Privacy, Algorithms and Artificial Intelligence

Catherine Tucker*

August 31, 2017

Abstract

Artificial intelligence can use an individual's data to make predictions about what they might desire, be influenced by, or do. The use of an individual's data in this process raises privacy concerns. This article focuses on what is novel about the world of artificial intelligence and privacy, arguing that the chief novelty lies in the potential for data persistence, data repurposing, and data spillovers.

*Catherine Tucker is the Sloan Distinguished Professor of Management Science at MIT Sloan School of Management, Cambridge, MA, and Research Associate at the NBER.

Imagine the following scenario. You are late for a hospital appointment and searching frantically for a parking spot. You know that you often forget where you parked your car. So you use an app you downloaded called 'Find my Car'. The app takes a photo of your car, and then geocodes the photo enabling you easily to find exactly the right location when you come to retrieve it and can predict accurately when it should provide the prompt. This all sounds very useful. However, this example illustrates a variety of privacy concerns in a world of artificial intelligence.

1. Data Persistence: This data, once created, may potentially persist longer than the human that created it, given the low costs of storing such data.

2. Data Repurposing: It is not clear how such data could be used in the future. Once created, such data can be indefinitely repurposed. For example, in a decade's time, parking habits may be part of the data used by health insurance companies to allocate an individual to a risk premium.

3. Data Spillovers: There are potential spillovers for others who did not take the photo. The photo may record other people and they may be identifiable through facial recognition; incidentally captured cars may be identifiable through license plate databases. These other people did not choose to create the data, but my choice to create data may have spillovers for them in the future.

This article will discuss these concerns in detail, after considering how the theory of the economics of privacy relates to artificial intelligence ('AI').

1 The theory of privacy in economics and Artificial Intelligence

1.1 Current models of economics and privacy and their flaws.

The economics of privacy has long being plagued by a lack of clarity about how to model privacy over data. Most theoretical economic models model privacy as an intermediate good (Varian, 1996; Farrell, 2012). This implies that an individual desire for data privacy will depend on how they anticipate that data will affect future economic outcomes. If for example, this data may lead a firm to charge higher prices based on the behavior they observe in the data, a consumer may desire privacy. If a datum may lead a firm to make intrusive intrusions on my time, then again a consumer may desire privacy.

However, this contrasts with, or at the very least has a different emphasis, to how many policy makers and even consumers think about privacy policy and choice.

First, much of the policy debate involves whether or not consumers are capable of making the right decision surrounding the decision to provide data, and on whether 'notice and consent' is sufficient to inform consumers to make the right choice. Work such as McDonald and Cranor (2008) emphasizes that even ten years ago, it was unrealistic to think that consumers would have time to properly inform themselves about how their data may be used through privacy policies as reading them would take an estimated 244 hours each year. Since that study, the amount of devices (thermostats, smart phones, apps cars) collecting data has increased dramatically, suggesting that it is if anything more implausible now that a consumer has the time to actually understand the choice they are making in each of these instances.

Relatedly, even if customers are assumed to have been adequately informed, a new 'behavioral' literature on privacy shows that well-documented effects from behavioral economics, such as the 'endowment effect' or 'anchoring,' may also distort the ways customers make decisions surrounding their data (Acquisti et al., 2016). Such distortions may allow for policy interventions of the 'nudge' type to allow consumers to make 'better' decisions (Acquisti, 2010).

Third, this theory presupposes that customers will only desire privacy if their data is actually used for something, rather than experiencing distaste at the idea of their data being collected. Indeed, some of the earliest work on privacy in the internet era, stated Varian

(1996), 'I don't really care if someone has my telephone number as long as they don't call me during dinner and try to sell me insurance. Similarly, I don't care if someone has my address, as long as they don't send me lots of official-looking letters offering to refinance my house or sell me mortgage insurance.'

