

SYMPOSIA PAPER

Organizations and Values in Science and Technology

Justin B. Biddle

School of Public Policy, Georgia Institute of Technology, Atlanta, GA 30332, USA
Email: justin.biddle@pubpolicy.gatech.edu

(Received 19 April 2023; revised 01 October 2023; accepted 05 October 2023; first published online 20 October 2023)

Abstract

This paper articulates a conceptual framework for examining philosophical issues such as the role of values in science at an organizational level. It distinguishes between three dimensions of organizations – organizational aims, organizational structure, and organizational culture – and it examines how these dimensions relate to values in research and development, with a focus on machine learning systems for predictive policing. This framework can be fruitful in identifying interesting and understudied philosophical problems – including those involving inter-organizational divisions of labor – that might otherwise be difficult to conceptualize.

1. Introduction

There are many issues in science and technology of epistemological and ethical significance that are helpfully addressed at the level of the organization. Questions concerning the role of values, for example, can be examined at an organizational level (e.g., that of a research institute, lab, or technology firm), and doing so can generate fruitful philosophical insights. Despite this, organization-level analyses have not been emphasized by philosophers of science. This article articulates a conceptual framework for examining the role of values in science and technology at an organizational level. I distinguish between three dimensions of organizations (organizational aims, organizational structure, and organizational culture), examine briefly how these dimensions relate to values in science and technology, and draw upon these dimensions to highlight some of the philosophical work that this framework can do. For example, issues surrounding data reuse and “data journeys” are being scrutinized by philosophers of science (e.g., Leonelli and Tempini 2020); organizational perspectives can help to identify additional problems in this area, including those that involve interorganizational divisions of labor.

Following a discussion in section 2 of the terminology of “organizations” and “values,” section 3 articulates the conceptual framework introduced in the preceding text and briefly highlights some respects in which the dimensions of organizations

can impact values in science and technology. Section 4 introduces a case that illustrates the importance of interorganizational divisions of labor in research and development (R&D)—namely, machine-learning (ML) predictive-policing systems that generate predictions influencing arrest, incarceration, and penal sentencing. This section serves as an illustration of the fruitfulness of the conceptual framework articulated earlier and as a conceptual resource for data provenance projects, such as “datasheets for data sets” (Gebru et al. 2021), providing further justification for the inclusion of questions about organizational features and alignment.

2. Organizations and values

I use the term “organizations” to mean “social structures created by individuals to support the collaborative pursuit of specified goals” (Scott and Davis 2007, 11). In this article, I’m concerned with scientific and technological organizations such as (again) research institutes, laboratories, or technology firms. The concepts of organizations and institutions can overlap, but I use the term “organization” to distinguish it from conceptions of institutions as general, established patterns of behavior such as human language, patriarchy, and money (c.f. Elliott 2023 and Fernandez-Pinto 2023).

My usage of the term “values” is indebted to work by Elizabeth Anderson (1993). I define values in terms of evaluative standards—the values of an individual are the standards that it employs for evaluating other individuals, actions, or things (where an individual might be a person, an organization, or some other entity) (Biddle 2023). If a researcher chooses one theory over another because it is simpler, then simplicity is operating as a value. If a designer criticizes an algorithm for being unfair to a demographic group, then fairness is operating as a value.

There are other legitimate ways to define values (c.f. Brown 2020), but I believe that the conception of values as evaluative standards is a useful one, and it is consistent with other prominent treatments of values in philosophy of science. Kuhn (1977), McMullin (1983), and others have theorized values in terms of evaluative standards; my treatment is similar, except that I do not assume (or believe) that “epistemic” and “nonepistemic” values can be sharply distinguished.

There are two implications of this conception of values that are worth highlighting here. The first is that values are not necessarily consciously held or endorsed. Individuals might evaluate things according to criteria that they do not endorse or about which they are unaware; unconscious biases, for example, can operate as values. The second has already been mentioned—namely that organizations, as well individual persons (and perhaps other entities), can have values. For example, if a university admissions department evaluates students according to standardized test scores, then scoring well on standardized tests is an organizational value. This is so, even if some or even most of the individuals who work in the admissions office do not value this personally.

