

Governing Algorithms: A Provocation Piece

Solon Barocas
sib237@nyu.edu

Sophie Hood
sah7@nyu.edu

Malte Ziewitz
mcz3@nyu.edu

New York University
March 29, 2013

*Paper prepared for the “Governing Algorithms” conference,
May 16-17, 2013, at New York University*

1. Algorithms have developed into somewhat of a modern myth. They “compet[e] for our living rooms” (Slavin 2011), “determine how a billion plus people get where they’re going” (McGee 2011), “have already written symphonies as moving as those composed by Beethoven” (Steiner 2012), “allow self-determined action on the Internet but also contain aspects of control over this action” (Institute for Media Archeology 2011), and “free us from sorting through multitudes of irrelevant results” (Spring 2011). How do they do all this, and more? What exactly are algorithms “doing,” and what are the implications? Can an algorithm “do” anything? And who or what are “algorithms” anyway?

2. This provocation piece addresses the recent rise of algorithms as an object of interest in research, policy, and practice. It does so through a series of provocations that aim to trouble the coherence of the algorithm as an analytic category and to challenge some of the assumptions that characterize current debates. The goal of this piece is thus to stimulate discussion and provide a critical backdrop against which the *Governing Algorithms* conference can unfold. It asks whether and how we can turn the “problem of algorithms” into an object of productive inquiry.

Algorithms, the very idea

3. To judge from the steady increase in algorithm-themed workshops, conference sessions, and media mentions over the past twelve months, algorithms have become a hot topic (indeed, this very conference bears testimony to that). Examples include a recent double session on “The Politics of Algorithms” at the Annual Meeting of the Social Studies of Science (4S); a session on “Automated Trading” that picked up on very similar issues in finance; related workshops on “Code and Control in Online Spaces” in Hamburg and “Algorithms and Markets” at the London School of Economics; blogs like “Culture Digitally” that regularly write about algorithms and “#algotop” that tracks their appearance in popular culture; a new research project on “Algorithmic Living” at the Intel Science and Technology Center for Social Computing; myriad press articles on everything from algorithmic matchmaking to algorithms’ human helpers; and a series of books, including “Automate This” (Steiner 2012) and “The Black Box Society” (Pasquale, forthcoming). How may we account for this growing interest in algorithms?

4. This question seems even more pressing in light of the well-known tendency in scholarship and practice to proclaim new “eras,” “ages,” or “societies” whenever new technologies allow (cf. Woolgar 2002). Earlier instantiations of such “cyberbole” comprised the “Dot com era,” “Web 2.0,” “user-generated content,” and “big data.” Are algorithms just another fad in a series of grandiose claims about purportedly “new” and “revolutionary” technologies? What explains their sudden claim to fame? Why have commentators rallied around this term?

5. In 2004, Lorraine Daston called for research “devoted to the history and mythology (in the sense of Roland Barthes) of the algorithm” (Daston 2004: 362). She herself has taken up this call, recently telling a history that shows how, “[i]n the models of game theory, decision theory, artificial intelligence, and military strategy, the algorithmic rules of rationality replaced the self-critical judgments of reason” (Daston 2013). Matthew Jones has also recently begun work on the history of data mining, examining “its growth as a fundamental new form of technical expertise in business and scientific research” (Jones 2013). These explorations of the genesis of algorithmic thinking have involved periodizing the algorithm, showing how its particular place in both practice and discourse has changed with time. What can we learn from this history and from the idea that algorithms have experienced distinct periods of change and consolidation?

6. Where some narrowly conceive of algorithms as technology (Clifton, Mulligan and Ramakrishnan 2006), others understand algorithms as a particular form of decision-making (Citron 2007; Gandy 2010) or an epistemology unto itself (Breiman 2001; Hey, Tansley, and Tolle 2009). Still others take a more expansive view, conceiving of algorithms as a particular form of rationality, symptomatic of a general mode of social ordering (Beer 2009; Lyon 2003). And then there are those who see algorithms as a sociotechnical process (Canhoto 2007). Rather than decry the ambiguity around algorithms as a term of art, perhaps we should embrace it. What do we gain by accepting that algorithms are, in fact, multiple?