However, there is evidence to suggest that people do care about the mere fact of collection of their data, to the extent of changing their behavior, even if the chance of their suffering meaningfully adverse consequences from that collection is very small. Empirical analysis of people's reactions to the knowledge that their search queries (Marthews and Tucker, 2014) had been collected by the US National Security Agency (NSA), shows a significant shift in behavior even when that data was not going to be used by the Government to identify terrorists as it was simply personally embarrassing. Legally speaking, the Fourth Amendment of the US Constitution covers the "unreasonable seizure" as well as the "unreasonable search" of people's "papers and effects", suggesting that governments, and firms acting on government's behalf, cannot entirely ignore seizure of data and focus only on whether a search is reasonable. Consequently, a growing consumer market has emerged for "data-light" and "end to end encrypted" communications and software solutions, where the firm collects much less or no data about their consumers' activities on their platform. These kinds of concern suggest that the fact of data collection may matter as well as how the data is used.

Last, often economic theory assumes that while customers desire firms to have information that allows them better match their horizontally diffferentiated preferences, they do not desire firms to have information that might inform their willingness to pay (Varian, 1996). However, this idea that personalization in a horizontal sense may be sought by customers, goes against popular reports of consumers finding personalizations repugnant or creepy (Lambrecht and Tucker, 2013). Instead, it appears tht personalization of products using horizontally differentiated taste information is only acceptable or successful if accompanied by a sense of control or ownership over the data used, even where such control is ultimately illusory (Tucker, 2014; Athey et al., 2017).

1.2 Artificial Intelligence and Privacy

Like 'privacy,' 'artificial intelligence' is often used loosely to mean many things. This article follows (Agrawal et al., 2016) and focuses on AI as being associated with reduced costs of prediction. The obvious effect that this will have on the traditional model of privacy is that more types of data will be used to predict a wider variety of economic objectives.

Again, the desire (or lack of desire) for privacy will be a function of an individual's anticipation of the consequences of using their data in a predictive algorithm. If they anticipate that they will face worse economic outcomes if the AI uses their data, they may desire to restrict their data sharing or creating behavior.

It may be the that the simple dislike or distaste for data collection will transfer to the use of automated predictive algorithms to process their data. The creepiness that leads to a desire for privacy which is attached to the use of data would be transferred to algorithms. Indeed, there is some evidence of a similar behavioral process where some customers only accept algorithmic prediction if it is accompanied by a sense of control (Dietvorst et al., 2016).

In this way, the question of AI algorithms seems simply a continuation of the tension that has plagued earlier work in the economics of privacy. So, a natural question is whether AI presents new or different problems. This article argues that many of the questions of AI and privacy choices will constrain the ability of customers in our traditional model of privacy to make choices regarding the sharing of their data. I emphasize three themes which I think may distort this process in important and economically interesting ways.

2 Data Persistence, AI and Privacy

'Data persistence' refers to the fact that once digital data is created it is difficult to delete completely. This is true from a technical perspective (Adee, 2015). Unlike analog records, which can be destroyed with reasonable ease, the intentional deletion of digital data requires resources, time and care.

2.1 Unlike previous eras, data created now is likely to persist

Cost constraints that used to mean that only the largest firms could afford to store extensive data, and that for a limited time, have essentially disappeared.

Large shifts in the data supply infrastructure have rendered the tools for gathering and analyzing large swaths of digital data commonplace. Cloud-based resources such as Amazon, Microsoft, and Rackspace make these tools not dependent on scale[1] and storage costs for data continue to fall, so that some speculate they may eventually approach zero.[2]This allows ever smaller firms to have access to powerful and inexpensive computing resources. This decrease in costs suggest that data may be stored indefinitely and be used in predictive exercises should it be thought to be a useful predictor.

The chief resource constraint on the deployment of 'big data' solutions, is a lack of human beings with the data science skills to draw appropriate conclusions from analysis of large data sets (Lambrecht and Tucker, 2015). As time and skills evolve, this constraint may become less pressing.

Digital persistence may be concerning from a privacy point of view because privacy preferences may change over time. The privacy preference that an individual may have felt when they created the data may be inconsistent with the privacy preference of their older self. This is something we documented in Goldfarb and Tucker (2012). We showed there that while younger people tended to be more open with data, as they grew older their preference for withholding data grew. This was a stable effect that persisted across cohorts. It is not the case that young people today are unusually casual about data; all generations when younger are more casual about data; but this pattern was simply less visible previously,

[1]`http://betanews.com/2014/06/27/comparing-the-top-three-cloud-storage-providers/`

[2]`http://www.enterprisestorageforum.com/storage-management/can-cloud-storage-costs-fall-to-zero-1.html`

because social media, and other ways of sharing and creating potentially embarrassing data, did not yet exist.

