

Reading the Tea Leaves: Ethnographic Prediction as Evidence

CLAIRE MAIERS

WillowTree, Inc.

Those who work in research know that we live in a world that is strongly influenced by what Tricia Wang has called the quantification bias. More so than other forms of information, numbers have incredible formative power. In our culture, numbers are seen as trustworthy representations of reality that are strongly associated with objectivity and untainted by human bias and shortcomings. Recently, data science, big data, algorithms, and machine learning have fueled a new wave of the quantification bias. One of the central fascinations of this wave has been the promise that humans now have the power of prediction at their fingertips. In this paper, I reflect on what it means to make predictions and explore the differences in how predictions are accomplished via quantitative modeling and ethnographic observation. While this is not the first time that ethnographic work has been put in conversation and in contrast with quantified practices, most theorists have framed the role of ethnography as providing context to that quantified work. Here, I argue that ethnographers produce predictions in their own right. I begin by discussing what it means to predict something, focusing on its function. This is followed by a discussion of the ways in which predictions are constructed through both machine learning and ethnographic work. In the course of this discussion I show the commonalities that exist between ethnographic work and machine learning, and I outline methodologies that claim that ethnographic work can make generalizable and accurate statements about the world, including predictive claims. I also point to some of the challenges in using machine learning as a means of producing predictions. This discussion is not meant to discredit these practices, but to demystify the process as a means of loosening quantification's authority, contextualizing its best applications, and putting the two approaches to knowledge production on even footing. Finally, I discuss circumstances in which qualitatively produced predictions may be most valuable, such as when dealing with emerging phenomena and unstable contexts.

INTRODUCTION

As ethnographers in industry, our work is increasingly combined with or compared against the perceived power of big data and data science. Anybody who works in research understands that we live in a world strongly influenced by what Tricia Wang has called the quantification bias (Wang 2016). More so than interpretive work, theoretical concepts, or narrative, numbers have incredible formative power. In our culture, numbers are seen as trustworthy representations of reality (Espeland and Stevens 2008) that are strongly associated with objectivity and untainted by human bias and shortcomings (Daston 1992; Jasanoff 2005). Data science, big data, algorithms, and machine learning not only fit neatly into an epistemological view in which numbers and metrics are seen as taken-for-granted representations of reality (Beer 2016; Espeland and Stevens 2009; Poovey 1998), but they also have fueled a new wave of the quantification bias.

One of the central fascinations of this wave has been the promise that humans now (finally) have the power of prediction at their fingertips. According to the tales told through the public discourse of big data, two key developments have delivered on this promise of science. First, the proliferation of data points provided by the expansion of digital sensors

has gifted us the ability to measure and capture the dynamics of a complex world without human interpretation and distortion. Second, the process of machine learning in general, and unsupervised machine learning in particular, has freed knowledge production from human-generated theories and concepts. Together, the narrative goes, these developments have made prediction a reality and revitalized our value of all things quantified.

To be sure, a great deal of this revitalized enthusiasm for numbers is inspired by material changes in our ability to record and create data, our capacity to store and move data, increased processing power, and greater ease of access to the tools to complete these tasks. But the enthusiasm for big data and prediction that stems from the narrative described above has generally outpaced, or at least out-performed, discussions of the epistemological reality of big data predictions among both the public in general and key decision-makers, such as chief marketing officers, policy makers, or even research directors, in particular.

As others have made clear, this new emphasis on data in recent years has provided both an opportunity to reflect on the distinctive value that qualitative and ethnographic work offers in contrast to data science (Wang 2016) and to draw some lessons on what we as qualitative researchers can learn from the practice of data science (Nafus 2016). In this paper, my intention is to add to this conversation by reflecting on what it means to make predictions and to explore the differences in how predictions are accomplished via quantitative modeling and ethnographic observation.

While this is not the first time that ethnographic work has been put in conversation and in contrast with the new quantified practices associated with datafication (van Dijck 2014), most theorists have framed the role of ethnography as providing context to that quantified work. In this paper, I make a slightly different argument, showing that ethnographers produce predictions in our own right. I begin by discussing what it means to predict something, focusing on its function. This is followed by a discussion of the ways in which predictions are constructed through both machine learning and ethnographic work. In the course of this discussion I elaborate on the methodologies that allow us to claim that ethnographic work can generate generalizable, causal, and accurate statements about the world, including predictive claims. I also discuss some of the shortcomings of using machine learning as a means of producing predictions. This discussion is not meant to discredit these practices, but to demystify the process as a means of loosening quantification's authority, contextualizing its best applications, and putting the two approaches to knowledge production on even footing.