7. To cite a recent essay, for example, Tarleton Gillespie identifies six dimensions of what he termed “public relevance algorithms”: “patterns of inclusion,” “cycles of anticipation,” “the evaluation of relevance,” “the promise of algorithmic objectivity,” “entanglement with practice,” and “the production of calculated publics” (Gillespie forthcoming). Are these dimensions particular to “public relevance algorithms”? Or are other types of algorithms similarly imbued with, for example, the “promise of algorithmic objectivity”? Furthermore, how might one conceptualize “public relevance algorithms” in contrast to other types of algorithms? If one were to create a typology of algorithms, what types would there be?

8. Gillespie focuses on the criteria that algorithms consider and the particular weights assigned to those criteria. This makes sense when the research question is how to compute relevance (Gillespie forthcoming). But this also may give the misleading impression that algorithms involve express criteria and explicit weighing, when oftentimes they only involve a preference for certain kinds of patterns or structure in the data. If relevance is determined by what counts as a pattern, then should we be studying the patterns that different algorithms prefer? Might we distinguish between different algorithmic genres, based upon these pattern differences?

Methods, knowledge, expertise

9. There is a puzzling tension at the heart of much current reasoning about algorithms. On the one hand, algorithms are invoked as powerful entities that govern, judge, sort, regulate, classify, influence, or otherwise discipline the world. On the other hand, algorithms are portrayed as strangely elusive and inscrutable, or—in fact—as virtually unstudyable (Van Couvering 2007; Roehle 2010). Taken together, these claims open up a tempting space in which algorithms can serve as placeholders or “blank figures” (Hetherington 1999). What is it that makes them so amenable to all kinds of claims, especially about technology and politics?

10. This brings up the closely related question of specificity. A simple test would go like this: would the meaning of the text change if one substituted the word “algorithm” with “computer”, “software”, “machine”, or even “god”? What specifically about algorithms causes people to attribute all kinds of effects to them?

11. Sometimes authors try to circumvent these tensions by prefacing their analyses with all kinds of disclaimers. A common strategy are statements that assure the reader that, of course, the text is not to be read as suggesting an all-too-simple view of algorithms, especially not a deterministic one. It does not take long, however, until these good intentions are jettisoned and various capacities are attributed to algorithms, which tend to form the subject of a sentence. Algorithms do this, algorithms do that. Why are we drawn into this way of speaking? How can we resist it? Against this backdrop, is it possible to study algorithms in a meaningful way – and if so, how?

12. Ian Bogost has asked us to imagine what it would mean to acknowledge that “the perception and experience of other beings [including computers] remains outside our grasp, yet available to speculation thanks to evidence that emanates from their withdrawn cores like radiation around the event horizon of a black hole” (Bogost 2012). In conceiving of algorithms in this way, we concede that the algorithms’ logic may not be available to us—not because it’s concealed, but because it’s entirely beyond our view. Contrast this with the traditional metaphor of a black box, whose inner workings are hidden from view even though the box itself may be visible. What would the study of algorithms look like if we accepted that algorithms are inscrutable because there’s actually nothing to see? What metaphors are most useful for making sense of this?

13. Often, these questions are framed in terms of disciplines and expertise. How do you study algorithms if you don’t know math? Shouldn’t you take a class in computer science first? Understanding algorithms as highly technical concepts suggests the existence of a certain kind of expert – or rather, algorithmic experts. What does such boundary-drawing do? What skills are needed to study algorithms? Who counts as an expert and who doesn’t?

14. Some suggest that there are different ways of studying algorithms. For example, at a recent workshop a scholar suggested that there could be a technical approach that studies algorithms as computer science; a sociological approach that studies algorithms as the product of interactions among programmers and designers; a legal approach that studies algorithms as a figure and agent in law; and a philosophical approach that studies the ethics of algorithms. Compartmentalizing the study of algorithms into disciplines may reduce complexity and, in fact, “discipline” the

discourse. But what are the risks involved in this “division of labor”? Would an interdisciplinary study of algorithms be fruitful? How would it work?

15. Further, questions abound about the distinction between “algorithms” and “data.” Access to an algorithm does not reveal what discoveries that algorithm might facilitate because those discoveries depend as much on the dataset’s features as on the pattern-discovery mechanisms. In other words, a better sense of *how* discoveries can be made would be of little help in assessing *what* discoveries can be made. This upends the common notion that access to the source code of software grants full access to the software’s functionality. Does the study of algorithms therefore require a kind of symmetry, looking at the discovery procedure and parameters (the software code) along with the full range of inputs to that procedure (the dataset)? How can we develop a method of analysis that takes this mutual dependence into account?