This implies that one concern regarding AI and privacy is that it may use data which was created a long time in the past, which in retrospect the individual regrets creating.

Data that was created at t=0 may have seemed innocuous at the time, and in isolation may still be innocuous at t=t+1; but increased computing power may be able to derive much more invasive conclusions from aggregations of otherwise innocuous data at t+1 relative to t. Second, there is a whole variety of data generated on individuals that individuals do not necessarily consciously choose to create. This not only includes incidental collection of the data such as being photographed by another party, but also data generated by the increased passive surveillance of public spaces, and the use of cellphone technology without full appreciation of how much data about an individual and location it discloses to third parties, including the government.

Though there has been substantial work in bringing in the insights of behavioral economics into the study of the economics of privacy, there has been less work on time-preference consistency, despite the fact that it is one of the oldest and most studied Strotz (1955); Rubinstein (2006) phenomena in behavioral economics. Introducing the potential for myopia or hyperbolic discounting into the way we model privacy choices over the creation of data seems therefore an important step. Even if the economist concerned rejects behavioral economics or myopia as an acceptable solution, at the very least it is useful to emphasize that privacy choices should be modeled not as something where the time between the creation of the data and the use of the data is trivial, but instead is more acceptably modeled as a decision that may be played out over an extended amount of time.

2.2 How long will data's predictive power persist?

If we assume that any data created will probably persist, given low storage costs, it may be that the more important question for understanding the dynamics of privacy is the question of how long data's predictive power persists.

It seems reasonable to think that much data created today does not have much predictive power tomorrow. This is something we investigated in Chiou and Tucker (2014) where we showed that the length of data retention period that search engines were restricted to by the EU did not appear to affect the success of their algorithm at generating successful search results. This is where the success of a search result was measured by whether or not the user felt compelled to search again. This may make sense in the world of search engines where many searches are either unique or focused on new events. On August 31st 2017, for example, the top trending search on Google is 'Hurricane Harvey,' something that could not have been predicted on the basis of search behavior from more than a few weeks ago.[3].

However, there are some forms of data where it is reasonable to think that their predictive power will persist almost indefinitely. The most important example of this is the creation of genetic digital data. As Miller and Tucker (2017) points out that companies such as 23andme.com are creating large repositories of genetic data, spanning more than 1.2 million people. As pointed out by Miller and Tucker (2017), genetic data has the unusual quality that it does not change over time.

While the internet browsing behavior of a 20 year old may not prove to be good at predicting their browsing behavior aged 40, the genetic data of a 20 year old will almost perfectly predict the genetic data of that person when they turn 40.[4]

[3] `https://trends.google.com/trends/`

[4] As discussed in articles such as `http://www.nature.com/news/2008/080624/full/news.2008.913.html` DNA does change somewhat over time, but that change is itself somewhat predictable.

3 Data Repurposing, AI and Privacy

The time frame that digital persistence of data implies increases uncertainty surrounding how the data will be used. This is because once created a piece of data can be reused an infinite number of times. As prediction costs are lower, this generally expands the number of circumstances and occasions in which data may be used. If an individual is unable to reasonably anticipate how their data may be repurposed or what the data may predict in this repurposed setting, this makes modeling their choices over the creation of their data more difficult and problematic than in our current very deterministic models which assume certainty over how data will be used.

3.1 Unanticipated Correlations

There may be correlations in behavior across users that may not be anticipated when data is created. And it is in these kind of spillovers it seems that the largest potential consequences for privacy of AI may be found.

One famous example of this is that someone liking (or disliking) 'Curly Fries' on Facebook would have been unable to reasonably anticipate it would be predictive of intelligence (Kosinski et al., 2013) and potentially therefore used as a screening device by algorithms aiming to identify desirable employees or students.[5]

3.2 Unanticipated Distortions in Correlations

In these cases, an algorithm could potentially make projection based on a correlation in the data, using data that was created for a different purpose. The consequences for models of economics of privacy is that they assume a singular use of data, rather than allowing for the potential of resuse in unpredictable contexts.

