based on demographic variables

First, we divided the sample using individual-level demographic variables—specifically gender, race-ethnicity, education, and partisanship. Table 2 shows summary statistics for these variables (and more, which we elaborate on in the next section). We then calculated the image-level win percentages for these demographic subsets, precisely as we did with the full sample.

Table 2 Descriptive statistics for individual-level variables

	N	Mean/prop	SD	Min	25p	Median	75p	Max
Variable								
Gender	914							
Female	447	0.489						
Male	467	0.511						
Race[a]	900							
White	624	0.693						
Black	90	0.100						
Hispanic/Latinx	49	0.054						
Asian	84	0.093						
Multiracial	53	0.059						
Education	919							
Less than a 4-year degree	360	0.392						
At least a 4-year degree	559	0.608						
Political party	919							
Democrat	374	0.407						
Republican	202	0.220						
Independent	313	0.341						
Other	30	0.033						
Meanings of the Presidency (1 = highest ranking, 7 = lowest ranking)	919							
Chief Legislator		4.62	1.51	1	1	5	6	7
Chief of Party		5.57	1.70	1	5	6	7	7
Chief Executive		3.01	1.53	1	2	3	4	7
Chief Citizen		2.61	2.10	1	1	1	4	7
Chief of State		4.02	2.04	1	2	4	6	7
Commander in Chief		3.69	1.84	1	2	4	5	7
Chief Diplomat		4.48	1.58	1	3	5	6	7

Proportions may not sum to 1 due to rounding.

[a]The "American Indian" and "Other" categories were left out of the analyses due to small sample size.

Subsets based on meanings of the presidency

We also subset the data according to how people understood the office of the presidency because how people make sense of the presidency might influence what representations of President Obama they view as more iconic. To measure meanings of the office of the presidency, we asked people to rank the different roles of the president in order of importance. The roles they rank (and their definitions) come from civics textbooks that describe the "seven hats of the presidency": Chief of State (ceremonial head of government), Chief Executive (head of the executive branch), Commander in Chief (head of the nation's armed forces), Chief Citizen (representative of all of the people), Chief Diplomat (architect of American foreign policy), Chief

Legislator (proposer of legislative agenda), and Chief of Party (Leader of the political party that controls the executive branch). We use their ranking of these seven hats as a measure of respondents' meanings of the presidency to determine whether perceptions about the office of the president influenced perceptions of relative iconicity. We then calculated the image-level win percentages for each of these subsets, placing a respondent in one and only one subset based on their top ranking.

Subsets based on micro-community clustering

Dividing the data by individual-level variables may miss hidden heterogeneity if there is variation that does not cleanly map onto demographics or perceptions of the Presidency. To address these limitations and maximize our ability to identify heterogeneity in salient image qualities affording people's perceptions of relative iconicity, we also took a more inductive and granular approach to divide the data using clustering methods. The aim was to infer taste communities based on similarities in judgments of relative iconicity rather than deductively sorting people according to pre-existing categories.

We needed individual-level judgments to do any clustering analysis based on individual judgments of relative iconicity. However, because most of the 190 possible image pairs were not seen by everyone, the pairwise data did not allow direct individual comparison. We supplemented the pairwise comparison task with another ranking task to address this limitation. After completing their pairwise comparisons, participants were shown all 20 images simultaneously and asked to choose the three "most iconic" images.

This "top-three" ranking task provided useful individual-level information, allowing us to create new aggregate measures of iconic potential across inductively produced categories. We assume that if people's sets of three "most iconic" images align, they likely share comparable understandings of what makes an image iconic. We, thus, assigned people into clusters based on their top-three choices, subsetting the data according to these cluster assignments. This permitted us to calculate and compare the iconic potential measures across these subsets, revealing previously missed differences in iconic tastes.

Different clustering methods are available, which are more or less suitable depending on the research goals. Our choice was informed by current theories of concept formation (Barsalou 2009, 2016). As an individual-level understanding, iconicity is tacitly built up from individual experience with images in the world rather than from clearly defined rules. Because of this, we suspect there may be multiple pathways through which an image may be read as iconic. As such, we wanted a clustering method that maximized communities of taste rather than one that privileged the most parsimonious clustering solution. As a test of the stability of iconic potential across groups, this offers a highly *conservative test.* If micro-community distributions of iconic potential are statistically similar to the full-sample distribution—even as we maximize the number of clusters—that suggests an image's iconic potential is *very* stable.

We clustered individual-level responses to the "most iconic" image task using Lizardo's (2024) method for identifying microgenre communities of taste with

survey data. Lizardo's approach permits increasingly thin but heterogeneous clusters of micro-communities. The microgenre community approach is designed to get past the limitation of clustering based on macro-genre labels. The problem, as Lizardo describes it, is that in the real world, genres are fuzzy and overlap, so clustering based on macro-genre labels misses hidden heterogeneity. Lizardo's method uses "link clustering," which takes the linkage between a person and a genre choice as its unit of analysis. These person-genre links are clustered with other links that have similar co-occurring genres. Thus, within any genre, such as "classical," there may be multiple overlapping "microgenres" based on similar sets of genre choices that include classical music and other distinct tastes.

The applicability of the microgenre community approach to the current study may not immediately be apparent, given we are dealing with selections of images rather than genres. However, just as two people expressing a taste in classical music may, in reality, prefer distinct microgenres of classical music, multiple people who select the same image as the most "iconic" might have different implicit understandings of iconicity that focus on different qualities of the same image. The microgenre approach allows us to identify distinct yet overlapping implicit understandings of iconicity. If we clustered people using more traditional methods based solely on their image choices, we would force oppositions between images and miss potential hidden heterogeneity. Thus, the microgenre community approach allows us to maximize our chances of identifying heterogeneity in people's perceptions of relative iconicity.

After identifying ten distinct micro-communities[7] based on the top-three ranking, we created subsets of our data based on micro-community membership. We then recalculated aggregate win percentages for each of these ten subsets, developing measures of iconic potential for each one.

Findings

Iconicity as affordance

Our first research aim is to determine whether people can evaluate the iconic potential of images based on image qualities or, said differently, whether the images afford perceptions of iconicity. Our survey's pairwise comparison task tests this using our primary outcome variable: iconic potential. Observing the distribution of iconic potential scores across a set of images allows us to determine whether images "afford" evaluations of iconicity without previously being labeled as "iconic." If iconicity were wholly a matter of labeling and "all hits are flukes," we would expect the pairwise results to be predominantly random, with win percentages converging

[7] We used the silhouette width statistic to assess cluster fit for many different cluster solutions. The 10-cluster solution fits the data better than any lower-level solutions; however, we found that, in general, the silhouette width increases monotonically with the number of clusters. As higher-level solutions are hierarchically nested within lower-level solutions (and not *mixes* of lower-level solutions), we chose the cluster solution where qualitative interpretability was still tractable.

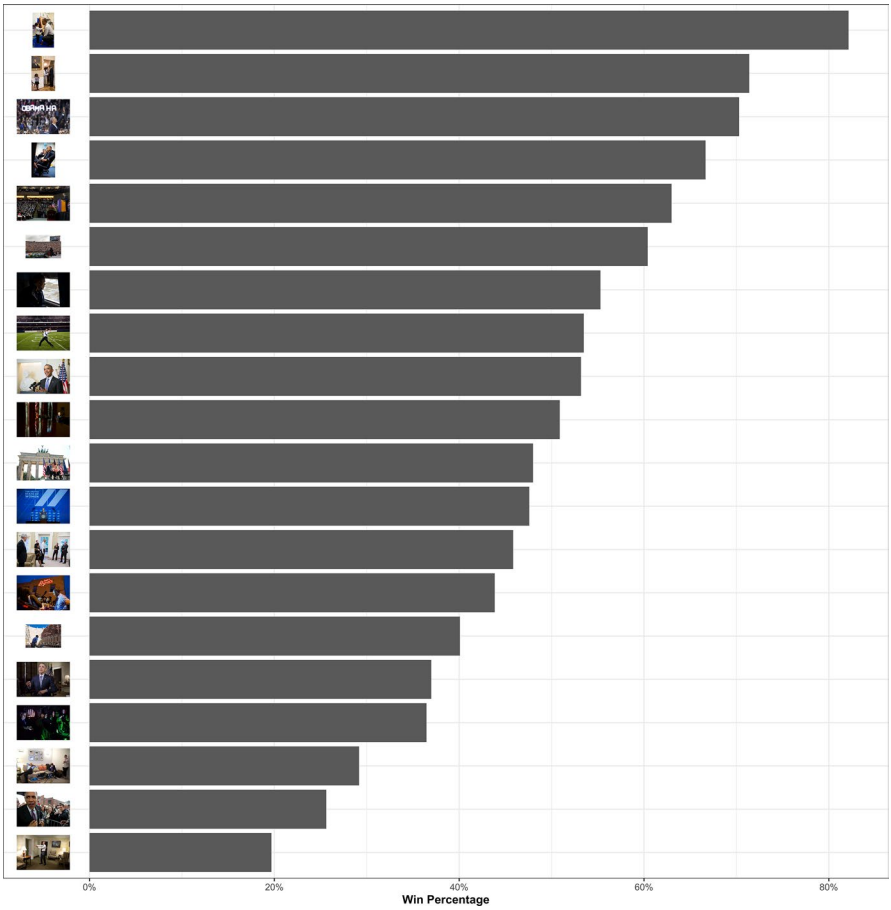

Fig. 22 Full-sample distribution of image iconic potential by win percentage

around the 50% mark because the images in the study are lesser-known images and unlikely to be familiar to audiences, and so do not come with iconic status signals from cultural gatekeepers. If, on the other hand, certain images rank consistently higher or lower than others in terms of iconic potential, this would suggest that some image qualities are more likely to afford judgments of iconicity.