Analyses of values in science and technology can occur at multiple levels. Many focus on the level of the individual researcher. These analyses might, for example, attempt to identify ethical norms that individual researchers should apply in their R&D activities. At the opposite end of the spectrum, analyses might occur at a societal level—for example, by examining the impacts of ideologies such as neoliberalism or White supremacy on research or technology. Individual-level and societal-level

analyses are both important and necessary. The framework articulated in this article represents a mid-level organizational perspective—which mediates between the individual and the societal—and I attempt to show that it can provide fruitful resources for philosophical investigation.

3. A conceptual framework for organizations and values

The importance of organizational features for research and practice deserves additional attention by philosophers of science. Longino (2002) has proposed norms that “communities” should satisfy if they are to achieve epistemic goals such as objectivity or knowledge—namely, public venues for criticism, uptake of criticism, public standards, and a tempered equality of intellectual authority. But there is still much to be learned about the relationships between organizations, values, and research and technological outputs. To understand these relationships more systematically, I distinguish between three dimensions of organizations: organizational aims, organizational structure, and organizational culture. There are other ways in which one might distinguish between different organizational features; I hope to show that this is a useful way.¹

By *organizational aims* (or goals), I do not mean the motives, preferences, or goals of the individuals making up those organizations. Organizations act in ways that are not necessarily reducible to individual-level characteristics (e.g., Mayo-Wilson et al. 2011). I also do not mean what organizations assert about their goals. Organizational aims are much more closely related to what organizations do—what criteria they use in making decisions.

Organizational aims can be understood in terms of “criteria for generating and selecting among alternative courses of action” (Scott and Davis 2007, 185; see also Simon 1964, 1). If an organization uses profitability as a significant constraint on decision making, then profitability is an organizational aim. If an organization uses profitability and social responsibility (understood in some specific sense) as significant constraints on decision making, then these are both organizational aims. These can be organizational aims, even if they do not personally motivate the individuals making up those organizations. Additionally, if organizational aims are understood in terms of constraints on decision making, they might bear little relation to reported organizational aims. For example, if considerations of social responsibility do not constrain a corporation’s behavior, then social responsibility is not organizational aim—whatever its advertising campaign might say.

Organizations make decisions by dividing labor into different activities and coordinating them. These two dimensions—division of labor and coordination—make up the *structure* of an organization. “The structure of an organization can be defined simply as the sum total of the ways in which it divides labor into distinct tasks and then achieves coordination among them” (Mintzberg, quoted in Maguire 2003, 11). To analyze this concept further, some theorists have broken down the complexity of an organization’s division of labor into functional differentiation (e.g., which tasks are undertaken in a research lab, how those tasks are divided), spatial differentiation

¹ Sources that have influenced and informed this conceptualization include Maguire (2003), Scott and Davis (2007), Walsh and Lee (2015), and Walsh et al. (2019). I’m also indebted to John Walsh for conversations on this topic.

(e.g., whether a research group's activities are done in the same room or in different cities or countries), and vertical differentiation (e.g., how many levels of hierarchy there are between the highest and lowest levels in a lab, and how many people occupy these various levels) (Maguire 2003). Coordination and control, moreover, can be analyzed into administration (e.g., the degree to which a research lab employs administrative assistants and public relations officers), formalization (e.g., the extent to which decision-making procedures are formalized in rules and regulations), and centralization (e.g., the degree to which decision-making power is concentrated in an individual person or small group of leaders) (ibid.). Group size is an important variable in explaining organizational structure; for example, larger organizations might require more complex structure and mechanisms of control than smaller ones (ibid.; Walsh and Maloney 2007).