However, even supposing that individuals were able to reasonably anticipate the repur-

[5]This study found that the best predictors of high intelligence include "Thunderstorms," "The Colbert Report," "Science," and "Curly Fries," whereas low intelligence was indicated by "Sephora," "I Love Being A Mom," "Harley Davidson," and "Lady Antebellum."

posing of their data, there are incremental challenges with thinking about their ability to project distortions that might come about as a result of the repurposing of their data.

The potential for distortions based on correlations in data is something we investigate in new research.[6]

In Miller and Tucker (2018) we document the distribution of advertising by an advertising algorithm that attempts to predict a person's 'ethnic affinity' from their data online. We ran multiple parallel ad campaigns targeted at 'African American', 'Asian American' and 'Hispanic' ethnic affinities. We also run an additional campaign which was targeted at those who were not judged to have any of these three ethnic affinities. These campaigns highlighted a Federal Program designed to enhance pathways to a Federal Job via Internships and Career guidance.[7] We ran this ad for a week, and collected data on how many people in which county the ad was shown to. We found that relative to what would be predicted by the actual demographic makeup of that county given census data, the ad algorithm tends to predict that more people are African American in states where there is a historical record of discrimination against African Americans. This pattern is true for states that allowed slavery at the time of the American Civil war, and also true for states that restricted the ability of African Americans to vote in the twentieth century. In such states, it was only the presence of African Americans that was overpredicted, not people with Hispanic or Asian American backgrounds.

We show that this cannot be explained by the algorithm responding to behavioral data in these states, as they were there was no difference in click through patterns across different campaigns across states with or without these history of discrimination.

We discuss how this can be explained by four facts about how the algorithm operates:

1. The algorithm identifies a user as having a particular ethnic 'affinity' based on their

[6]This new research will be the focus of my presentation at the NBER meetings.

[7]For details of the program see `https://www.usajobs.gov/Help/working-in-government/unique-hiring-paths/students/`.

liking of cultural phenomena such as celebrities, movies, TV shows and music.

2. People who have lower incomes are more likely to use social media to express interest in celebrities movies, TV shows and music.

3. People who have higher incomes are more likely to use social media to express their thoughts about the politics and the news.[8]

4. Research in economics has suggested that African Americans are more likely to have lower incomes in states which have exhibited historic patterns of discrimination (Sokoloff and Engerman, 2000; Bertocchi and Dimico, 2014).

The empirical regularity that an algorithm predictive of race is more likely to predict someone to be black in geographies which have historic patterns of discrimination matters because it highlights the potential for historical persistence in algorithmic behavior. It suggests that dynamic consequences of earlier history may affect how artificial intelligence makes predictions. When that earlier history is repugnant, it is even more concerning. In this particular case the issue is using a particular piece of data to predict a trait with, when the generation of that data is endogenous.

This emphasizes that privacy policy in a world of predictive algorithms is more complex than in a straightforward world where individuals make binary decisions about their data. In our example, it would seem problematic to bar low-income individuals from expressing their identities via their affinity with musical or visual arts. However, their doing so could lead them to be more likely to be predicted to be of a particular ethnic group, and they may not be aware *ex ante* of the risk that disclosing a musical preference may cause Facebook to infer an 'ethnic affinity' and advertise to them on that basis.

[8]One of the best predictors of high income on social media is a liking of 'Dan Rather.'

3.3 Unanticipated Consequences of unanticipated repurposing

In most economic models, a consumer's prospective desire for privacy in the data depends here on the consumer being able to accurately forecast the uses to which the data is put. One problem with data privacy is that AI/algorithmic use of existing data sets may be reaching a point where data can be used and recombined in ways that people creating that data in, say, 2000 or 2005, could not reasonably have foreseen or incorporated into their decision-making at the time.