We begin by presenting the distributions of the full-sample aggregate measure of iconic potential (see Fig. 22). If the perception of iconicity is wholly dependent on arbitrary labeling and not associated with image qualities, then we would expect the responses to the pairwise tasks to be random, yielding a flat distribution of iconic potential across the different images. However, as Fig. 22 shows, this is not the case. Instead, we find clear, non-random patterns of evaluation. Collectively, participants consistently evaluated specific images as more iconic than others and other images as consistently less iconic. As these images have yet to be widely circulated, these differences arise due to the visual qualities of those images. Furthermore, at the

extreme ends of the distribution, we see a remarkable level of agreement, with the most iconic image winning 82.1% of its total comparisons and the least iconic image winning only 19.7% of the time.

These findings suggest that, first, the images varied in their ability to afford evaluations of iconicity, and second, iconic potential likely had a high degree of stability across persons. That some images win or lose consistently suggests iconic potential is a plausible mechanism explaining part of what makes images iconic because images vary in their capacity to afford iconicity. We would not expect to see these patterns if achieving public iconic status was purely a function of social construction and labeling. Some images have qualities that enhance their iconic potential.

Iconic potential: image-level analysis

Having established a high degree of agreement in our sample about which images of President Obama have more iconic potential, we now turn our attention to identifying patterns in the *qualities of images* that enhance an image's iconic potential. If people's perceptions of iconic potential do depend on certain image qualities, our next research aim is to identify which image qualities matter. We took two approaches to this task: a statistical approach using a set of measures of image qualities and a qualitative, interpretive approach.

Quantitative analysis of image qualities and relative iconicity

We sought to assess whether specific image qualities correlated with higher relative iconicity. Since our image codes used in the analysis were almost exclusively dichotomous, we used a series of two-sided independent samples t tests to assess which image qualities best predict differences in means in relative iconicity (using $\alpha = 0.05$). The exception here was the set of low-level image feature scores derived from the imagefluency package, which are continuous variables. We used correlation t tests to assess the relationships between these variables and the image win percentages. Given our small image sample size ($N = 20$), not every variable had enough variation to evaluate. Those we were able to test are in bold (see Table 1). In addition, our small sample size at the image level ($N = 20$) makes it challenging to establish statistical significance.[8] *Do any image-level qualities predict higher iconicity?*

Of those we could test, only one image-level variable predicted higher iconicity for the entire sample.[9] That was depicting Obama in *profile* (see the "Full Sample"

[8] For that matter, many of these independent variables did not have enough variation to include in the analysis, as when only one or two images in the sample had an image quality. For instance, only one image included "other world leaders" (e.g., Angela Merkel) and only two had depictions of former presidents.

[9] This includes all variables assessed with independent-sample t tests (image codes) and correlation t tests (image features).

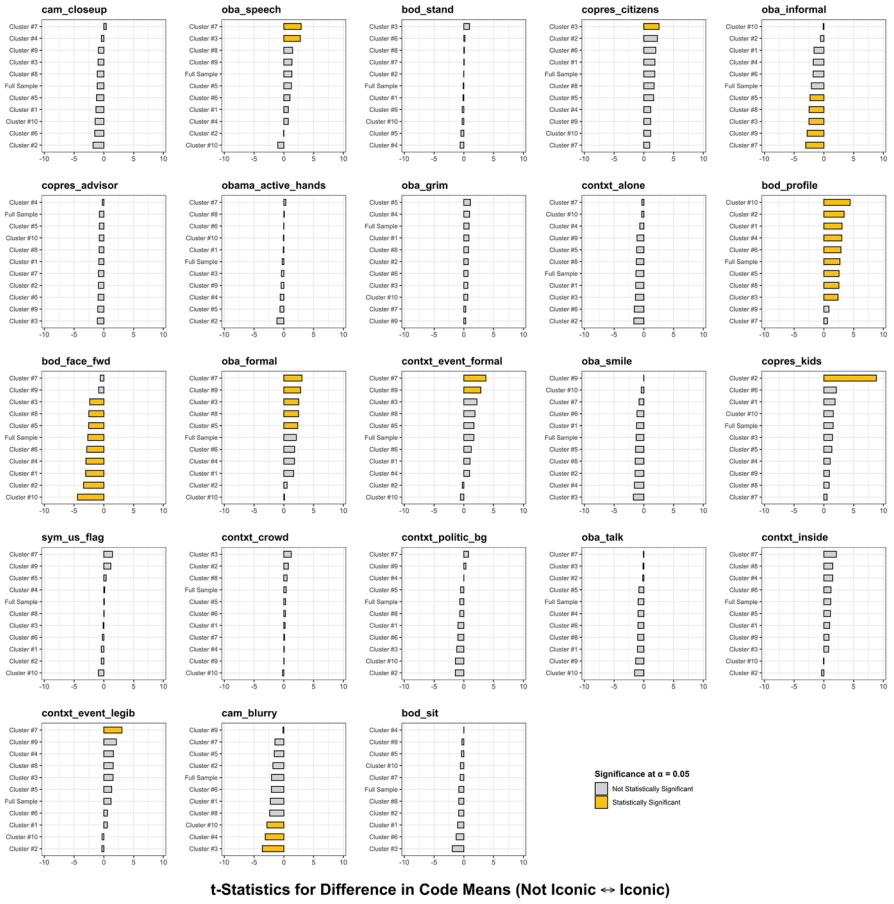

Fig. 23 *t* Tests for differences in image ranking means across image codes, micro-community subsets

rows in Fig. 23).[10] At first glance, it may be surprising that depicting Obama in profile would be significant. People may find President Obama more appealing from the side than facing front. That said, we suspect this pattern has a reason: U.S. coins. Coins are a common and available representation of U.S. presidents and the presidency that have historically depicted former Presidents *in profile*. How are people first exposed to presidents? Coins are a common, if not the most common, way young people first encounter depictions of presidents. Before entering educational environments, most children will have handled and examined coins. Coins are everyday objects that kids play with (spin, slide, stack) that historically have depicted presidents in profile. If we think about how schemas are established, the qualities

[10] "Body Facing Forward" (bod_face_fwd) is also a significant predictor of differences in mean image win percentages, but this is simply the inverse of body in profile. None of the correlation *t* tests were statistically significant for the low-level image feature variables.

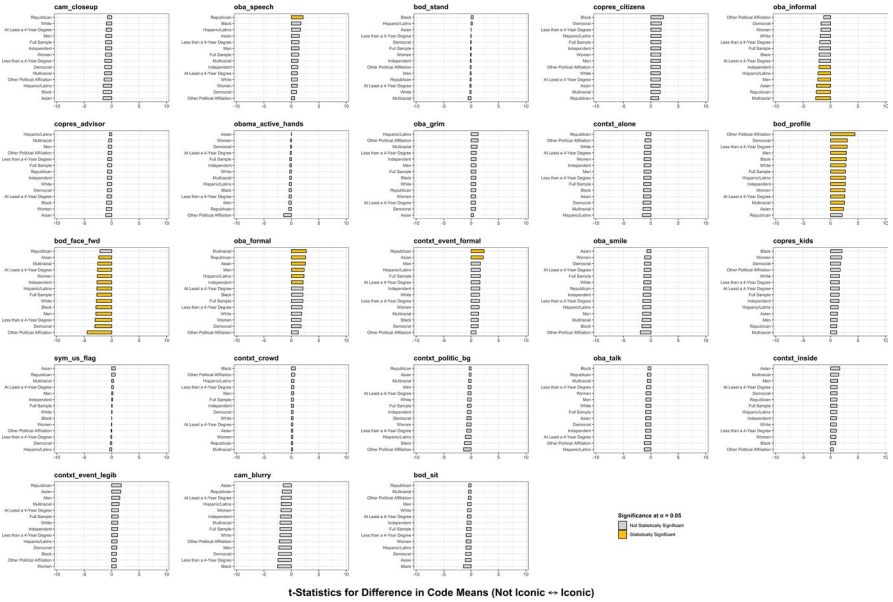

Fig. 24 *t* Tests for differences in image ranking means across image codes, sociodemographic subsets

present in initial exposures to a type-token relationship tend to set the type for how people imagine and think through later tokens (Taylor et al. 2019; McDonnell et al. 2021). It makes sense, then, that seeing a depiction of a president in profile might make them appear more *presidential.*

Figure 23 shows differences in mean relative iconicity between image quality levels (that is, absence versus presence of a quality) and how those differences look across different micro-community subsets of respondents. Figures 24 and 25 show this same information, but for the sociodemographic (Fig. 24) and "meanings of the presidency" (Fig. 25) subsets. As the figures show, the "profile effect" is pronounced, statistically significant, and in the same direction for the majority of subsets. This speaks to the power of image affordances for constructing iconic potential. This is a point on which we elaborate with more detail in a later section.