Some of Longino's norms of ideal scientific communities are norms about organizational structure. The first—"publicly recognized forums for the criticism of evidence, of methods, and of assumptions and reasoning" (Longino 2002, 129)—is a requirement that ideal scientific communities be structured to have venues for deliberation that are public. Ideal communities should not have only back-channel forums that are hidden or inaccessible to some. To state this in terms of division of labor and coordination, communities should have publicly recognized mechanisms for the coordination of research communication—and in some cases, these mechanisms might impact the division of labor by requiring that individuals or offices be responsible for overseeing these mechanisms. Another of Longino's norms requires that communities be structured so that research is evaluated according to standards that are public. These standards, according to Longino, might be held "explicitly or implicitly" (ibid.). I take it that they might also be formal or informal. For example, a regulatory organization might allow a technology to be evaluated according to health safety criteria but not socioeconomic impact. This standard might be formalized in rules and policies—as in U.S. regulatory agencies such as the FDA and EPA—or it might be informal, as in the case of scientific labs that, as a matter of practice, do not concern themselves with socioeconomic impact.

Organizational structure is related to, but distinguishable from, organizational aims. If an organization is acting rationally, then it will be structured in such a way as to promote its aims. For example, if a firm aims to develop and license ML systems in a socially responsible manner, it might create an office of ethics that significantly influences organizational behavior. Relatedly, if a firm advertises that it has social responsibility as an aim, and if it is structured so that its office of ethics has no real power—or if it shuts down that office altogether—then this might indicate that it does not aim at social responsibility. The influence can also go in the other direction; changes in organizational structure (e.g., the creation of an ombuds office due to grassroots activism by employees) might lead to changes in organizational aims (e.g., promotion of equity).

The third dimension of organizations is *organizational culture*, by which I mean the norms, practices, habits, and assumptions that are commonly held in an organization (e.g., Armacost 2004, 493). The concept of organizational culture overlaps significantly with aims and structure, but I include it as a distinct dimension because it is possible for different organizations to have the same aims and structures with different cultures. For example, two organizations might have the same aims and structures

but have different cultural norms about who may speak up or give their honest opinion without fear of negligent misinterpretation or retribution. Similarly, two organizations might have the same aims and structures but display different levels of toleration of bad behavior or ethical violations (e.g., *ibid.*, 494).

Longino's norm of a tempered equality of intellectual authority is primarily a cultural norm. Its requirements include "that every member of the community be regarded as capable of contributing to its constructive and critical dialogue" (*ibid.*, 131–32). This norm does not specify organizational aims or decision-making processes; it is rather a norm about who should be taken seriously in critical discursive interactions, and how seriously they should be taken. In other words, it is a norm about practices, habits, and expectations (i.e., culture). The last of Longino's norms—uptake of criticism—might also be viewed as a cultural norm, though particular structures might be more or less effective in fostering it.

Each of these dimensions of organizations relates to values. Organizational aims are values—where values, again, are understood as criteria for evaluation. Moreover, organizational aims can impact the values that are reflected in the scientific and technological outputs of those organizations (c.f. Elliott and McKaughan 2014; Intemann 2015). For example, whether financial profit is an aim and, if so, whether there are additional aims such as social responsibility can impact the framing and design of a research project or technological system, as well as how epistemic risks are managed (Biddle and Kukla 2017). Organizational structure and culture can be similarly impactful. For example, the division of labor within a research organization can impact the coordination of how epistemic risks are managed across that organization (Walsh et al. 2019; Winsberg et al. 2014), and how inclusive a culture is can impact the values brought to bear on problem framing and model appraisal (Longino 2002).

Organizations have dimensions that either function as or impact values, and the values of organizations are distinguishable from the values of the individuals making up those organizations. Because of this, if we wish to affect the values embedded in research or technology, we might look to these dimensions for points of intervention. Moreover, an awareness of these dimensions, and their relation to values, can help to illuminate issues of epistemological and ethical import. The next section will provide an example of this by examining the development of predictive-policing systems from an organizational point of view.

4. Predictive-policing systems, division of labor, and data reuse

ML systems are increasingly used in police departments and judicial systems in many countries. Predictive policing systems forecast criminal activity and are used to allocate police resources. Recidivism-prediction systems assess the risk that an "offender" will "reoffend," and they are used by many judges to influence penal sentencing decisions. The design, use, and regulation of these systems have generated significant controversy, especially in the United States (Angwin et al. 2016; Biddle 2020; Richardson et al. 2019).