Again, this brings up legal concerns where an aggregation, or 'mosaic', of data on an individual is held to be sharply more intrusive than each datum considered in isolation. In *US v. Jones (2012)*, Justice Sotomayor wrote in a well-known concurring opinion, "It may be necessary to reconsider the premise that an individual has no reasonable expectation of privacy in information voluntarily disclosed to third parties [...] This approach is ill suited to the digital age, in which people reveal a great deal of information about themselves to third parties in the course of carrying out mundane tasks." AI systems have shown themselves to be able to develop very detailed pictures of individuals' tastes, activities and opinions based on analysis of aggregated information on our now digitally intermediated "mundane tasks." Part of the risk in a 'mosaic' approach for firms is that data previously considered to not be 'personally identifiable' or 'personally sensitive information', such as your five-digit zipcode, your gender, or your age to within ten years, when aggregated and analyzed by today's algorithms, may in combination suffice to identify you as an individual.

This general level of uncertainty surrounding the future use of data, coupled with certainty that it will be potentially useful to firms, affects the ability of a consumer to be able to clearly make a choice to create or share data. With large amounts of risk and uncertainty surrounding how private data may be used, this has implications for how an individual may process their preferences regarding privacy.

4 Data Spillovers, AI and Privacy

In the US, privacy has been defined as an individual right, specifically an individual's right to be left alone (Warren and Brandeis, 1890) (in this specific case, from journalists with cameras).

Economists' attempts to devise a utility function which reflects privacy have reflected this individualistic view. A person has a preference for keeping information secret (or not) because of the potential consequences for their interaction with a firm. So far, their privacy models have not reflected the possibility that another person's preferences or behavior could have spillovers on this process.

5 Some types of data used by algorithms may naturally generate spillovers

For example, in the case of genetics the decision to create genetic data has immediate consequences for family members, since one individual's genetic data is significantly similar to the genetic data of their family members. This creates privacy spillovers for relatives of those who upload their genetic profile to 23andme. Data that predicts I may suffer from bad eyesight or macular degeneration later in life could be also use to reasonably predict that those who are related to me by blood may also be more likely to share a similar risk profile.

Of course, one hopes that an individual would be capable of internalizing the potential externalities on family members of genetic data revelation, but it does not seem far-fetched to imagine situations of estrangement where such internalizing would not happen and there would be a clear externality.

Outside the realm of binary data, there are other kinds of data which by their nature may create spillovers. These include photo, video and audio data taken in public places. Such data may be created for one purpose such as the result of a recreational desire to use video to capture a memory or to enhance security, but may potentially create data about other individuals who voices or images are captured without them being aware that their data is being recorded. Traditionally, legal models of privacy have distinguished between the idea of a private realm where an individual has an expectation of privacy and a public realm where an individual can have no reasonable expectation of privacy. For example in the Supreme Court case *California v. Greenwood (1988)*, the court refused to accept that an individual had a reasonable expectation of privacy in garbage he had left on the curb.

However, in a world where people use mobile devices and photo capture extensively, facial recognition allows accurate identification of any individual while out in public, and individuals have difficulty avoiding such identifications. Encoded in the notion that we do not have a reasonable expectation of privacy in the public realm are two potential errors: That one's presence in a public space is usually transitory enough to not be recorded, and that the record of one's activities in the public space will not usually be recorded, parsed an exploited for future use. Consequently, the advance of technology muddies the allocation of property rights over the creation of data. In particular, it is not clear how video footage of my behavior in public spaces which can potentially accurately predict personal outcomes such as health outcomes can be clearly dismissed as being a context where I had no expectation of privacy, or at least no right to control the creation of data. In any case, these new forms of data, due in some sense to the incidental nature of data creation seem to undermine the clear cut assumption of easily definable property rights over the data which is integral to most economic models of privacy.

5.1 Algorithms themselves will naturally create spillovers across data

One of the major consequences of AI and its ability to automate prediction is that there may be spillovers between individuals and other economic agents. There also may be spillovers across a person's decision to keep some information secret, if such secrecy predicts other aspects of that individual's behavior that AI might be able to project from.