The consistency of this finding across subsets is overwhelming (Figs. 23, 24, 25), especially when one considers all the different ways the tastes of particular subgroups might inflect their perception of iconicity. For most of these image qualities, people share a *consistent* direction in their orientation toward that quality across sub-groups (i.e., associating it with iconicity or not), even when most image qualities are not statistically significant. Associations are either positive or negative for all (or most) groups, and very few image qualities have a mix of divergent associations. For instance, in addition to our "profile" finding, across groups, people share a perception that Obama in "formal" clothing, at "formal events," "copresent with citizens," and "copresent with kids" are more iconic—even if not statistically significant, the directionality here is telling. Some groups feel these more strongly. For instance, Republicans see images of President Obama in formal clothing, at formal

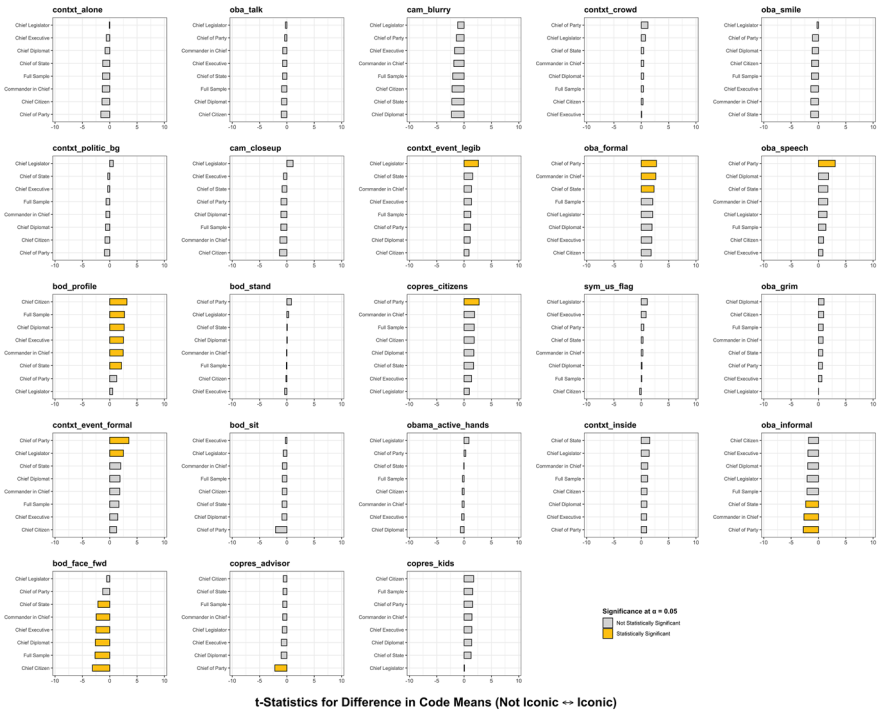

Fig. 25 *T* Tests for differences in image ranking means across image codes, meanings of the presidency subsets

events, and giving a speech as more iconic—suggesting when they don't like the person, they focus on the representation of the *Office* of the President. Additionally, some qualities are consistently negatively associated with iconicity (even if not statistically significant): blurry images or Obama dressed informally. We take from this that Americans have a consistent sense of what makes images of the President or Obama iconic.

Our twenty-image sample size limits the patterns in image qualities that our statistical analysis can reveal. We suspect with a larger sample, several qualities might have become significant. We now complement our image-level analysis by digging deeper into patterns in image qualities with a qualitative analysis.

Qualitative analysis of image qualities and iconic potential

In our qualitative analysis, we focused on the three images with the highest win percentages and the three with the lowest win percentages. We examined these images for common qualities. What qualities might indicate higher iconic potential or negate iconic potential? We engage in interpretive analyses to assess why the images that consistently win or lose better *resemble* or *represent* shared meanings of President Obama or the Office of the President (Smith 2012).

A close reading of the top- and bottom-three images and comparisons of their qualities reveal several distinctive patterns that shed light on how iconic potential operates—namely, capturing a prototypical ideal essence of President Obama and the presence or absence of situational legibility. In the pairwise comparison exercise, the most-winning image was evaluated as more iconic than its competitor image an impressive 82.1% of the time (Fig. 2), and the number two and three ranked images 71.4% and 70.3% of the time (Figs. 3 and 4, respectively). What do these images have in common?

First, all three topmost images show Obama interacting with citizens: playing with a child, greeting citizens to the Oval Office, and waving to a crowd. These images depict Obama as *engaging,* capturing a prototypical ideal for presidents. Much has been made of the importance of presidents (and politicians) engaging the public and appearing likable—someone you'd like to have a beer with. Looking comfortable and engaging, rather than uncomfortable or awkward, can be essential for political success. Speaking of Obama in particular, the press talked about struggles with such interactions. "After examining recent photos from the campaign trail, 'a body language and behaviour expert' finds President Barack Obama to be 'fatally ill-at-ease' while holding an infant."[11] For that matter, Obama had a reputation for being intellectual and aloof. As the New York Times Magazine noted: "Obama does not relish glad-handing. That's what he has Vice President Joe Biden for."[12] In contrast to this reputation, the top images depict Obama comfortably engaging people in ways that align with shared ideals of a personable, engaging man of the people.

Physical touch communicates this sense of being engaging. These top two images are the *only images* in the sample in which President Obama is making direct physical contact with citizens. Physical access to presidents is often limited due to their status and intense security. Rather than appearing at a distance from people, these images depict President Obama touching and directly interacting with citizens. One quality that distinguishes the top-ranked image is that President Obama is allowing *himself* to be touched—indicative of openness. More than most images in this data set, the top image humanizes President Obama. The third-ranked image (Fig. 4)—depicting President Obama waving at supporters—also aligns with this pattern. While not *physically* touching citizens, he engages and interacts with the public.

More than just depicting a President as engaging, the historic importance of Obama as the first Black President of the United States appears to matter for iconic potential. The top two images (Figs. 2 and 3) capture President Obama directly interacting with Black citizens. In the most-winning image (Fig. 2), Obama as role model and inspiration are clear. As we see this child playing doctor juxtaposed with the most powerful man in the country, it calls to mind the hope and aspirations Obama represents for African Americans. That President Obama is in the White House makes it seem like anything is possible for this child—becoming a doctor or becoming president is made manifest in Obama's being. Obama metonymically stands for hope, possibility, and opportunity for all Americans. The second-ranked

[11] https://www.motherjones.com/media/2012/01/politicians-kissing-babies-brief-history/

[12] https://www.nytimes.com/2010/10/17/magazine/17obama-t.html

image (Fig. 3) has President Obama welcoming an African-American family into the Oval Office. With the portrait of Lincoln looking on, this image celebrates possibility and captures the advances toward equality we've made while also calling out how far we've had to come. In both images, Obama is more than just a President but symbolizes something more profound. Contrasted with the bottom-three images, there is a clear difference in how the qualities of these two images call upon specific aspects of the historic nature of Obama's presidency.

Children also emerge as a common quality of these two most iconic images. Our sample has three images of President Obama co-present with children, and two of these had the highest iconic potential. We can imagine several explanations for the bump in iconicity around images of President Obama with kids. Just as politicians on the campaign trail receive positive attention for "kissing babies," interacting with kids may humanize the President and make him appear more relatable.

Another image in the set (Fig. 15) also captures President Obama interacting with a group of Black citizens—including kids—outside Arthur Bryant's BBQ restaurant in Kansas City. This image had less iconic potential (43.8% win rate). As it captured many qualities of the top two images (presence of children and Black Americans), one might expect it to have similar levels of iconic potential. Why not? Compared to the top two images, the interactions seem less poignant. Unlike the White House image, President Obama towers over the children in this image. Rather than meaningful back-and-forth interactions, it appears that Obama is talking to the boy in the image, who is dutifully listening. While the boy is attentive, the girl is not, which may undermine the iconic power of Obama in this photograph.