WHAT ARE PREDICTIONS?

In order to make claims about the role that ethnographic work plays in generating predictions, we first need to come to terms with what we mean when we use the word "prediction." I discuss both colloquial and technical definitions and then suggest that we utilize a definition that focuses on the function of predictive claims in practice.

Colloquially, we think of a prediction as a claim about an event or a state that will occur in the future. It is this very general, and yet powerful, conception that most of us rest upon when using the term. Those more versed in statistical practices, machine learning, or data science may have a nuanced definition in mind: a prediction includes an assessment of the likelihood of such states or events actually manifesting. For example, SAS, a company that is in the business of producing predictive analytics, describes prediction as "the use of data,

statistical algorithms and machine learning techniques to identify the likelihood of future outcomes based on historical data” (SAS Institute Inc). This statement of likelihood takes the form of mathematical measurements, such as confidence intervals.

Neither of these definitions serves us well for thinking about the possibility of ethnographic predictions. The colloquial definition says nothing about the origins or production of prediction, while the definition provided by SAS is already infused with the assumptions of statistical sampling methods and inferences. Although references to ethnographic and qualitatively-based predictions can be found in discussions of ethnographic methodology (Burawoy 1998; Small 2009), technical definition for these kinds of predictions are not usually part of these discussions.

However, if we observe the production and application of predictions, we can construct a working understanding. First, the kinds of predictions we talk about with regard to research are linked to empirical data. Second, in applied settings such as marketing agencies, hospitals, consulting firms, or government, predictions are used to support decision-making. At this point, we could say that predictions are empirically-supported claims that are believed to reduce uncertainty about future events or states that are used to buttress decision-making. This definition encompasses a variety of empirical approaches without implying a particular method for generating these claims, and points toward their use in applied settings.

However, careful observation of the problems to which predictions are put and the ways in which they influence decisions allows us to refine this definition further. In my own observations of predictive medical algorithms (Maiers 2017), nurses and doctors used predictive risk scores designed to predict the likelihood of future infection. Despite the theoretical indication of the future onset of infection, these predictions were used to determine if the patient needed antibiotics immediately, suggesting that the prediction factored into the decision-making process by affecting the clinicians’ assessment of the patient’s current condition.

As I discuss in more detail below, using predictive claims to better recognize current states is a common practice with regard to statistically and machine-learning derived predictions. In order to continue to keep the application of machine-learning predictions within our definition, we must alter it slightly: *predictions are empirically-supported claims believed to reduce uncertainty about current states, future events, or future states that are used to buttress decision-making*. This aspect of the definition is not simply an accommodation for the sake of giving ethnographic work some authority over the realm of predictive research. Rather, it is based on both the function and application of claims that are already called “predictions” in the context of big data and machine learning practices. By this definition, ethnographers and qualitative researchers frequently engage in predictive work. We produce descriptions and claims about the world that help our clients and stakeholders make better decisions. These claims may not always be explicitly future oriented.

HOW PREDICTIONS ARE MADE

To show how ethnographers make predictions in their own right, the remainder of this paper will touch on methodological aspects of both ethnography and machine learning. I focus on machine learning because it is an analytical technique associated with big data and because of its growing application as a quantitative approach to generating predictions. After pointing to some of the commonalities between ethnography and big data and

machine learning practices, I outline the process for generating predictions through machine learning. Through this discussion, I hope to arm readers with a better understanding of these practices and the ability to ascertain when and where they work best. I then discuss the scientific status of ethnographic work. Predictions are often seen as belonging to the territory of positive science and seemingly depend upon definitively measured phenomena and the development of covering laws or models. As a result, the interpretive endeavor of ethnographic work may appear to preclude the possibility of prediction. In an effort to show how ethnography can be used in predictive work, I share an alternative framework upon which to base the accuracy of ethnographic claims in general.