Problems and solutions, success and failure

16. Are we studying the transposition of certain questions into problems that can be solved by formalized rules? Is our analytic object the practice of making something computable? Should we scrutinize that practice? And should algorithms be subject to more or less scrutiny in different contexts, such as high frequency trading, predictive policing, retail marketing, political campaigns, and medical diagnosis?

17. The payoff and promise of algorithms has been the overriding focus of the business and popular writing on the topic. But a focus on success often obscures the difficulties with which those successes were earned and the uncertain process by which they were achieved. Instances of failure challenge what this literature often describes as the inherent potential of algorithms to act on big data, but they also reveal the challenges of getting algorithms to work in practice. What if the very possibility of addressing a problem algorithmically were not taken for granted? What might we learn from research that instead focused on poorly formed questions and failed implementations—problems and projects that don’t go anywhere? While some have already challenged the way that data mining companies present their claims to efficacy (Danna and Gandy 2002), new research might look at the way these companies frame the problems to which they then promise a solution. Rather than challenging the claims to efficacy, this work could reveal how problems are made suitable for an “algorithmic solution” in the first place. What could be gained from such an approach?

18. A good example of the challenges of thinking about algorithms in terms of problems and solutions is the field of data mining. Where data mining algorithms now figure as a response to information overload, they actually predated their current demand (Frawley et al. 1996). In fact, this demand had to be created by announcing and explaining the promise of such algorithmic procedures. Thus, ironically, the techniques that many now see as the logical response to the data deluge far preceded the apparent flood. In many cases, the basic algorithms that power even contemporary data mining applications date to as early as the post-war period. What do we gain by rejecting tidy accounts that suggest that the data motivate the mining?

19. We seem to have pivoted from a discourse of overload to a discourse of opportunity. Algorithms have turned an unmanageable flood into a reservoir of ideas. When, why, and how does this happen?

Agency, automation, accountabilities

20. In 1989, Peter Slezak announced the advent of an era of “scientific discovery by computer” that “is totally isolated from all social and cultural factors whatever” (Slezak 1989: 563). In a much discussed article, he suggested:

The work I will describe involves computer programs being developed in the burgeoning interdisciplinary field of cognitive science, and specifically within ‘artificial intelligence’ (AI). The claim I wish to advance is that these programs constitute a ‘pure’ or socially uncontaminated instance of inductive inference, and are capable of autonomously deriving classical scientific laws from the raw observational data. For example, given the rough or ‘noisy’ data used by Robert Boyle 1660, one program is reported to have no difficulty finding the law relating gas pressure and volume. (Slezak 1989: 563-564)

Not surprisingly, Slezak’s claims met fierce resistance. Even though this debate took place more than 20 years ago, it points to an interesting dynamic. Specifically, there appears to be a long-standing fascination with mechanizing discovery (for a recent example, see Anderson 2008). The idea of delegating the production of knowledge to machines has repeatedly attracted the attention of researchers and analysts. What is it about these operations that make them seem so “smart” and “autonomous”? Are we witnessing just another iteration of this dream and its debunking, or is there something different this time around?

21. Algorithms seem to involve two related, but very different forms of automation. In the first instance, they automate the process of subjecting data to analysis, undertaking tasks that would be impossible to perform manually. But the results of these analyses help to automate a second and very different set of operations: decision-making. For example, in sussing out what seem like the signs of creditworthiness through rote computation, algorithms also make the decision to issue credit amenable to automation. What can we learn by paying closer attention to the differences between these kinds of automation and the relationship between the two?

22. This also raises questions of agency and control. Who are the arbiters of these (potentially) brave new algorithms? If Google engineers design a search algorithm in a certain way, do they thereby assert authority over more than the algorithm? Are engineers arbiters of algorithms? Are algorithms arbiters of how information flows within public spheres? And might the transitive property apply, i.e. does that make engineers of algorithms the arbiters again? Or do some algorithms have a species of autonomy? Generally speaking, to what extent can we regard accountability as transitive and mobile across these various entities? Where should accountability lie when an algorithm goes awry and how could this accountability be engineered (Felten 2012)?