Legibility is another factor that may explain the difference between Fig. 15 and Figs. 2 and 3. Figures 2 and 3 depict legible settings (e.g., classroom, Oval Office) that lend important context to understanding this picture and deepening the meaning. Figure 15 shows Obama in a crowd, but many respondents may not reflexively be able to 'place' the event. Arthur Bryant's is known in Kansas City as a local institution that welcomes politicians and gives them a platform for engaging the public. That is not knowledge most Americans have; therefore, the image is less legible than the top two. This suggests legibility matters for iconicity—or rather, lack of legibility may be anti-iconic.

Contrasting the top-three images to the bottom images (see Figs. 19, 20, and 21), a lack of legibility and engagement may explain the lower iconic potential. These images share specific qualities: Obama is often looking away from the people with whom he is co-present, undermining the prototypical ideal of a President engaging with citizens or in situations that are less immediately legible.

The 3rd lowest-ranked image (Fig. 19) shows Obama distanced from his staff, on the phone and staring into space. Like the rally outside Arthur Bryant's, this image also has relatively low legibility—the situation is unclear, and it is unknown who Obama is speaking with on the phone. Obama is not engaging the people in the room, who are looking at their phones and are not drawn in by the president's aura, as their interactions with the President are more mundane than everyday citizens. The 2nd lowest-ranked image (Fig. 20) has Obama with a crowd, but he appears to actively ignore them by looking away from them. A fence physically separates him from the crowd. Obama's distinct lack of engagement, along with a blurry image

and an awkward facial expression, likely undermine the iconic potential here. Even though the crowd appears to show him reverence, Obama's orientation to them may undermine the iconicity. The legibility of the situation is also undermining this image—he seems in transit, between situations. The bottom-ranked image (Fig. 21) has a distinct lack of legibility—is he interacting with someone? If so, who? The location is vague (is it the White House? Or a hotel?), and the social context is a mystery. In each of these situations, there is a lack of interaction and limited contextual information to improve one's confidence in what is happening, undermining the legibility of these images.

Image qualities and iconic potential

What do these qualitative and quantitative assessments of image qualities tell us about iconic potential? These findings tell us something about iconicity as "representation" and "resemblance" (Smith 2012) and what we call "anti-iconic qualities."

What meanings does an image signify, and how does an image afford the "feeling" of those ideas (Alexander 2008)? This is the question of how image qualities come to prototypically *represent* those ideas that bind people around shared social values. The patterns in our data suggest the images with more iconic potential are those with qualities that afford prototypically ideal meanings of President Obama as engaged, gregarious, inspirational, and humble. What is more difficult to parse is whether these are qualities people value in Barack Obama the person, or Barack Obama as metonymically standing in for the Office of the President, or both—which is likely given how tightly imbricated Obama the man and the President are. Given that iconicity has been theorized as a feeling, people are not likely deliberately thinking through these possibilities of iconic *in relation to what,* even as their pre-existing ideas of Obama and the Presidency (and other ideas like masculinity, race, etc.) inform their interpretations of whether these images are iconic. Instead, these images offer a *gestalt*—a sense of iconicity—that aligns with people's sense of President Obama.

This analysis also reveals there are associational meanings at the level of *resemblance*. In this domain of images of the President, our respondents appear to strongly associate iconicity with depictions of Obama in profile. This suggests we have unconscious and inarticulable preferences for this quality, which we suggest is driven by everyday exposure to coins or other depictions of presidents in profile. This preference is likely unconscious and hard to articulate, based upon a widely shared prototype developed from repeated exposure to tokens of iconic images of presidents that share that quality (Taylor et al. 2019).

Notably, some qualities of images of President Obama appear *anti-iconic.* For instance, images have qualities that violate expectations for what people imagine as Presidential. Images that depict the president interactionally disengaged with co-present others. Additionally, images that are less legible in their social situation lead people to do too much cognitive work. If iconicity works through its automaticity, having to think too much to get it to "fit" one's idea of iconicity pushes it out of contention. Similarly, a lack of *legibility* is cognitively burdensome and undermines the ability of an image to *feel* iconic. Finally, images with technical problems also

have anti-iconic qualities. We find that for some micro-communities, out-of-focus images are negatively correlated with iconic potential (Figure 23 shows that blurry images were negatively correlated with win percentage for micro-communities 3, 4, and 10), and we imagine that an image with Obama's head out of frame would be less iconic than any of the images we have in our sample.

We have established some patterns in the qualities of images and the mechanisms through which they afford (or not) assessments of iconicity. In the next section, we ask the subsequent question: How stable are the affordances of these images across people? Said differently, do most people's evaluations of iconicity appear to be based on similar visual qualities?

Stability of iconic potential

How stable are the affordances of iconicity across people? To assess stability, we examined whether the most and least iconic images are consistently toward the top of the distribution, regardless of micro-communities, demographic categories, or how respondents define the role of the president. This offers a measure of the "stability" of iconic potential of an image across groups. If an image is stable in its relative position regardless of sub-group, that suggests that the image affordances do more to stabilize iconic potential than the cultural tastes of the audience. Inversely, if an image's relative position is unstable, meaning its ranking shifts considerably from sub-group to sub-group, then the unique tastes of that sub-group would drive evaluations of iconic potential.

Figures 26, 27, and 28 depict the stability of the top- and bottom-three images across micro-communities, demographic categories, and groups partitioned by who ranked which presidential role as most central. In each figure, the leftmost column is the full-sample pairwise ranking distribution, with the distribution for each subset moving from the second column to the right (sorted by the win percentage of the topmost image). We see that the top image in the full sample remains among the top three for *all but one* subset—and that one exception (people who see the president first and foremost as the Chief Legislator) place that image fourth. Similarly, the bottom-ranked image in the full sample remains among the bottom three for all but one subset (again, the Chief Legislator subset). Overall, the top- and bottom-three ranked images in the full sample hold steady in the top and bottom of the subset distributions. This strongly suggests that the qualities of these images drive stability at the top and bottom of the distribution.[13] Looking at the middle of the distribution, we see less stability.

Looking at Figs. 29, 30, and 31 gives us a higher resolution analysis of stability. Figure 29 compares the pairwise ranking of the full sample (the leftmost distribution of each pair) to the rankings for each micro-community to the right of each pair

[13] This observation matches the Salganik et al. observation in their study of social influence in the popularity of songs in a cultural market: "In general, the "best" songs never do very badly, and the "worst" songs never do extremely well" (2006, p. 855)'s. Despite finding strong evidence for social influence in their study, they see qualities acting back on people's assessments.

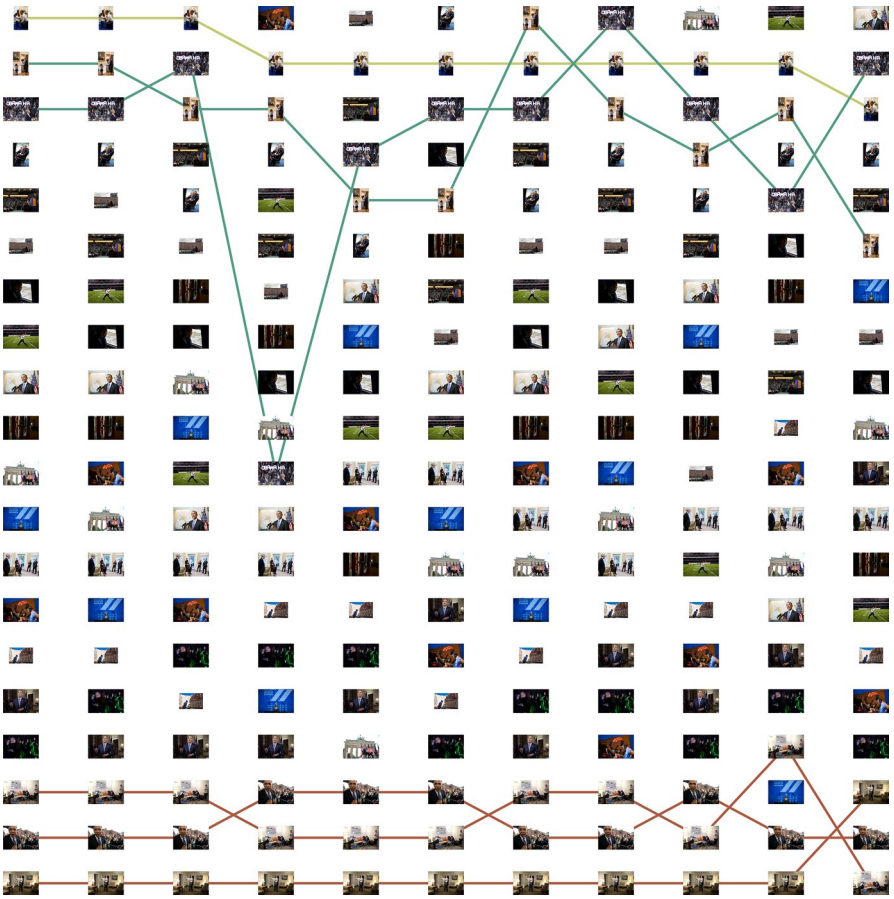

Fig. 26 Visual representation of the stability of the top- and bottom-three images across micro-communities. *Note*: The leftmost column is the full-sample pairwise comparison rankings from 1 to 20. The other columns represent the 10 micro-communities pairwise rankings