While a discussion of the internal workings of these systems is beyond the scope of this article, it is crucial to emphasize the significance of how concepts such as "crime," "offender," and "reoffender" are operationalized. While there are many ways

in which this might be done—for example, one might operationalize “crime” in terms of conviction by a jury or settlement with admission of guilt—it is common to do so in terms of arrest (Angwin et al. 2016; Biddle 2020). Thus, systems that are supposed to predict where and when “crimes” will likely occur are trained on arrest data, which reflects not only criminal activity (very imperfectly) but also decisions by police departments about which neighborhoods should be subject to strict monitoring and enforcement. Similarly, many recidivism-prediction systems are trained on arrest data, with the consequence that a “recidivist” is operationalized as someone who has been arrested and then rearrested—whether or not they have committed a crime. The controversial recidivism-prediction system COMPAS, for example, assesses the risk that someone who has been arrested will be arrested again within two years (*ibid.*).

Decisions about which data to use to train ML systems involve epistemic risk (Biddle 2020), and concerns about bias in training data have justifiably received significant attention. Across the United States, Black, Indigenous, and People of Color (BIPOC) are disproportionately targeted and subject to arrest (Alexander 2012). These decisions lead to biased arrest data, and when these data are used to train predictive-policing systems, these biases become encoded into them. Richardson et al. (2019) highlight this problem in their study of “dirty data” used to train predictive policing systems. They show, for example, how the Chicago Police Department developed its Strategic Subject List based on data sets that reflect a pattern of unlawful stop and frisk practices that disproportionately harmed Black residents.

One might think that a “solution” to this problem would be to eliminate, or at least reduce, racial bias in police departments. This could reduce racial bias in arrest data, which could (it might be argued) lead to the production of more neutral—or at least less biased—predictive-policing systems. Reducing racial bias in police departments is a laudable goal that should be pursued. What is less clear is the effect that this would have on predictive-policing systems. Examining this case from an organizational perspective suggests that the process of designing these systems might still be fraught with epistemological and ethical problems, even if we could significantly reduce racial bias in police departments.

The design and development of predictive-policing systems is organizationally complex. Some of the complexities are evident by comparing two very different types of organizations that are intimately involved in the creation of these technologies. The first and most obvious is the ML organization. Predictive-policing systems are created (in part) by teams of computer scientists, engineers, social scientists, statisticians, and others who are highly trained in technical fields. There is, of course, variation in the ML organizations that produce these systems. Some are for-profit firms that develop the systems and then license them for use by states and local governments. Equivant, for example, is a private, for-profit company that created and licenses COMPAS. Academic research labs in universities also produce these technologies. The aims, structures, and cultures that characterize these organizations are diverse—but despite this diversity, there are commonalities. Most of them share the epistemic aim of developing models that fit their data. For example, many data analytics firms that develop recidivism-prediction systems aim to develop models that are accurate with respect to arrest data.

There is another, very different type of organization that is heavily involved in the creation of predictive-policing systems. Local police departments produce the arrest data on which these systems are trained. For example, the Broward County Sheriff's Office in Florida generated the data on which the COMPAS algorithm was trained (Angwin et al. 2016). Local police departments vary, just as ML organizations do—particularly in terms of their organizational structure (Maguire 2003). But despite this variation, there are organizational features that are shared by most police departments, such as similarities in organizational culture (Armacost 2004). Furthermore—and importantly for the argument of this article—there are significant differences between the organizational features of police departments and those of ML organizations, and these differences impact the ways in which these organizations manage epistemic risks. They impact, for example, which epistemic aims are prioritized and which are deemphasized, and which epistemic failures or inadequacies are tolerated and to what degree.