Research has documented algorithmic outcomes that appear to be discriminatory, and has argued that such outcomes may occur because the algorithm itself will learn to be biased on the basis of the behavioral data that feeds it (O'Neil, 2017). Documented alleged algorithmic bias spans charging more to Asians for test-taking prep software[9], to black names being more likely to produce 'criminal record' check ads (Sweeney, 2013), to women being less likely to seeing ads for an executive coaching service Datta et al. (2015).

Such 'data-based discrimination' is often held to be a privacy issue (Custers et al., 1866). The argument is that it is abhorrent for a person's data to be used to discriminate against them - especially if they did not explicitly consent to its collection in the first place. However, though not often discussed in the legally orientated data-based discrimination literature, there are many links between the fears expressed for the potential of data-based discrimination and the earlier economics literature on statistical discrimination literature. In much the same way that some find it distasteful when an employer extrapolates from general data on fertility decisions and consequences among females, to project similar expectations of fertility and behavior onto a female employee, an algorithm making similar extrapolations is equally distasteful. Such instances of statistical discrimination by algorithms may reflect spillovers of predictive power across individuals, which in turn may not be necessarily internalized by each individual.

However, as of yet there have been few attempts to try to understand why ad algorithms can produce apparently discriminatory outcomes, or whether the digital economy itself may play a role in apparent discrimination. I argue that above and beyond the obvious similarity to the the statistical discrimination literature in economics, sometimes apparent discrimination can be best understood as spillovers in algorithmic decision making. This makes

[9] `https://www.propublica.org/article/asians-nearly-twice-as-likely-to-get-higher-price-from-princeton-review` - in this case, the alleged discrimination apparently stemmed from the fact that Asians are more likely to live in cities which have higher test prep prices.

the issue of privacy not just one of the potential that an individual's data can be used to discriminate against them.

In Lambrecht and Tucker (2017), we discuss a field study into apparent algorithmic bias. We use data from a field test of the display of an ad for jobs in the Science, Technology, Engineering and Math fields (STEM). This ad was less likely to be shown to women. This appeared to be a result of an algorithmic outcome, as the advertiser had intended the ad to be gender-neutral. We explore various ways that might explain why the algorithm acted in an apparently discriminatory way. An obvious set of explanations is ruled out. For example, it is not because the predictive algorithm has fewer women to show the ad to, and it is not the case that the predictive algorithm learns that women are less likely are to click the ad, since women are more likely to click on it - conditional on being shown the ad - than men. In other words this is not simply statistical discrimination. We also show it is not that the algorithm learned from local behavior that itself may historically have been biased against women. We use data from 190 countries and show that the effect we measure does not appear to be influenced by the status of women in that country. Instead, we present evidence that the algorithm is reacting to spillovers across advertisers. Women are a prized demographic among advertisers, both because they are often more profitable and because they control much household expenditure. Therefore, profit-maximizing firms pay more to show ads to female eyeballs than male eyeballs, especially in younger demographics. These spillovers across advertisers and the algorithms' attempt to cost-minimize given these spillovers, explain the effect we measure. Women are less likely to see an intended gender-neutral ad due to crowding out effects.

To put it simply, our results are the result of these factors:

1. The ad algorithm is designed to cost-minimize so that advertisers advertising dollars will stretch further.

2. Other advertisers consider female eyeballs to be more desirable and deliver a higher return on investment and therefore are willing to pay more to have their ads shown to women than men.

Lambrecht and Tucker (2017) explores apparent algorithmic bias which is the consequence of clear economic spillovers between the value of a pair of eyeballs for one organization to another. Beyond ensuring that, for example, firms advertising for jobs are aware of the potential consequences, it is difficult to know what policy intervention is needed or the extent to which this should be thought of as a privacy issue rather than analyzed through the already established policy tools set up to address discrimination.

This kind of spillover, though, is another example of how in an interconnected economy, models of privacy which stipulate privacy as an exchange between a single firm and a single consumer may no longer be appropriate for a connected economy. Instead, the way any piece of data may be used by a single firm may itself be subject to spillovers from other entities in the economy, again in ways that may not be easily foreseen at the time of data creation.