(micro-community one through ten, moving upper right to bottom left). Figures 30 and 31 report the same information but with the demographic and presidential role subsets. Blue lines represent the movement of the top-three images from the full sample to the subset distribution, and yellow lines make the movement of the bottom-three images visible. The red lines depict the three images whose winning percentages move the most (up or down the distribution) for this subset compared to the full-sample distribution. If the biggest movers are also a top or bottom-ranked image from the full sample, we color it red (rather than blue or yellow). Light gray lines represent the movement of all other images.

Looking at these blue lines, we see consistency among the top images. Looking at the micro-community subsets (Fig. 29), the top image stays in the top position for two micro-communities and is in the top three of all ten. The top image never dips below an 80.4% winning percentage in any micro-community and goes

Fig. 27 Visual representation of the stability of the top- and bottom-three images across demographic groups. *Note*: The leftmost column is the full-sample pairwise comparison rankings from 1 to 20. The other columns represent the 13 demographic subsets (women, men, Asian respondents, Black respondents, Hispanic/Latinx respondents, multiracial respondents, white respondents, democrats, respondents with less than a four-year degree, respondents with at least a four-year degree, republicans, independents, and respondents who have a different political affiliation)

as high as 92.8% for micro-community one. This suggests that even with different audience tastes, some images consistently rise to the top. Across groups, the more iconic images are stable, and for some communities, the top wins virtually every time (90.8% for Community One, 91.7% for Community Eight, 90.9% for Community Three).

Images at the bottom of the distribution are similarly stable. Across micro-communities, people consistently identify similar images as less iconic. The bottom-ranked image in the full sample remains ranked in the bottom three for each of the ten micro-communities. Its highest win percentage is 24.5% in micro-community ten, and for micro-community three, it won only 13.5% of the time.

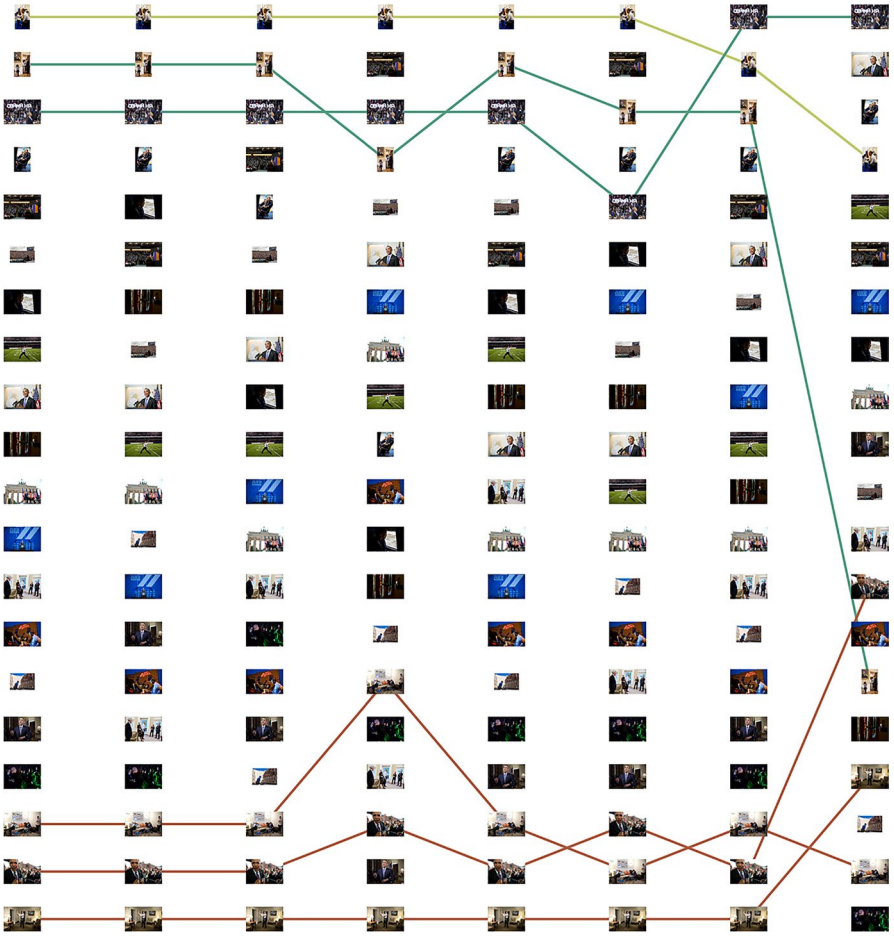

Fig. 28 Visual representation of the stability of the top- and bottom-three images across groups sorted by top ranking of presidential roles. *Note*: The leftmost column is the full-sample pairwise comparison rankings from 1 to 20. The other columns represent the 7 presidential role subsets (Chief Citizen, Chief Legislator, Chief of Party, Chief Executive, Chief of State, Commander in Chief, and Chief Diplomat)

We see evidence of relative stability across the sample for most images. Looking at the slopes of the gray lines, it appears most images in the middle stay within a narrow range of movement. Across micro-communities, for example (Fig. 29), images remained within one rank of their full-sample position 62.5% of the time, within two ranks 77.5% of the time, and within five ranks 94% of the time.

The red lines draw our attention to the images with the largest movement. Even with the three biggest movers, they moved more than three ranked positions only 16 times (53.3% of the time). Five images account for more than half of those 16 instances of big jumps between micro-communities (images 6, 8, 11, 12, 14). These five images had a substantially greater degree of movement than the rest of