23. One way to look at these issues is in terms of “accountability relations” (Woolgar & Neyland, forthcoming). Rather than taking the identity of actors for granted, the focus here is on how the objects and subjects of accountability come about. Accountability, thus, appears as a relational phenomenon that does not only establish obligations, but also establishes those who carry them. Algorithms are prime candidates for such attributions. For example, by portraying them as autonomous decision-makers, their operators can defer accountability. Similar to invocations of “technical failure,” responsibility and blame tend to be put on “the algorithm”. So how do algorithms become “accountabilia” (Ziewitz 2011) – objects for the deferral of accountability? How do they shift and change existing networks of accountability?

Secrecy, obscurity, inscrutability

24. Users encounter algorithms every day in locations far less obvious than a search bar. However, they are often unaware that algorithms are at work. For example, one may request one’s credit score without being privy to the way it was produced. But even if these algorithms were somehow more manifest, would we find that they are nonetheless inscrutable? What if some algorithms are so complex that they are inherently unknowable? What would it mean to even claim to understand an algorithm: knowledge of its foundational assumptions, its inner logic, its likely effects?

25. Unsurprisingly, a persistent complaint among those studying settings in which algorithms appear salient are the difficulties involved in knowing them (cf. Hargittai 2000, Introna and Nissenbaum 2000, Machill, Neuberger and Schindler 2003, Van Couvering 2007, Roehle 2010). Often, these difficulties are expressed in terms of secrecy. Algorithms, the argument goes, are closely guarded trade secrets. “Revealing” the secret sauce would either open doors for manipulation or competition. Interestingly, for those who work with algorithms on a day-to-day basis, secrecy presents itself as a slightly different problem. Search Engine Optimization (SEO) consultants, for example, differentiate themselves from their competitors by claiming “exclusive knowledge” about the inner workings of the engine. At the same time, however, an often voiced opinion goes as follows: “Even if you had Larry and Sergey at this table, they would not be able to give you a recipe for how a specific search results page comes about” (cf. Ziewitz 2012). So what if the secret behind all this is that there is no secret? How can we rethink the so-called “secrecy” problem to make the issue more productive?

26. The inscrutability of certain algorithms presents interesting legal conundrums, too. Although trade secret definitions differ from jurisdiction to jurisdiction, in most of the United States, for information to be considered a trade secret, that information must be “not ... generally known” or “readily ascertainable” (Uniform Trade Secrets Act, Section 1.4). The information must also be “the subject of [reasonable] efforts ... to maintain its secrecy” (Uniform Trade Secrets Act, Section 1.4). Yet if information (here, an algorithm) is inherently secret, how does that alter the application of this definition? Is an inscrutable algorithm necessarily unknown—and is this the same as being kept secret? Is a company making “reasonable efforts under the circumstances to maintain [an algorithm’s] secrecy” if the algorithm is automatically secret?

27. The path by which algorithm-wielding actors arrive at conclusions can be circuitous in the extreme: it involves inducing a rule from an entire population’s behavior that can then be applied

to specific individuals. Further, in detecting subtle patterns within a dataset, algorithms can reveal the many different variables—and the particular weighing of those variables—that correlate with characteristics that cannot be learned directly. What can be known about the individual now depends as much on what others say and do as what the individual might be willing to reveal. How does this alter our understandings of privacy?

Normativity, values, bias

28. Is there such a thing as algorithmic neutrality or impartiality? Or is such neutrality necessarily illusory? Why might the public desire “neutral” or “impartial” algorithms? What values animate the urge for such neutrality? Fairness? Justice? What do those terms mean within an algorithmic context? Yet how does one ascertain whether an algorithm is “good” or “bad,” “fair” or “unfair,” “just” or “unjust”? Good for whom? Fair to what end? Just according to whom? Who—or what—decides? May these values be embodied in an algorithm’s design, however dynamic? Or must one assess the outcomes produced by an algorithm in assessing the algorithm itself?

29. Additionally, how do technical understandings of good and bad interface with philosophical understandings of good and bad? Are the problems that get solved by algorithms those to which algorithms present an obvious solution? If so, does this crowd out efforts to address other issues of social or political significance that are less amenable to algorithmic solutions? Might technical and non-technical understandings of what is normatively “good” become conflated?