Common Ground

When it comes to social data, the processes by which predictions are made in machine learning and big data are similar to the ethnographic process in its basic approach. The value of ethnographic observation is our ability to process and synthesize a complex set of data points and relationships. Similarly, the advantage of new data collection practices and the proliferation of sensors is their ability to capture a wide range of data points. On this front, qualitative and quantitative work are increasingly in conversation as data scientists and ethnographers collaborate and utilize new digital tools in the research process (Rattenbury and Nafus 2018; Anderson, Rattenbury, and Nafus 2009).

Although the kinds of data points created through big data and ethnography take different forms, they often describe something similar, namely behavior in context rather than the lab. This is a relatively new application for statistical inference. When it comes to data related to the social world, much of the data that fueled quantitative analysis of social and behavioral data in the past was sourced through surveys. Its application to behavioral data has been expanded thanks to the many sensors and practices that leave “digital sweat,” or record of human behavior (Gregg 2015). Whether it is social interactions, unintended uses of technology, purchasing patterns, or twitter traffic, both ethnography and big data work with representations of human behavior in context (Golias 2017; Ladner 2014).

From an analysis of this data, both ethnographers and data scientists extrapolate generalizable claims that help us to better understand phenomena and the relationships between phenomena, thereby reducing our uncertainties about the world. To be sure, the process for extrapolating those claims differs a great deal in process. I now want to walk in a little bit more detail through those processes.

How Machine Learning Makes Predictions

The following section, I describe, in very basic terms, the process by which machine learning produces predictions. This description draws upon my work as a sociologist of knowledge, in which I studied the cultural and epistemological dynamics that both promote and result from quantified practices of knowledge production. I used observations of and conversations with data scientists to examine the cultural assumptions about how legitimate knowledge is produced and the ways in which various methods and claims interact with these cultural assumptions. As part of that work, I frequently asked data scientists to describe the process of machine learning. The resulting description is based, in part, on those conversations.

In the simplest terms, machine learning is a process that allows computers to develop methods for making predictions and inferences. The first step is to provide the computer with a data set. This is often called training data. The learning process may be supervised, in which case the computer is given a data set with labeled or classified phenomena, such as a collection of photos of pets that have been labeled as either “cat” or “dog” and told to develop a method for telling those phenomena apart. Note that this requires the human work of assigning labels to phenomena at some point in the data collection process. Or it may be unsupervised, meaning that the computer defines the categories by which data are described. When it comes to pictures of pets, an unsupervised process could result in categories that humans find meaningful, such as brown pets versus spotted pets, but it might also develop categories that are less salient or even noticeable to humans, such as a mathematical relationship between tail length and ear shape. Once there is an algorithm or model for identifying which images are of cats and which are of dogs, the model will be tested. Often it is tested on a subset of the original data set which has been intentionally set aside for these purposes. If the model fails to successfully predict the known outcomes, the model can be adapted and tested again in an iterative process.

Despite this appeal, there are some limitations in this process worth noting. First, even though the resulting model may be great at predicting which pictures are cats and which are dogs in the original training data set and the test data sets, it may not be very good at making similar predictions on new data sets. In other words, the model may not be very generalizable to new settings and contexts. It is impossible to know in advance how the model will perform in truly novel settings that may have slightly different variables at play. When it comes to social data, this is a particularly difficult problem given that social data are endlessly complex and shaped by both macro structures and local contexts. Furthermore, once these models are applied to novel settings, it can be almost impossible to evaluate their accuracy. The only way to know if predictions are correct is to measure them against the real outcomes, and in many cases that may not be possible.

In order to better clarify this problem, let’s consider some cases in which it *is* possible to compare predictions to actual outcomes. The infamous case of the Google Photos algorithm from 2015 that identified and labeled people of color as gorillas is one such instance. This offensive and problematic misrecognition by the algorithm may have stemmed, in part, from a data set trained on photos with an insufficient amount of diversity, bringing attention to issues surrounding bias in data sets and algorithms. It also shows, more generally, how algorithms can fail to be accurate when released from their testing environment on the boarder world. However, we were only aware of the prediction’s flaws because we could see and compare the algorithm’s prediction to our own assessment. Similarly, in my own work with predictive medical algorithms, I watched intensive care unit (ICU) clinicians develop a critical assessment of algorithmic predictions (Maiers 2017). Over time, they were confronted with the corporeal reality of their patients in contrast to the algorithm’s claims. They saw that the algorithm tended to be successful in some cases and less reliable in others, allowing them to build rules of thumb for contextualizing and sometimes discounting these predictions. However, in many cases, predictions will be used to make crucial decisions long before their accuracy beyond a testing environment can be assessed.