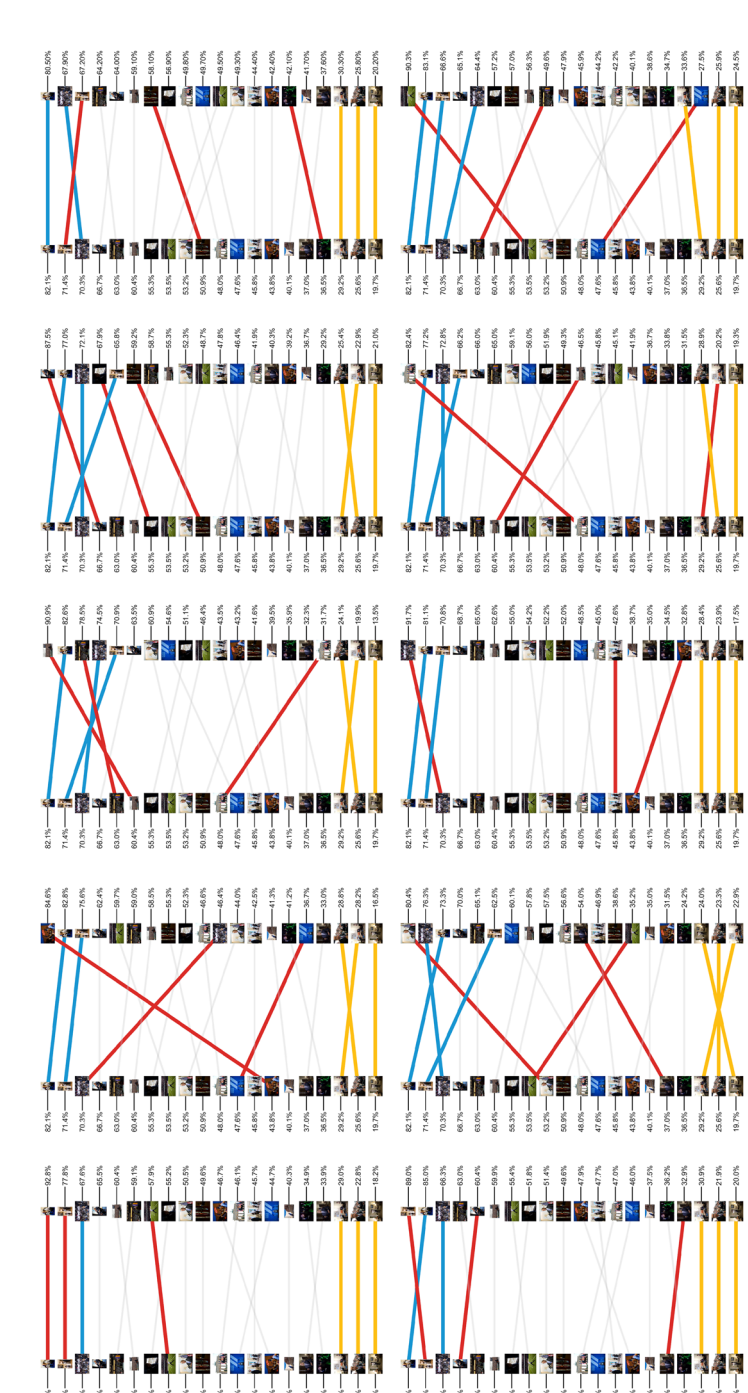

Fig. 29 Comparing the pairwise ranking in full-sample rankings (left) to each micro-community's rankings (right of each pair). *Note*: Full sample v. Micro-community 1 is the upper left, moving in numerical order from right to left, with Full sample v. Micro-community 5 in the upper right position. Blue lines mark movement of the top-three images from the full sample. Yellow lines mark the movement of the bottom-three images from the full sample. Red lines represent the three biggest movers in win percentage from full sample to micro-community

the sample, with an average range of movement across micro-communities of 12.6 and a win percentage range from 32.6 to 55.1%. For instance, Image 8 (Obama throwing the football) had the largest range, with a win percentage of 90.3% for one micro-community and a low of 35.2% for another. This suggests great variation in how people read the iconic potential of this image. Put differently, these five images were somewhat less stable and more susceptible to differences across audience tastes. That said, even the most swingy image (Obama throwing the football) remained within three rankings of the overall position for seven out of the ten micro-communities.

Beyond these five, with large variations in iconic potential depending on which micro-community is viewing it, there is much consistency and stability at the image level. This pattern of stability drives the high overall correlations among the micro-communities, demographic categories, and presidential role groups. People appear to know iconicity when they see it (or when they do not). This finding suggests that people have a *widely shared iconic taste* in the context of images of President Obama.

Discussion

Summary of findings

In this study, we developed a novel approach to studying iconicity. In contrast to previous work that analyzes already established iconic images, we sought to elicit people's understandings of relative iconicity by measuring their perceptions of not yet consecrated images. While such an approach may not comprehensively account for a particular image's trajectory to iconic status, it sheds light on "iconic potential." Iconicity is not merely a label that arbitrarily sets apart images as "iconic" regardless of image qualities, but an affordance—an implicit perception emerging from the interplay between image qualities and audience capacities for a particular subject (e.g., Barack Obama, or "the Presidency").

Our first analytic goal was to test the iconicity-as-affordance hypothesis by ascertaining whether people could consistently evaluate the relative iconicity of images to which they would likely have had little exposure. The strong patterns in the results of the pairwise task support the idea that iconicity is indeed an affordance.

If iconicity emerges from an interaction between certain image qualities and audience capacities, what are the salient object qualities in this interaction? Our second analytic goal was to begin to inductively identify image qualities associated with higher iconic potential around President Obama. In the whole-sample image-level statistical analysis, we found that images with Obama facing in profile were more likely to have higher win percentages. Our qualitative analysis showed that images of Obama interacting with citizens—especially children—tended to have more iconic potential.

Our third analytic goal was to test the stability of these images' iconic potential across distinctive sub-samples. Do people share implicit understandings of iconicity, or do different sub-samples have distinct understandings of what makes images

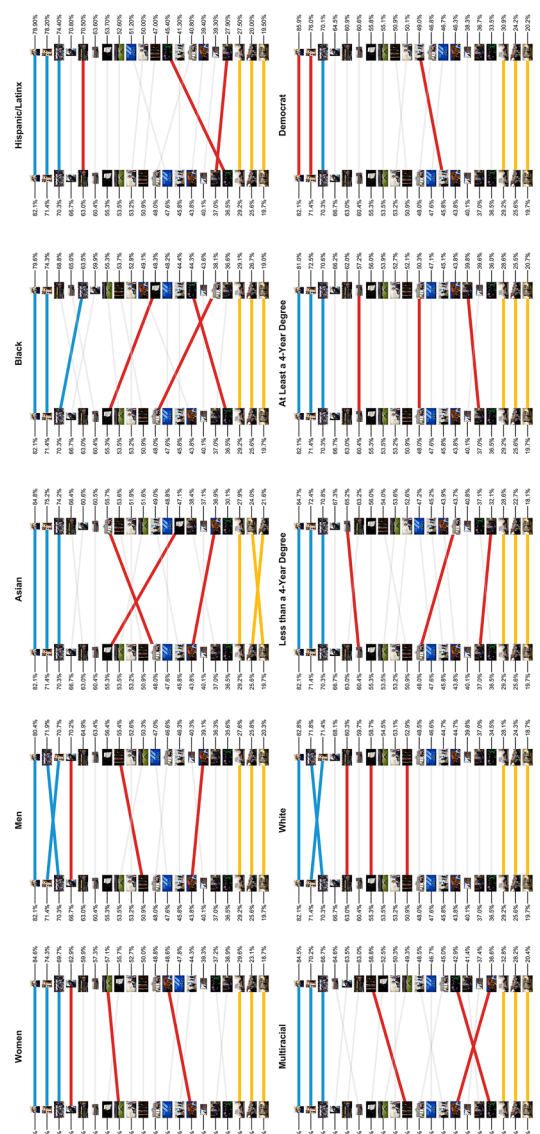

Fig. 30 Comparing the pairwise ranking in full-sample rankings (left) to each demographic group's rankings (right of each pair). *Note*: Blue lines mark movement of the top-three images from the full sample. Yellow lines mark the movement of the bottom-three images from the full sample. Red lines represent the three biggest movers in win percentage from full sample to demographic subset

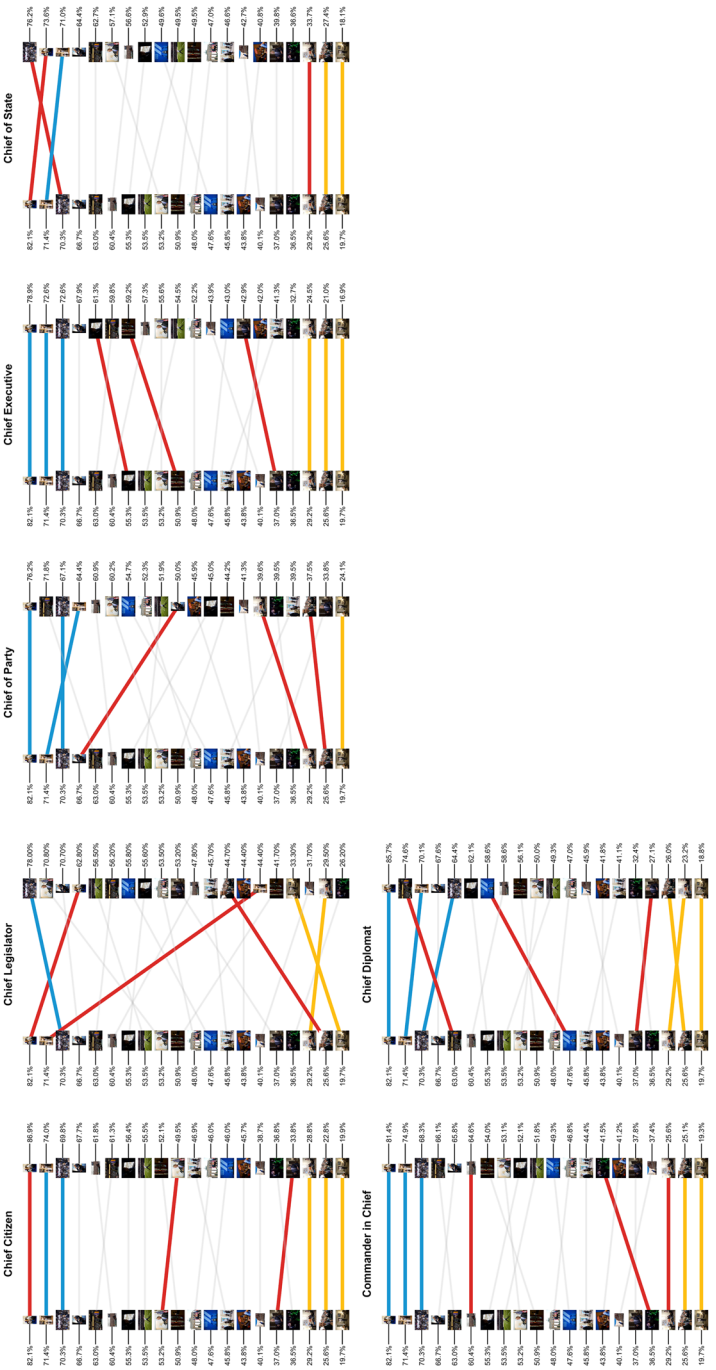

Fig. 31 Comparing the pairwise ranking in full-sample rankings (left) to each presidential role group's rankings (right of each pair). *Note*: Blue lines mark movement of the top-three images from the full sample. Yellow lines mark the movement of the bottom-three images from the full sample. Red lines represent the three biggest movers in win percentage from full sample to role subset

of President Obama more or less iconic? We evaluated the stability of images with higher (and lower) iconic potential by comparing the distributions across our ten micro-communities, demographic groups, and groups with shared understandings of the presidency. We found that (1) the top and bottom images in the distribution held steady, (2) the position of most images in the middle of the distribution rarely moved much from micro-community to micro-community, and (3) only five images seemed to swing widely in their iconic potential depending on which audience was looking at them. While small differences in evaluations reshuffled a few images, the overall distributions were not statistically significantly different from the full sample. This stability suggests a strong association between object qualities and audiences' perceptions of what makes an image iconic.