30. There is a history of diagnosing “bias” in computer systems. Batya Friedman and Helen Nissenbaum (1996), for instance, consider algorithms as a possible source of technical bias and suggest that “freedom from bias should be counted among the select set of criteria—including reliability, accuracy, and efficiency—according to which the quality of systems in use in society should be judged” (330). But what is it to say that an algorithm is biased—and how would we know? What counts as “unbiased”? And how is bias diagnosed in practice? Can we abandon the idea that there’s a way to avoid bias while still insisting that it’s worth diagnosing bias?

31. Can algorithms enhance or detract from human morality? Can the automation of algorithms help humans with bounded rationality? Might we adopt algorithms to “bind ourselves to the mast” in a variety of contexts? Some have argued that algorithms will know us better than we know ourselves and that they will allow us to be more true to ourselves:

If we have given up autonomy, it was only the autonomy to make poor choices, go to bad restaurants with people we turn out not to like much, buy boring novels, listen to ear splitting music, engage in activities where costs outweigh benefits. I am actually more free now than ever before because my true self—the self that labored under misconceptions, limited information and emotional noise—is now facilitated by powerful and benevolent technology (Ford 2005, 1578-1579)

But in predicting more and more of our interests, will algorithms start to change them? Will we lose control of our own morality by relying on algorithms?

32. Might algorithms have their own morality? In the case of self-driving cars, for example, scholars and the public have started to ask about the moral choices those cars will face (Marcus 2012; Walker Smith 2012). Philosophers have long pondered the moral autonomy and culpability of machines, but they have tended to focus on extreme situations, often of life and death. What kind of work is there to be done on the morality of far more mundane algorithms? What might a “moral” algorithm look like?

33. One could also argue that algorithms embody a profound deference to precedent. According to this line of reasoning, algorithms draw on the past to act on (and enact) the future. Does this result in a self-reinforcing and self-perpetuating system, where individuals are forever burdened by a history that they are encouraged to repeat and from which they are unable to escape? Does deference to past patterns augment path dependence, reduce individual choice, and result in “cumulative disadvantage” (Gandy 2009)? On a larger scale, does this deference to precedent entrench economic stratification and decrease social mobility?

Rules, regulations, resistance

34. The idea of decision rules has long been subject to critical enquiry in a variety of fields. In an old, but still fascinating study, Lawrence Wieder examined the “convict code” operative in a California halfway house and tried to understand how rules like “do not snitch” come to work in everyday practice (Wieder 1974). Perhaps not surprisingly, it was during moments in which these rules were mobilized that order became recognizable—and not the other way around. Instead of simply guiding determinations of good and bad, the code enabled a range of interactions about these very questions. How might these insights apply to predominant conceptions of algorithms as “technologically-embedded rules”? What would it mean to study algorithms in practice (as Wieder studied rules in practice), and not just in theory?

35. A focus on practice squares with a recent move by a number of legal and policy scholars who advocate a move away from the regulation of the collection, circulation, and crunching of data to a regulation of the uses to which data and algorithms are put. Tal Zarsky (2004) and others have argued that we need to regulate use because we can’t regulate collection or analysis, whose implications or effects can’t be known in advance. Goldman makes the most forceful argument to this effect:

Even the worst motivations do not make inconsequential behavior consequential [...] Data mining is not the problem; the problem is bad uses of mined data, and that is where our focus should lie. (Goldman 2006, 229)

On this account, the moral standing of a project that relies on algorithmic processes is entirely independent of the project’s motivations and even the algorithmically derived insights that make the project possible. Is this approach sensible? Or might we be equally concerned with the unintended consequences of algorithmic processes? Or with, perhaps, the consequences of making something amenable to an algorithmic process in the first place?

36. Seemingly, algorithms could be a boon to due process because they formalize decision-making procedures. Consistently rendered decisions, on the basis of articulable and fixed criteria,

are the hallmark of the kind of fairness that due process attempts to ensure. And even where the move to algorithms fails to purge bias from the decision-making procedure, it may at least have the salutary effect of making bias more evident. At the same time, algorithms may involve rules of such complexity that they defy attempts to trace their reasoning. Is this the perfect perversion of due process: the uniform application of an inarticulable rule?