6 Conclusion

This essay is a short introduction into the relationship between artificial intelligence and the economics of privacy. It has emphasized three themes: data persistence; data repurposing; and data spillovers. These three areas may present some new challenges for the traditional treatment of privacy within an individual's utility function as they suggest challenges for how we model how an individual may make choices of the creation of data about themselves which can later be used to inform an algorithm. At the highest level, this suggests that future work on privacy in economics may focus on the dynamics of privacy considerations amid data persistence and repurposing and the spillovers that undermine the clarity of property rights over data, rather than the more traditional atomistic and static focus of our economic models of privacy.

References

Acquisti, A. (2010). From the economics to the behavioral economics of privacy: A note. *Ethics and Policy of Biometrics*, 23–26.

Acquisti, A., C. R. Taylor, and L. Wagman (2016). The economics of privacy. *Forthcoming, Journal of Economic Literature*.

Adee, S. (2015). Can data ever be deleted? *New Scientist 227*(3032), 17.

Agrawal, A., J. Gans, and A. Goldfarb (2016). The simple economics of machine intelligence. *Harvard Business Review 17*.

Athey, S., C. Catalini, and C. Tucker (2017). The digital privacy paradox: Small money, small costs, small talk. Technical report, National Bureau of Economic Research.

Bertocchi, G. and A. Dimico (2014). Slavery, education, and inequality. *European Economic Review 70*, 197–209.

Chiou, L. and C. E. Tucker (2014). Search engines and data retention: Implications for privacy and antitrust.

Custers, B., T. Calders, B. Schermer, and T. Zarsky (1866). Discrimination and privacy in the information society. *Studies in applied philosophy, epistemology and rational ethics 3*.

Datta, A., M. C. Tschantz, and A. Datta (2015). Automated experiments on ad privacy settings. *Proceedings on Privacy Enhancing Technologies 2015*(1), 92–112.

Dietvorst, B. J., J. P. Simmons, and C. Massey (2016). Overcoming algorithm aversion: People will use imperfect algorithms if they can (even slightly) modify them. *Management Science*.

Farrell, J. (2012). Can privacy be just another good. *J. on Telecomm. & High Tech. L. 10*, 251.

Goldfarb, A. and C. Tucker (2012). Shifts in privacy concerns. *American Economic Review: Papers and Proceedings 102*(3), 349–53.

Kosinski, M., D. Stillwell, and T. Graepel (2013). Private traits and attributes are predictable from digital records of human behavior. *Proceedings of the National Academy of Sciences 110*(15), 5802–5805.

Lambrecht, A. and C. Tucker (2013). When does retargeting work? Information specificity in online advertising. *Journal of Marketing Research 50*(5), 561–576.

Lambrecht, A. and C. Tucker (2017). Algorithmic discrimination? apparent algorithmic bias in the serving of stem ads. *Mimeo, MIT*.

Lambrecht, A. and C. E. Tucker (2015). Can big data protect a firm from competition?

Marthews, A. and C. Tucker (2014). Government surveillance and internet search behavior. *Mimeo, MIT*.

McDonald, A. M. and L. F. Cranor (2008). Cost of reading privacy policies, the. *ISJLP 4*, 543.

Miller, A. and C. Tucker (2017). Privacy protection, personalized medicine and genetic testing. *Management Science*.

Miller, A. and C. Tucker (2018). Historic patterns of racial oppression and algorithms. *Mimeo, MIT*.

O'Neil, C. (2017). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Broadway Books.

Rubinstein, A. (2006). Discussion of behavioral economics.

Sokoloff, K. L. and S. L. Engerman (2000). History lessons: Institutions, factors endowments, and paths of development in the new world. *The Journal of Economic Perspectives 14*(3), 217–232.

Strotz, R. H. (1955). Myopia and inconsistency in dynamic utility maximization. *The Review of Economic Studies 23*(3), 165–180.

Sweeney, L. (2013). Discrimination in online ad delivery. *ACMQueue 11*(3), 10.

Tucker, C. (2014, October). Social networks, personalized advertising, and privacy controls. *Journal of Marketing Research 51*(5), 546–562.

Varian, H. R. (1996). Economic aspects of personal privacy.

Warren, S. D. and L. D. Brandeis (1890, December). The right to privacy. *Harward Law Review 4*(5), 193–220.