Furthermore, the predictions themselves may be “performative,” meaning that the very act of predicting shapes the outcomes that are observed (Callon 1998). This makes it

difficult to know what the outcomes would have occurred in absence of such a prediction. Think, for example, of credit scores which are used to assess the likelihood of someone defaulting on a loan. Given that these scores preclude many individuals from taking out a loan in the first place, the algorithm is shaping the very outcomes which it aims to predict, making it ever more difficult to know if the assessments of one's likelihood to default on a loan was accurate in the first place.

The fact that predictive algorithms work best when applied within the same system or domain in which they were trained and tested leads to a second issue. Predictive algorithms are less well-suited for dealing with emerging phenomena, rare events, and unprecedented events. As the definition from SAS reminds us, machine learning predictions are dependent on historical data. This eliminates the possibility of novel events or factors from being included in the model and greatly reduces the chances that the model will sufficiently account for rare occurrences. In addition, depending on the chosen model and method, rare events may be labeled as "outliers" and intentionally eliminated during the data cleaning process. This means that although machine learning may be great at making predictions with stable systems and conditions, it is more likely to mis-predict outcomes in unstable or changing contexts such as social systems or globalizing markets.

Finally, throughout this section of the paper readers may have noticed that the example I used was not about future states at all, but about estimating the likelihood of current states. Though the language of prediction is used to talk about these processes, the actual results are far from our colloquial definition of predictions: they are not about the future. In fact, the same webpage from which I quoted the technical definition of prediction provides many examples of the kinds of predictions machine learning can provide. The first of these predictions is fraud detection. This is not a prediction about the future at all, but an assessment of the likelihood that certain claims *are* fraudulent. In other words, it is reducing uncertainty about a *current* state. This is also the case with the predictive medical algorithms that I mentioned earlier. These predictions are used to identify patients that *are* developing blood infections. In each of these cases, the "prediction" is not about a future state, but about current states. This is not to say that predictive algorithms are not put to uses that are about the future. Models for setting ticket prices, assigning credit scores, or determining how to stock store shelves are future oriented. The point I want to convey is that in their application and function, machine learning predictions are sometimes about the present.

The point of the previous paragraphs has not been to undermine the legitimacy of machine learning claims. Indeed, data scientists have methods in place to mitigate some of the issues discussed here. My hope instead is to demystify how these algorithms work in order to lay the foundation for claiming that ethnography is also a legitimate way to reduce uncertainty about the future.

Foundations for Making Ethnographic Claims

Most of the work that we do in industry is aimed at reducing uncertainty when making decisions (Dourish and Bell 2014). In reviewing her work on the home in several European countries, Genevieve Bell (2001) describes the job of her team as "understanding people and their daily practices with an eye toward finding new users and uses of technology." Not only does this task suggest that her team's work was aimed at reducing uncertainty, their search for the new suggests a future orientation in which they aim to identify which technologies

and applications might be developed and successfully adopted by consumers. Given this demand for reducing uncertainty, ethnographers have employed novel qualitative methods designed to better derive insights about potential futures (Dourish and Bell 2014, Forlano 2013, Lindley, Sharma and Potts 2014).

But what is the epistemological framework that allows us to claim that our thick descriptions and qualitative inquiries can reduce uncertainty about consumers and how they will behave and react to products? While some ethnographers might take this capability for granted, a suspicion of qualitative and interpretive work is part and parcel of the quantitative bias in our culture. Depending on the home discipline of the ethnographer, she may not have had to question the validity of her sampling, her analytical methods, or her conclusions based on their epistemological legitimacy before entering into industry. Luckily, this is not the case in my home discipline of sociology, where quantitative sociologists sit on the dissertation committees, editorial boards, and grant committees that review and therefore pass judgement on ethnographic and qualitative work.

In the following paragraphs, I draw on the work of qualitative sociologists who have explored the methodological foundations for claiming that ethnographic work is as justly suited for reducing uncertainty as quantitative research. Though there are many issues I could cover here, I focus primarily on the question that colleagues and clients most frequently ask me when I present my work: how can we be sure that our findings are generalizable? This is also closely related to our ability to establish correlations or causal connections between observed phenomena and consumer or user behaviors. By exploring the epistemological foundations of our work, my hope is to bolster our faith in ethnographic claims and to help put the predictions of machine learning and ethnography on even footing.