37. This also raises important questions about the relationship between algorithms and rules. May decision-making ever be reduced to the rote application of particular rules? Or is legal interpretation necessarily more complex (Dworkin, 1986)? Legal scholars and judges have argued persuasively that legal interpretation does not entail straightforward textual analysis; decisions are more than mere calls of “ball” or “strike” (e.g. Posner, 2012, 2006). Similarly, wouldn’t decision-making algorithms entail more than straightforward coding, thus timeshifting key interpretive questions? If algorithms were to render decisions, would the locus of legal reasoning shift to the coding of those algorithms?

38. Nevertheless, these ideas underscore the fact that algorithms are deployed in a variety of contexts, some more intuitively troubling than others. The potential and pitfalls of an increasingly algorithmic world—from financial markets to facial recognition software—beg the question of whether legal and policy changes are needed to regulate our changing environment. Should we regulate, or further regulate, algorithms in certain contexts? What would such regulation look like? Is it even possible? What ill effects might regulation itself cause? Given the ubiquity of algorithms, do they, in a sense, regulate us? And what would it mean to resist them?

Conclusion

39. What, then, do we talk about when we talk about “governing algorithms”?

References

- Beer, David (2009). "Power through the algorithm? Participatory web cultures and the technological unconscious," *New Media & Society*, 11, no. 6: 985-1002.
- Breiman, Leo (2001). "Statistical Modeling: The Two Cultures," *Statistical Sciences* 16, no. 3: 199-231.
- Clifton, Christopher W., Dierdre K Mulligan and Raghu Ramakrishnan (2006). "Data Mining and Privacy: An Overview," in Katherine J Strandburg and Daniela Stan Raicu eds, *Privacy and Technologies of Identity: A Cross-Disciplinary Conversation* (New York, NY: Springer).
- Canhoto, Ana Isabel (2007). "Profiling Behaviour: The Social Construction of Categories in the Detection of Financial Crime," London School of Economics and Political Science, London.
- Citron, Danielle Keats (2007). "Technological Due Process," *Washington University Law Review* 85: 1249–1313.
- Danna, Anthony, and Oscar Gandy (2002). "All That Glitters is Not Gold: Digging Beneath the Surface of Data Mining," *Journal of Business Ethics*, 40, no. 4: 373-386.
- Daston, Lorraine J. (2013). "How Reason Became Rationality," Max Planck Institute for the History of Science website, http://www.mpiwg-berlin.mpg.de/en/research/projects/DeptII_Daston_Reason/index_html (last visited March 21).
- (2004). "Whither Critical Inquiry?" *Critical Inquiry* 30, no. 2 (2004): 361-364.
- Dworkin, Ronald (1986). *Law's Empire*, Cambridge, MA: Harvard University Press.
- Frawley, William J, Gregory Piatetsky-Shapiro, and Christopher J Matheus (1992). "Knowledge Discovery in Databases: An Overview," *AI Magazine* 13, no. 3: 57-70.
- Friedman, Batya, and Helen Nissenbaum (1996). "Bias in Computer Systems," *ACM Transactions on Information Systems* 14, no. 3: 330–347.
- Ford, Richard T. (2000). "Save the Robots: Cyber Profiling and Your So-Called Life," *Stanford Law Review* 52, no. 5: 1573-1584.
- Gandy, Oscar (2010). "Engaging Rational Discrimination: Exploring Reasons for Placing Regulatory Constraints on Decision Support Systems," *Ethics and Information Technology* 12, no. 1: 29–42.
- Gillespie, Tarleton (2011). "Can an algorithm be wrong? Twitter Trends, the specter of censorship, and our faith in the algorithms around us," *Culture Digitally*, Oct. 19, available at: <http://culturedigitally.org/2011/10/can-an-algorithm-be-wrong/> (last visited Dec. 10, 2012).

— (forthcoming). The Relevance of Algorithms, in Tarleton Gillespie, Pablo Boczkowski, and Kirsten Foot (eds), *Media Technologies* (Cambridge, MA: MIT Press).

Goldman, Eric (2006). "Data Mining and Attention Consumption," in Katherine J Strandburg and Daniela Stan Raicu eds, *Privacy and Technologies of Identity: A Cross-Disciplinary Conversation* (New York, NY: Springer).

Graham, Scott D. N. (2005). "Software-sorted geographies", *Progress in Human Geography*, vol. 29, no. 5: 562-580.

Hargittai, Esther (2000). "Open portals or closed gates? Channeling content on the World Wide Web", *Poetics*, vol. 27, no. 4: 233-253.