Measurement and modeling

One potential concern is whether our method is "really" measuring iconicity. However, this concern stems from a misunderstanding of the analytic goals. Respondents were instructed to make judgments about the relative iconicity of a set of images, and they did so in systematic, non-arbitrary ways. It is possible that when faced with this task, respondents consciously or unconsciously made evaluations according to a related concept, like "likability." However, whether a person's subjective perception of iconicity is parasitic on concepts like "likeability" is an orthogonal concern. The fact remains that participants responded in patterned ways, suggesting that people do have the hypothesized capacity to perceive and make systematic judgments about iconicity. To suggest that respondents were not really making judgments about iconicity implies that they had a secret "iconicity" concept that they chose to withhold for arbitrary reasons. This seems implausible.

These findings offer an initial step toward understanding *how* image qualities matter for iconicity. Many questions remain unanswered and will require further data collection. We have sought to isolate the effects of image qualities on iconic potential and assess the capacity of images to afford stable assessments of iconicity, regardless of audience. By identifying qualities associated with iconic potential, we do not argue that such qualities perfectly predict the likelihood that an image will become iconic. Social influence and consecration processes also shape what images ultimately rise to the level of iconic (Alexander 2012). We argue that image qualities are necessary but insufficient to explain iconicity.

We may have missed important qualities as we privileged the coding of more "objective" measures—more subjective assessments of qualities of emotionality and aesthetic style may matter as much or more than the qualities we've measured and analyzed. We also warn against extending these results to other domains given that the image qualities matching the prototypical ideal for "Presidential" or Obama may not generalize beyond this domain. We imagine qualities that enhance the iconic potential of images of Obama likely enhance iconic potential in comparable domains: Other presidents, political leaders, etc. Alternatively, if this were a study of Marilyn Monroe, we wouldn't expect that speaking at a podium or engaging with kids would match the prototypical ideal. The image qualities we've uncovered

as aligned with iconic potential may extend to cases with a family resemblance to Presidents. Our account of iconic potential is not "universal," nor do we argue that every image is idiosyncratic. Domain specificity matters.

We've asked about which qualities seem to afford iconic potential, but we have not explicitly engaged the question of how relations among qualities might shape iconicity. How are image qualities related? Are qualities "additive" such that an image of Obama in profile + children + giving a speech + interaction with citizens = more iconic potential? Or are qualities "multiplicative," such that images of Obama "with children" x "appearing informal" amplify the effect of either quality, humanizing him more than each alone (Seguin 2023)? Alternatively, a particularly excellent depiction of an emotion—Obama smiling with joy—may have such power that it can draw an audience's focus away from the anti-iconic qualities of the image. More work needs to be done to theorize the different mechanisms through which qualities work in tandem to increase (or decrease) iconic potential.

The inarticulable sense of what makes an image iconic means that our data capture people's implicit, personal culture resulting from enculturation (Lizardo 2017). One benefit of our approach is that it draws out implicit patterns from the ground up. Given the way iconic images communicate a feeling, people are unlikely to actively and deliberately weigh these qualities against each other. We argue it is more likely that people assess iconicity through a *gestalt*. In this sense, we caution against reading these findings as "hard and fast" rules for what makes an image iconic—like other work on categorization, people's evaluation of iconicity likely engages prototypes and works through fuzzy logic but requires a baseline sense of legibility.

While we have shown broad consensus among the Americans we sampled for what images have iconic potential, we have yet to delve into the question of multiple pathways—how the same quality might afford different meanings and, therefore, different routes to iconicity. Even as the *outcome* of relative iconic potential remains stable, we don't know whether there are multiple ways people could reach the same outcome. It could be that Fig. 13 is iconic because President Obama is standing on a podium (matching those with a taste for formal representations of the President) or because President Obama is speaking on women's issues (satisfying feminist audiences). Griswold's (1987) theory of cultural power is particularly relevant for thinking about such competing paths to iconicity. Different audiences may view the same image (and its qualities) as iconic for different reasons, putting that image to different ends and renewing its iconic power through debate and dialogue.

Future research

Though we established *that* visual qualities matter, the task of nailing down precise explanations for *how image qualities matter* for iconicity remains incomplete. Two hypotheses require greater investigation:

First, there is the *Captured Essence Hypothesis*. Some work argues that an image is more likely to be perceived as iconic if it evokes meanings consistent with the perceiver's conceptual understanding of the subject. We call this the "captured essence"

hypothesis. Notably, this explanation of "how image qualities matter" relies less on a fixed concept(s) of "iconicity" and more on conceptual understandings of the subject. In this sense, the nature of the implicit understanding of "iconicity" is more procedural (i.e., does this image match my understanding of the subject?) and less substantive. Future research could test this hypothesis more directly by pairing subjective evaluations with open-ended questions to measure people's webs of meaning about the subject.

Second, there is the *Domain-Specific Iconicity Schema Hypothesis*. This hypothesis captures the possibility that people have a general concept of "iconicity" such that "iconic" images directly evoke the concept. This concept would entail associations with specific kinds of image qualities. If people possess an "iconicity schema," it could be because they have schematized from past experiences with images that others have labeled "iconic." Field theorists have made similar arguments about how people come to grasp what is valuable in a field (Green. 2013; Martin 2015). Attempts to identify a general theory that accounts for qualities of images across domains that make them iconic have proven untenable. Therefore, the qualities that activate this schema are likely *domain-specific*—otherwise, it would have been much easier to identify shared qualities of the genre. This hypothesis suggests that people schematize iconicity differently across different domains. For example, if this hypothesis is true, then the image qualities that are salient when evaluating the relative iconicity of images of presidents may *not* be salient when evaluating different types of subjects.

What do our findings suggest about the relative plausibility of these hypotheses? The finding that children appear to predict higher relative iconicity might support the "captured essence" hypothesis. If people value a more humanized president, they might rate images of Obama with kids as more iconic. Or, those who view Obama as a symbol of the advancement of civil rights or as a role model for African Americans might prefer images of Obama engaging with Black Americans.

Alternatively, our finding that pictures of Obama in profile position were associated with higher relative iconicity could be interpreted as supporting the domain-specific iconicity schema hypothesis, especially if people do, in fact, schematize iconicity in the domain of presidential images from their experiences with coins. Moreover, certain micro-communities appear to have unique domain-specific preferences that likely emerge from exposure, as when we see a preference for commonly circulated images of presidents standing before a podium at a formal event.

Ultimately, however, our findings cannot definitively judge which hypothesis of "how image qualities matter" better explains iconicity. Given that we have evidence for both, we suspect both processes are at work. Future studies would need more targeted measures to refine our understanding of how people's understandings of a subject's essence and domain-specific iconicity schemas operate, separately or in tandem.

Conclusion

These findings lead us to some conclusions. First, iconicity is not a taste independent of subject or domain such that a "universal" theory is possible. How people understand the subject and their experience with the domain influence the kinds of qualities that afford iconicity. Second, the stability of an image's iconic potential across demographics, political positions, and micro-communities is an *accomplishment.* This finding is unexpected given reception theory (Griswold 1987; Rawlings and Childress 2019), which would predict that variation in experience and dispositions across groups would lead people to have more widely divergent tastes. Such an accomplishment is first a matter of the stabilizing power of image qualities. If people's judgments of the relative "iconicity" of images display systematic patterns, then image qualities must matter because they are the basis of judgment. However, which material qualities matter is also a matter of enculturation into shared cultural understandings of Obama and the Presidency and through the circulation and familiarity with iconic images in the domain. The material qualities significant to a person's judgment of iconicity will be influenced by their lived experiences with iconic images. Individuals likely develop an iconic taste based on their aggregated experiences with images publicly labeled as "iconic" within a given domain. The circulation of images (and their qualities) is foundational to this process of developing such a taste. In this sense, the affordances for iconicity that emerge from the interaction of image qualities and people's enculturated tastes are likely driven by experiences with images and domains that cut across identity. That iconicity appears to be widely shared suggests how powerfully iconic images can bring people together and do important cultural work.