Mimicking Statistical Inference or Rejecting Generalizable Claims - The nature and status of ethnographic and qualitative work is one that has been greatly debated within the social sciences (Reed 2011; Pugh 2013; Vaisey 2009). Following Geertz (1973), many of us have stressed the value and legitimacy of interpretation as a type of knowledge production. We have seen first-hand how these thick descriptions illuminate everything from the inner workings of high energy physics (Knorr Cetina 1999) to seemingly paradoxical political phenomena (Hochschild 2016). While this function, in and of itself, is sufficient for work that aims to add to the accumulation of knowledge or to provide better understandings of our fellow humans, it is not a sufficient perspective for those aiming to make broad empirical claims or predictions about entire regions, markets, or segments of consumers or users. This is the case for two reasons. First, interpretation often fails to hold authority in a world saturated in quantification bias and the epistemological mental models that accompany this bias. This is due, in part, to the common complaint that qualitative work fails to meet the standards of statistical representativeness and inference. Second, the ways in which qualitative and interpretive work have become associated with methodological programs that emphasize local particularities over generalized patterns or the creation of shared understanding over relational claims makes it difficult to extend the claims of qualitative work beyond its immediate cases and contexts. Both of these can be overcome through an exploration of the methodological approach to qualitative work.

There have been a variety of attempts to remedy this situation. First, we might try to solve this problem by mimicking the assumptions of positivist quantitative work. The idea is that in choosing the right case we can mimic the assumptions of statistical representativeness

upon which many quantitative claims are based. We look for cases and locations that best represent a broader population or we do comparative ethnography as a way to isolate and identify causal relationships and to “control” for confounding variables. This is inevitably problematic. As Small (2009) makes clear, this process often mistakes the concept of representativeness with that of averages. For example, a study of social engagement in a mid-sized town in the mid-west which matches the national average of income or education levels cannot somehow represent this process across most American communities, even though it may statistically resemble the nation as a whole. Furthermore, by intentionally excluding rare or unique cases, we miss out on the opportunity to learn about emergent and developing phenomena or to observe the effects of interactions that may be difficult to observe in average cases.

As an alternative, many of us have been trained in a perspective of interpretive empiricism (as discussed in Reed 2012), in which social knowledge is created through inductive processes that stress the locality of social investigation. Under this epistemological framework, we are not trying to make generalizable, objective claims about covering laws or models at all, but to best explain the social world by articulating the dynamics of the particular. This focus on locality, alongside a resistance to theory and macro social constructs, makes it all but impossible to generalize the findings from one case out to a broader population. It is particularly a problem for industry work in which we hope to use in-depth qualitative analysis of a small sample to inform decisions about entire markets or populations. Is there a way forward in which interpretive and qualitative analysis can make generalized inferential claims without trying to wedge itself into the assumptions of statistical inference?

Finding Our Own Footing - First, we should recognize that statistical inference, and its reliance on large, representative samples, is not the only way to generalize claims. In examining the extended case method (Burawoy 1998), Mario Small (2009) suggests that ethnographers use logical inference instead. This means that the inferences refer to situations rather than populations. As Small explains, in a statistical inference, we hypothesize that populations with a given set of characteristics will display the same set of corresponding characteristics or properties observed in a sample. An example is to say that active adults in the Charlottesville, Virginia are more likely to purchase a gym membership if their household income is over \$60,000. These kinds of claims require some sort of instrument for establishing representativeness and therefore the accuracy of claims. With logical inference, the focus shifts to processes and mechanisms of a situation. We might hypothesize that when offered a free trial at a new gym, the consumer’s decision to book or ignore the offer depends partially on perceptions of cultural fit between the gym and the potential customer. This statement is based on our ability as ethnographers to observe a chronology of events and therefore make causal links between behaviors, decisions, feelings, and events.

This kind of logical inference is particularly good for making what Small calls “ontological statements” or “the discovery of something previously unknown to exist” (2009: 24). In the example I have given, logical inference allows me to make a claim about the relationship between cultural fit and gym membership purchases. This is a great advantage of ethnographic work. Big data does not capture what it does not measure. In addition to the challenges presented in measuring emotions and things like “perception of

cultural fit,” a quantitative study could not take such phenomena into account without someone determining it was a variable worth measuring. All phenomena must be known, at least in the form of measurable data points, ahead of the machine learning process.