Hey, Tony, Stewart Tansley, and Kristin Tolle (2009). *The Fourth Paradigm: Data Intensive Scientific Discovery*, Redmond, WA: Microsoft.

Institute for Media Archeology (2011). Symposium | insight-turm, available at: http://insight-turm.ima.or.at/?page_id=12&lang=en&lang=en (last visited Sept. 28, 2012).

Jones, Matthew (2013). Matthew Jones - Faculty - Department of History - Columbia University, <http://history.columbia.edu/faculty/Jones.html> (last visited March 21).

Lash, Scott (2007). "Power after hegemony: Cultural studies in mutation?," *Theory, Culture & Society*, vol. 24, no. 3: 55-78.

Lind, E. Allan, and Tom R. Tyler (1988). *The Social Psychology of Procedural Justice*, New York, NY: Springer.

Lohr, Steve (2013). "Algorithms Get a Human Hand in Steering Web," *The New York Times*, Mar. 10, 2013.

Lyon, David (2003). "Surveillance as Social Sorting: Computer Codes and Mobile Bodies," in David Lyon ed, *Surveillance as Social Sorting: Privacy, Risk, and Digital Discrimination* (New York, NY: Routledge).

Marcus, Gary (2012). "Moral Machines," *The New Yorker*, Nov. 27, 2012, available at: <http://www.newyorker.com/online/blogs/newsdesk/2012/11/google-driverless-car-morality.html>

McGee, Matt (2011). SMX East Keynote: A conversation with Eli Pariser, available at: <http://searchengineland.com/smx-east-liveblog-keynote-conversation-with-eli-pariser-92782> (last accessed Sept. 28, 2012).

Posner, Richard A. (2006). "The Role of the Judge in the Twenty-First Century," *Boston University Law Review*, vol. 86: 1049-1068.

— (2012). “The Incoherence of Antonin Scalia (reviewing Reading Law: The Interpretation of Legal Texts, By Antonin Scalia and Bryan A. Garner),” *The New Republic*, Aug. 24.

Rawls, John (1971). *A Theory of Justice*, Cambridge, MA: Harvard University Press.

Roehle, Theo (2010). *Der Google-Komplex: Über Macht im Zeitalter des Internets*, Bielefeld: Transcript.

Slavin, Kevin (2011). “How algorithms shape our world,” *TED Talk*, available at: http://www.ted.com/talks/kevin_slavin_how_algorithms_shape_our_world.html (last visited Sept. 26, 2012).

Smith, Bryant Walker (2012). “How does a traffic cop ticket a driverless car?” *New Scientist*, Dec. 24, available at: <http://www.newscientist.com/article/mg21628966.400-how-does-a-traffic-cop-ticket-a-driverless-car.html> (last visited January 20, 2013).

Spring, Tom (2011). “How Google, Facebook and Amazon run the internet,” *PC Advisor*, available at: <http://www.pcadvisor.co.uk/features/internet/3304956/how-google-facebook-and-amazon-run-the-internet/> (last visited Sept. 28, 2012).

Steiner, Christopher (2012). *Automate This: How Algorithms Came to Rule Our World*, New York, NY: Portfolio/Penguin.

Uniform Trade Secrets Act, Section 1.4.

Van Couvering, Elizabeth (2007). “Is relevance relevant? Market, science, and war: Discourses of search engine quality,” *Journal of Computer-Mediated Communication*, vol. 12, no. 3: 866-887.

Winner, Langdon (1980). “Do Artifacts Have Politics?” *Daedalus* 109, no. 1: 121–136.

Woolgar, Steve and Daniel Neyland (forthcoming). *Mundane Governance*, Oxford: Oxford University Press.

Zarsky, Tal (2004). “Desperately Seeking Solutions: Using Implementation-Based Solutions for the Troubles of Information Privacy in the Age of Data Mining and the Internet Society,” *Maine Law Review* 56, no. 1: 13-59.

Ziewitz, Malte (2012). *Evaluation as Governance: The Practical Politics of Reviewing, Rating and Ranking on the Web*. DPhil Thesis, University of Oxford.

— (2011). “How to think about an algorithm? Notes from a not quite random walk,” Discussion paper for Symposium on “Knowledge Machines between Freedom and Control”, 29 September.