Consider, for example, the epistemic aim of empirical accuracy. Most ML organizations, again, share the epistemic aim of developing models that are accurate with respect to their data. They might or might not be concerned about the provenance of these data and how biases in these data impact model performance over different demographic groups. Police departments also aim for empirical accuracy—in rather different respects. When police departments share data about arrests with city governments (e.g., that x number of felony arrests were made), they typically attempt to ensure that these data are accurate (e.g., that x number of felony arrests were made). Whether arrest data is an accurate measure of criminality, however, is another question—and while many police might assume that it is an accurate measure, police departments do not aim to validate this empirically.

For a police department to aim to ensure that arrest is an adequate measure of criminality, it should attempt to avoid both Type I and Type II errors (arresting someone for a crime that they did not commit, and failing to arrest someone for a crime that they did commit). It is typically not possible to minimize both types of error—the inductive risk literature has shown there is typically no way to minimize both simultaneously and to balance the risk of both in a neutral way (e.g., Douglas 2000; Wilholt 2009). Because police departments make little to no effort to prevent at least one of these types of error, then (I suggest) they cannot be said to aim to ensure that arrest is an accurate measure of criminality.

Many have argued that police departments in the United States aim to perpetuate the subjugation of marginalized groups, especially Black Americans (e.g., Alexander 2012). On this account, police departments do not aim to prevent either Type I or Type II errors—their aim is to protect the interests of privileged groups, not to collect empirically accurate data. However, even if we reject this account and assume that police departments aim to protect and serve their communities—and thus attempt to avoid arresting innocent people—it is still difficult to argue that they aim to control Type II errors. For many reasons, police departments cannot arrest everyone who commits a crime. Moreover, for most crimes, police departments need not and do not aim to do so. While some states have laws requiring police officers to make arrests under certain conditions for certain crimes (e.g., mandatory arrest laws for domestic violence), police officers in most cases may exercise their discretion in deciding whether to make an arrest (Goldstein 1963; Huff 2021). Furthermore, even if police

officers were required to make an arrest in case a crime were reported, they are not required to search for all criminal activities, and police departments must, given resource constraints, make judgments about which crimes to prioritize. Police departments may decide that some crimes—for example, white-collar financial crimes—are simply not priorities. In doing this, they are deciding that Type II errors with respect to these crimes, at whatever rates they are occurring, are acceptable.

In some cases, we might even think it praiseworthy to refrain from arresting those who have broken the law. For example, we might think it laudable for police officers to follow a practice of refraining from arresting youth from marginalized communities who have committed only minor offences, if their arrest would result in detention in violent juvenile facilities. Such a practice would systematically increase the rate of Type II errors—and despite this, we might have strong social and ethical reasons for believing that it should be pursued. Police departments, whether they are systemically racist or not, do not aim for “full enforcement” or anything close to it (Goldstein 1963), and as a result, they do not aim to ensure that arrest is an accurate measure of criminality. Because of this, there might be problems associated with the use of arrest data to train predictive-policing systems, even if police departments were to be reformed and cleansed of racism.

From this discussion, it is evident that the organizational divisions of labor involved in the production of predictive-policing systems are especially complex. Not only do the various organizations involved—research firms, academic labs, police departments—have their own aims, cultures, and structures, including respective intraorganizational divisions of labor but also there is a stark *interorganizational division of labor*. ML organizations develop tools based on data collected from police departments, which have radically different organizational features than the ML organizations that reuse these data. Interorganizational divisions of labor can pose interesting and challenging problems that are epistemically and ethically significant. While a thorough treatment of these issues is beyond the scope of this article, I will briefly highlight two conclusions that are suggested by this case.

First, interorganizational divisions of labor can involve misalignment of epistemic aims, values, and cultures, as well as misalignment of how epistemic risks are managed, which can pose epistemic and ethical challenges for data reuse. While police departments might assume that arrest is an adequate measure of criminality, they do not aim to validate this measure empirically. Because of this, the use of arrest data as a measure of criminality is particularly fraught with epistemic risk. For organizations with different epistemic aims—especially epistemic aims that are incompatible with those of the data-producing organization, such as aims that require validated measures of criminality—the reuse of arrest data is both epistemically and ethically risky. For example, ML organizations that aim to reduce racial bias in criminal sentencing should avoid, or be extremely cautious about, reusing arrest data to train their ML systems.