Another option offered by Small is to take a different approach to sampling. Rather than look for a representative case, we use “case study logic” to sample. With case study logic, instead of relying on the representativeness of a sample, each additional iteration of investigation brings the researcher closer to an accurate understanding of the area under investigation. As such, this is a sequential process that ends only when the researcher is able to accurately predict the dynamics of the next case and no new phenomena or relationships have emerged. Rather than validating a claim or relationship by statistically showing that our sample would be highly unlikely to contain such a correlation or causal relationship when there is not one in the population, the hypothesis is validated through continual testing that challenges and refines the claim. Interestingly, the iterative nature of case study logic as a means of validation is somewhat similar to the process used in machine learning. Where ethnographers return again and again to the field to test and refine hypotheses, machine learning processes also refine models and algorithms through iterative testing with test data sets. Though the processes may look quite different, it is through repeated exposure to data that both data scientists and ethnographers gain confidence in their conclusions.

So far, I have discussed the ability of ethnography to make accurate and generalizable claims that reach beyond the immediate location of our observations. In the course of this discussion, I have also suggested that we can identify causal connections between phenomena through observation. These are important pieces in understanding why ethnographic work can make predictions. Our work is predictive insofar as it is used to reduce the uncertainty about future states and events, such as changes in markets or the reactions and decisions of users.

As I indicated at the start of this section, ethnographers in industry regularly engage in work that serves the function of prediction. We also use analytical and sampling methods that are similar to those offered by Small. In the next portion of the paper, I talk through an example, pointing out ways in which we might frame the epistemological legitimacy of our predictive work to stakeholders along the way.

APPLICATIONS OF AND ARGUMENTS FOR ETHNOGRAPHIC PREDICTIONS

Ethnographic work is particularly apt at reducing uncertainties for several reasons. As others have pointed out (boyd and Crawford 2012; Seaver 2015; Wang 2016), ethnographic and qualitative work captures different information than that captured by quantitative data sets. Ethnography and qualitative work is particularly advantageous for dealing with the connection between meaning and behavior (Reed 2012), for unearthing meta-feelings and cultural schema (Pugh 2013), and for illuminating subjective experiences (Ladner 2012). These are aspects of the human experience that are difficult to capture and surface in quantitative data sets.

In addition to these well-known advantages, ethnography is well-suited to predicting the emergence and implications of new phenomena. A widely cited example of this is Tricia Wang’s (2016b) ethnographic work on technology usage in China where she observed that low-income individuals were eager to gain access to the technology afforded by smart

phones. Framed slightly differently, I would suggest that her research allowed her to *predict* that low-income consumers in China would purchase affordable smartphones. As such, she advised Nokia to move their business model in that direction. As she describes, her insights fell on deaf ears; Nokia was not convinced by her argument. Their orientation toward quantified metrics and the epistemological models that accompany quantification left them unreceptive to ethnographic data. In particular, she tells us that they resisted her findings because they were not from a large enough sample to suggest representativeness and reliability and that they could not corroborate her findings with their large quantified data sets.

As Wang rightly puts it, “what is measurable isn’t the same as what is valuable.” This is one of the values of ethnographic work. We are able to observe and capture phenomena for which there are no quantified metrics and for which there may not be a name or label. We are also able to predict outcomes that are related to and intertwined with emerging phenomena. Recall that this is one of the challenges of relying upon machine learning for predictive claims; Predictive algorithms cannot account for forces about which they do not know. This means that we do more than provide context for quantified information. Ethnographers can quickly make adaptations to what they are observing and actively generate and test hypotheses about these dynamics during their work to accommodate rare and unprecedented events. In other words, ethnography is well suited to reducing uncertainties in systems and contexts that are changing and unstable. There is an enormous amount of business value in capturing the emerging needs of smartphone users in the changing digital environment of the gig economy, for example.