Acknowledgements This work was presented at Northwestern University's colloquium and at a Sociology of Culture regular session at the American Sociological Association meetings. This work received essential feedback from the Culture Workshop at the University of Notre Dame. Many thanks to Omar Lizardo for guidance with his pathbreaking methods. Thanks to our readers Erin McDonnell, Ann Mische, and David Gibson. Special thanks to the AJCS anonymous reviewers' developmental suggestions and the expert guidance of editor Jeffrey Alexander. The data collection for this paper was funded through a Small Grant for Research and Creative Work from Notre Dame's Institute for Scholarship in the Liberal Arts.

Data availability A replication repository for this article can be found at: https://github.com/MarshallSoc/iconicity_ajcs

References

Alexander, J., et al. 2012. *Iconic Power: Materiality and Meaning in Social Life*. New York: Springer.

Alexander, J.C. 2008. Iconic Experience in Art and Life: Surface/Depth Beginning with Giacometti's Standing Woman. *Theory, Culture & Society* 25 (5): 1–19.

Alexander, J.C. 2012. Iconic Power and Performance: The Role of the Critic. In *Iconic Power: Materiality and Meaning in Social Life*, ed. J.C. Alexander, D. Bartmański, and B. Giesen, 25–35. New York: Palgrave Macmillan US.

Barsalou, L.W. 2009. Simulation, situated conceptualization, and prediction. *Philosophical Transactions of the Royal Society of London Series b, Biological Sciences* 364 (1521): 1281–1289.

Barsalou, L.W. 2016. Situated Conceptualization: Theory and Applications. In *Foundations of Embodied Cognition: Perceptual and Emotional Embodiment*, ed. Y. Coello, 11–37. New York, NY: Routledge/Taylor & Francis Group.
Barthes, R. 1981. *Camera Lucida: Reflections on Photography*. New York: Macmillan.
Bartmanski, D. 2015. Modes of seeing, or, iconicity as explanatory notion: Cultural research and criticism after the iconic turn in social sciences. *Sociologica*, 9 (1).
Becker, H.S. 2008. *Outsiders*. New York: Simon and Schuster.
Benjamin, W. 1968. *Illuminations: Essays and Reflections,* ed. Hannah Arendt. Schocken Books.
Bielby, W.T., and D.D. Bielby. 1994. "all hits are flukes": Institutionalized Decision Making and the Rhetoric of Network Prime-Time Program Development. *The American Journal of Sociology* 99 (5): 1287–1313.
Domínguez Rubio, F. 2014. Preserving the Unpreservable: Docile and Unruly Objects at MoMA. *Theory and Society* 43 (6): 617–645.
Gibson, J.J. 1979. *The Ecological Approach to Visual Perception*. Hillsdale: Lawrence Erlbaum Associates.
Goodman, J. K., & Paolacci, G. 2017. Crowdsourcing consumer research. Journal of Consumer Research 44 (1):196–210.
Green, A.I. 2013. *Sexual Fields: Toward a Sociology of Collective Sexual Life*. Chicago: University of Chicago Press.
Griswold, W. 1987. The Fabrication of Meaning: Literary Interpretation in the United States, Great Britain, and the West Indies. *The American Journal of Sociology* 92 (5): 1077–1117.
Kemp, M. 2011. *Christ to Coke: How Image Becomes Icon*. Oxford: OUP.
Lizardo, O. 2017. Improving Cultural Analysis: Considering Personal Culture in its Declarative and Nondeclarative Modes. *American Sociological Review* 82 (1): 88–115.
Lizardo, O. 2024. From Macrogenres to microgenres via relationality Working Paper. *Poetics.* 102: 101868.
Martin, J.L. 2015. *Thinking Through Theory*. New York: WW Norton, Incorporated.
Mayer, S. 2020. imagefluency: Image Statistics based on Processing Fluency. R package version 0.2.3. https://CRAN.R-project.org/package=imagefluency.
McDonnell, T.E. 2016. *Best Laid Plans: Cultural Entropy and the Unraveling of AIDS Media Campaigns*. Chicago: University of Chicago Press.
McDonnell, T.E. 2023. Cultural Objects, Material Culture, and Materiality. *Annual Review of Sociology* 49 (1): 195–220.
McDonnell, T.E., A. Jonason, and K. Christoffersen. 2017. Seeing Red and Wearing Pink: Trajectories of Cultural Power in the AIDS and Breast Cancer Ribbons. *Poetics* 60: 1–15.
McDonnell, T.E., D.S. Stoltz, and M.T. Taylor. 2021. Revision, Reclassification, and Refrigerators. *Sociological Forum* 36: 1316–1344.
McDonnell, E.M., D.S. Stoltz, and M.A. Taylor. 2022. Multiple market moralities: identifying distinct patterns in how consumers evaluate the fairness of price changes. *Socio-Economic Review* 20(3) 883–914.
Norman, D.A. 1988. *The Design of Everyday Things*. Basic Books.
Peirce, C.S. 1955. *Philosophical Writings of Peirce*. Dover: Courier Corporation.
Rawlings, C.M., and C. Childress. 2019. Emergent Meanings: Reconciling Dispositional and Situational Accounts of Meaning-Making from Cultural Objects. *American Journal of Sociology* 124 (6): 1763–1809.
Salganik, M.J., P.S. Dodds, and D.J. Watts. 2006. Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science* 311 (5762): 854–856.
Salganik, M.J., and K.E.C. Levy. 2015. Wiki Surveys: Open and Quantifiable Social Data Collection. *PLoS ONE* 10 (5): e0123483.
Salganik, M.J., and D.J. Watts. 2008. Leading the Herd Astray: An Experimental Study of Self-Fulfilling Prophecies in an Artificial Cultural Market. *Social Psychology Quarterly* 74 (4): 338.
Schudson, M. 1989. The Sociology of News Production. *Media Culture & Society* 11 (3): 263–282.
Seguin, C. 2023. The origins of hits: Cumulative advantage vs. multiplicative returns in cultural markets'. *Poetics* 97: 101766.
Smith, P. 2012. Woodstock and Bayreuth: From ritual encounter to iconic event. In *Iconic Power: Materiality and Meaning in Social Life*, ed. J. Alexander, B. Giesen and D. Bartmanski, 171–183. New York. Palgrave Macmillan US.

Sonnevend, J. 2020. A Virus as an Icon: The 2020 Pandemic in Images. *American Journal of Cultural Sociology* 8 (3): 451–461.
Sturken, M. 2015. The Continued Relevance of the Icon: A Comment on the Symposium. *Sociologica*, 9 (1).
Taylor, M.A., D.S. Stoltz, and T.E. McDonnell. 2019. Binding Significance to Form: Cultural Objects, Neural Binding, and Cultural Change. *Poetics*. https://doi.org/10.1016/j.poetic.2019.01.005.
Zubrzycki, G. 2016. *Beheading the Saint: Nationalism, Religion, and Secularism in Quebec*. Chicago: University of Chicago Press.

Terence E. McDonnell is an Associate Professor of Sociology at the University of Notre Dame. McDonnell studies the role of objects in shaping belief and behavior. His research has been published in the American Journal of Sociology, Sociological Theory, Annual Review of Sociology, Social Problems, Sociological Forum, Poetics, Theory and Society, Qualitative Sociology, and more. His book, Best Laid Plans: Cultural Entropy and the Unraveling of AIDS Media Campaigns was published in 2016 with the University of Chicago Press. He is also a co-author of Measuring Culture, Columbia University Press. He received his Ph.D. from Northwestern University.

Marshall A. Taylor is Assistant Professor of Sociology at New Mexico State University. His research revolves around questions of cognition and measurement in the sociology of culture. He is author, with Dustin S. Stoltz, of Mapping Texts: Computational Text Analysis for the Social Sciences (Oxford University Press, 2024). His work has been published in peer-reviewed outlets such as Sociological Theory, Poetics, Political Behavior, Sociological Methods & Research, Journal of Computational Social Science, and others.

Michael Lee Wood is Assistant Professor of Sociology at Brigham Young University. His research focuses on culture, cognition, networks, morality, and religion. His work appears in Sociological Theory, Social Psychology Quarterly, Journal for the Social Scientific Study of Religion, and Socius, among others.