Second, epistemic risks are organizationally transmissible. As a result, even in cases of organizational alignment, epistemic and ethical problems can arise, including positive feedback loops that compound ethical harms. While some ML organizations are concerned with data provenance and exercise appropriate caution in reusing arrest data (e.g., Rodolfa et al. 2020), others aim to develop models that are accurate with respect to their data—and exhibit a lack of concern about the provenance of

these data. The data produced by an organization reflects the values of that organization, which in turn can be transmitted to organizations that reuse these data—whether or not they endorse these values. In the case of predictive-policing systems, which impact who is more or less likely to be arrested or incarcerated, the organizational transmissibility of epistemic risk can lead to the unintentional compounding of ethical harms and social injustice to both current and future generations.

5. Conclusion

In this article, I have articulated a conceptual framework for examining philosophical issues such as the role of values in science and technology at an organizational level, and I have argued that this framework can be fruitful in identifying interesting philosophical problems—including those involving interorganizational divisions of labor—that might otherwise be difficult to conceptualize. This framework can aid analysts in critiquing and improving the values embedded in science and technology, and it highlights the importance of data provenance efforts to document, and facilitate critical reflection on, the origins of data used to produce ML systems, including questions about organizational aims and alignment.

Acknowledgments. Early versions of this paper were presented at the University of Minnesota, the University of Toronto, PSA 2022, SPSP 2022, VMST 2022, and the Georgia Tech Philosophy Club. I'm grateful to the participants of those events and especially to Joyce Havstad, Kevin Elliott, and John Walsh for reading and providing valuable feedback on early drafts of the paper.

References

- Alexander, Michelle. 2012. *The New Jim Crow*. New York: The New Press.
- Anderson, Elizabeth. 1993. *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press.
- Angwin, Julia, Jeff Larson, Surya Mattu, and Lauren Kirchner. 2016. "Machine Bias." *ProPublica*. May 23. <https://www.propublica.org/article/machine-bias-risk-assessments-in-criminal-sentencing>
- Armocost, Barbara. 2004. "Organizational Culture and Police Misconduct." *George Washington Law Review* 72 (3):453–546.
- Biddle, Justin. 2020. "On Predicting Recidivism: Epistemic Risk, Tradeoffs, and Values in Machine Learning." *Canadian Journal of Philosophy* 52 (3):321–41. <http://doi.org/10.1017/can.2020.27>
- Biddle, Justin. 2023. "Values in AI Systems." In *Technology Ethics: A Philosophical Introduction and Readings*, edited by G. Robson and J. Tsou, 132–40. New York: Routledge.
- Biddle, Justin, and Quill R. Kukla. 2017. "The Geography of Epistemic Risk." In *Exploring Inductive Risk*, edited by K. Elliott and T. Richards, 215–37. Oxford: Oxford University Press.
- Brown, Matthew. 2020. *Science and Moral Imagination*. Pittsburgh: University of Pittsburgh Press.
- Douglas, Heather. 2000. "Inductive Risk and Values in Science." *Philosophy of Science* 67 (4): 559–79. <https://doi.org/10.1086/392855>
- Elliott, Kevin. 2023. "A Framework for Studying Institutional Roles in Science and Values." Presentation at PSA 2022.
- Elliott, Kevin, and Daniel McKaughan. 2014. "Nonepistemic Values and the Multiple Goals of Science." *Philosophy of Science* 81 (1):1–21. <https://doi.org/10.1086/674345>
- Fernández Pinto, Manuela. 2023. "Pandemic Science and Commercial Values: An Institutional Account of Values in Science." *Philosophy of Science*. Accepted manuscript. <https://doi.org/10.1017/psa.2023.125>
- Geburu, Timnit, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, Kate Crawford. 2021. "Datasheets for Datasets." *Communications of the ACM* 64:84–92. <https://doi.org/10.1145/3458723>
- Goldstein, Herman. 1963. "Police Discretion: The Ideal Versus the Real." *Public Administration Review* 23 (3): 140–48.