In addition, ethnographers employ alternative methods for hypothesis generation and validation, thereby drawing insights from contexts and areas of inquiry where large samples and data sets may not be possible or practical. This means that the best qualitative studies do not simply mimic statistical inference, but are conducted according to validation techniques appropriate for qualitative data. With regard to objections to the small sample from Wang’s Nokia study, we might point out that she was not using a statistical sampling logic and that the accuracy and inferential power of her claims are not dependent upon sample size. Instead, she likely used something closer to case study logic to sample, noting that her sample was saturated by consistent and predictable patterns. Instead of identifying relationships between phenomena through statistical inference, Wang may have used logical inference to develop and then verify her hypothesis. She may have seen that in the context of a changing technological and economic landscape, individuals were shifting their financial priorities. To be clear, I do not know the details of how Wang came about her sampling process or her choice of method for conducting analysis. However, many ethnographers and qualitative researchers work with processes that closely resemble the methodologies of case study logic and logical inference.

CONCLUSION

The discussion presented in this paper is meant to give researchers a foundation for arguing that ethnographic and qualitative work are as equally capable of reducing uncertainty surrounding future-oriented decisions as quantitative methods. In doing so, I have not intended to delegitimize research that relies upon machine learning or big data, but to provide a very basic understanding of these practices as a way to demystify them and show

qualitative researchers where there is room and need for qualitative observations. As we work to collaborate with data scientists and computer scientists, qualitative researchers would be well served by deepening their understanding of these processes beyond what I have described here in order to make such partnerships more fruitful.

In addition, I have suggested that there are instances in which ethnography may be particularly well-suited to predictive work. Instances in which we do not yet know what categories and phenomena will be relevant may be less well-served by quantitative work which depends upon known categories and mechanisms being in place to measure such phenomena. Similarly, ethnographic work has a particular advantage in unstable systems and contexts where new phenomena and patterns may be emerging.

Finally, I have shared an alternative framework on which to base the epistemological authority of ethnographic claims. Our stakeholders often evaluate and question our work based on the assumptions and processes of statistical inference. Though these criteria do not apply to ethnographic work, we can base our claims and ability to generalize in logical inference and case study logic. By giving these processes a name and educating our stakeholders in their assumptions and applications, we loosen the hold of the quantitative work and be free to employ the best method or mix of methods for the problem at hand.

Whether or not we should also use the term “prediction” to describe our research that is aimed at reducing uncertainty and supporting decision-making remains an open question. However, my hope is that the discussion provided here has made it clear that we have as much of a case for claiming that we do predictive work as those who use statistical inference or machine learning to produce knowledge.

NOTES

1. I want to emphasize that the definition of prediction that I construct here is not meant to be normative or to precisely represent formal definitions of prediction that might be found in mathematical or philosophical treatise. Entire books could and have been written on the topic. Instead, as an observer of knowledge practices, this definition is based on what people actually do when they make a prediction, the expectations they have of this construct, and the uses to which it is put.

2. Small takes the idea of case study logic from Yin (2002) and adapts it for interview-based research.

REFERENCES CITED

- Anderson, Ken, Dawn Nafus, Tye Rattenbury, and Ryan Aipperspach
2009 “Numbers Have Qualities Too: Experiences with Ethno-Mining.” *Epic*. 123-140.
- boyd, danah and Kate Crawford
2012 “Critical Questions for Big Data.” *Information, Communication & Society* 15(5): 662–79.
- Burawoy, Michael
1998 “The Extended Case Method.” *Sociological Theory* 16 (1): 4-33.
- Callon, Michel, editor
1998 *The Laws of the Markets*. London: Blackwell Publishers.
- Daston, Lorraine
1992 “Objectivity and the Escape from Perspective.” *Social Studies of Science* 22(4): 597–618.