- Huff, Jessica. 2021. "Understanding Police Decisions to Arrest: The Impact of Situational, Officer, and Neighborhood Characteristics on Police Discretion." *Journal of Criminal Justice* 75:101829. <https://doi.org/10.1016/j.jcrimjus.2021.101829>
- Intemann, Kristen. 2015. "Distinguishing between Legitimate and Illegitimate Values in Climate Modeling." *European Journal for Philosophy of Science* 5 (2):217–32. <https://doi.org/10.1007/s13194-014-0105-6>
- Kuhn, Thomas. 1977. Objectivity, Value Judgment, and Theory Choice. In *The Essential Tension: Selected Studies in Scientific Tradition and Change*, 320–39. Chicago: University of Chicago Press.
- Leonelli, Sabina and Niccolo Tempini, eds. 2020. *Data Journeys in the Sciences*. Cham, Switzerland: Springer.
- Longino, Helen. 2002. *The Fate of Knowledge*. Princeton, NJ: Princeton University Press.
- Maguire, Edward. 2003. *Organizational Structure in American Police Agencies: Context, Complexity, and Control*. Albany, NY: SUNY Press.
- Mayo-Wilson, Conor, Kevin Zollman, and David Danks. 2011. "The Independence Thesis: When Individual and Social Epistemology Diverge." *Philosophy of Science* 78:653–77. <https://doi.org/10.1086/661777>
- McMullin, Ernan. 1983. "Values in Science." In *PSA 1982, Vol. 2*, edited by P. D. Asquith and T. Nickles, 3–28. East Lansing, MI: Philosophy of Science Association.
- Richardson, Rashida, Jason Schultz, and Kate Crawford. 2019. "Dirty Data, Bad Predictions: How Civil Rights Violations Impact Police Data, Predictive Policing Systems, and Justice." *New York University Law Review* 94:192–233.
- Rodolfa, Kit, Erika Salomon, Lauren Haynes, Ivan Higuera Mendieta, Jamie Larson, Rayid Ghani. 2020. "Case Study: Predictive Fairness to Reduce Misdemeanor Recidivism Through Social Service Interventions." In Conference on Fairness, Accountability, and Transparency (FAT* '20), January 27–30, 2020, Barcelona, Spain. ACM, New York. <https://doi.org/10.1145/3351095.3372863>
- Scott, W. Richard, and Gerald F. Davis. 2007. *Organizations and Organizing*. Upper Saddle River, NJ: Pearson Prentice Hall.
- Simon, Herbert. A. 1964. "On the Concept of Organizational Goal." *Administrative Science Quarterly* 9 (1):1–22.
- Walsh, John P., and You-Na Lee. 2015. "The Bureaucratization of Science." *Research Policy* 44:1584–1600. <https://doi.org/10.1016/j.respol.2015.04.010>
- Walsh, John P., and Nancy Maloney. 2007. "Collaboration Structure, Communication Media, and Problems in Scientific Work Teams." *Journal of Computer-Mediated Communication* 12:712–32. <https://doi.org/10.1111/j.1083-6101.2007.00346.x>
- Walsh, John, You-Na Lee, and Li Tang. 2019. "Pathogenic Organization in Science: Division of Labor and Retractions." *Research Policy* 48:444–61. <https://doi.org/10.1016/j.respol.2018.09.004>
- Wilholt, Torsten. 2009. "Bias and Values in Scientific Research." *Studies in History and Philosophy of Science* 40:92–101. <https://doi.org/10.1016/j.shpsa.2008.12.005>
- Winsberg, Eric, Bryce Huebner, and Quill R. Kukla. 2014. "Accountability and Values in Radically Collaborative Research." *Studies in History and Philosophy of Science Part A* 46:16–23. <https://doi.org/10.1016/j.shpsa.2013.11.007>