- Dourish, Paul and Genevieve Bell
 2014 “Resistance is Futile”: Reading Science Fiction Alongside Ubiquitous Computing.” *Personal and Ubiquitous Computing* 18 (4): 769–778.
- Espeland, Wendy Nelson and Mitchell L. Stevens
 2008 “A Sociology of Quantification.” *European Journal of Sociology* 49 (3): 401–36.
- Forlano, Laura
 2013 “Ethnographies from the Future: What can ethnographers learn from science fiction and speculative design?” *Ethnography Matters*. Accessed July 30, 2018.
<http://ethnographymatters.net/blog/2013/09/26/ethnographies-from-the-future-what-can-ethnographers-learn-from-science-fiction-and-speculative-design/>
- Geertz, Clifford
 1973 *The Interpretation of Cultures: Selected Essays*. New York: Basic Books.
- Gregg, Melissa
 2015 “Inside the Data Spectacle.” *Television & New Media* 16 (1): 37–51.
- Golias, Christopher A.
 2017 “The Ethnographer’s Spyglass: Insights and Distortions from Remote Usability Testing.” *Ethnographic Praxis in Industry Conference Proceedings*, 247-261.
- Hochschild, Arlie Russell
 2016 *Strangers in Their Own Land: Anger and Mourning on the American Right*. New York: The New Press.
- Jasanoff, Sheila
 2005 *Designs on Nature: Science and Democracy in Europe and the United States*. Princeton, NJ: Princeton University Press.
- Knorr Cetina, Karin
 1999 *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge and London: Harvard University Press.
- Ladner, Sam
 2012 “Ethnographic Temporality: Using time-based data in product renewal.” *EPIC 2012 Proceedings*, 30-38. Accessed June 21, 2018. <https://anthrosource.onlinelibrary.wiley.com/doi/pdf/10.1111/j.1559-8918.2012.00005.x>
- 2014 *Practical Ethnography: A Guide to Doing Ethnography in the Private Sector*. Walnut Cree: Left Coast Press.
- Lindley, Joseph, Dhruv Sharma, and Robert Potts
 2014 “Anticipatory Ethnography: Design Fiction as an Input to Design Ethnography.” *Ethnographic Praxis in Industry Conference Proceedings*, 237-253.
- Maiers, Claire
 2017 “Analytics in Action: Users and Predictive Data in the Neonatal Intensive Care Unit.” *Information, Communication & Society* 20: (6): 915-929.
- Nafus, Dawn
 2016 “The Domestication of Data: Why Embracing Digital Data Means Embracing Bigger Questions.” *Ethnographic Praxis in Industry Conference Proceedings*, 384–399.
- Nafus, Dawn and Jamie Sherman
 2014 “Big Data, Big Questions | This One Does Not Go Up To 11: The Quantified Self Movement as an Alternative Big Data Practice.” *International Journal of Communication* 8(0):1784–94.

- Poovey, Mary
1998 *A History of the Modern Fact: Problems of Knowledge in the Sciences of Wealth and Society*. 1st edition. Chicago: University Of Chicago Press.
- Pugh, Allison J.
2013 “What Good Are Interviews for Thinking about Culture? Demystifying Interpretive Analysis.” *American Journal of Cultural Sociology* 1(1):42–68.
- Rattenbury, Tye and Dawn Nafus
2018 “Data Science and Ethnography: What’s Our Common Ground, and Why Does It Matter?” May 7, 2018. Accessed June 1, 2018. <https://www.epicpeople.org/data-science-and-ethnography/>
- Reed, Isaac Arial
2011 *Interpretation and Social Knowledge: On the Use of Theory in the Human Sciences*. Chicago; London: University of Chicago Press.
2012 “Cultural Sociology as a Research Program: Post-positivism, Meaning, and Causality.” *The Oxford Handbook of Cultural Sociology*. Jeffrey C. Alexander, Ronald N. Jacobs, and Philip Smith (eds). Oxford and New York: Oxford University Press.
- SAS Institute Inc.
“Predictive Analytics: What it is and Why it Matters.” Sas.com. Accessed June 4th, 2018. https://www.sas.com/en_us/insights/analytics/predictive-analytics.html
- Seaver, Nick
2015 “The Nice Thing about Context Is That Everyone Has It.” *Media, Culture & Society* 37(7): 1101–9.
- Small, Mario
2009 “‘How Many Cases do I Need?’: On Science and the Logic of Case Selection in Field-based Research.” *Ethnography* 10 (5): 5-37.
- van Dijck, J.
2014 “Datafication, Dataism and Dataveillance: Big Data Between Scientific Paradigm and Ideology.” *Surveillance & Society* 12(2): 197-208.
- Vaisey, Stephen
2009 “Motivation and Justification: A Dual-Process Model of Culture in Action.” *American Journal of Sociology* 114(6): 1675–1715.
- Wang, Tricia
2016 “The Human Insights Missing from Big Data.” TEDxCambridge. Accessed June 3rd, 2018. https://www.ted.com/talks/tricia_wang_the_human_insights_missing_from_big_data
- Wang, Tricia
2016b “Why Big Data Needs Thick Data.” *Medium*. Accessed June 2nd, 2018 <https://medium.com/ethnography-matters/why-big-data-needs-thick-data-b4b3e75e3d7>
- Yin, R.
2002 *Case Study Research*. Thousand Oaks, CA: Sage.