

## The Economics of Privacy



Following a special workshop convened by the Media Policy Project on ‘Automation, Prediction and Digital Inequalities’, **Alessandro Acquisti**, Professor of Information Technology and Public Policy at Carnegie Mellon University, asks what the economic argument is for and against privacy.

It may sound surprising that a review of economics and privacy can be provocative, but it actually is: some of the findings reported in our **recent review** of the economics of privacy are counterintuitive, or at least may challenge what is considered the conventional wisdom pertaining to how economists think about privacy.

Privacy first entered the economic discipline in the late 1970s and early 1980s with the work of seminal scholars such as Richard Posner and George Stigler, who were operating at the intersection of privacy and the law. Privacy was presented as a source of economic inefficiency and a source of redistribution. Either way, privacy was considered bad for the economy.

Nowadays, **a number of claims** can be found in the public debate over privacy and economics:

- *Privacy does not make economic sense* – the market will produce the optimal equilibrium between information protection and sharing.
- *Loss of privacy is the price to pay for the technological advances and benefits of big data* – by imposing regulation, we could create market failures where currently there are none.
- *Sharing personal data is an economic win-win* – for instance, in the case of targeted ads, customers may see adverts for products they might be interested in, advertisers therefore have a greater chance of selling, and the value of the ad is increased which benefits the publisher.

Privacy may be considered redistributive, in the sense that if an employer is not allowed to drug test potential employees, for example, this may cause the employer to hire the wrong person for the job. This causes harm to the employer, and therefore the cost of the employee’s drug habit is actually borne by the employer, rather than the employee. There is therefore a degree of distribution of costs to privacy.

Interestingly, however, more recent research has shown that a *lack* of privacy is also redistributive, as it affects the distribution of economic surplus and economic welfare. Varian, a world renowned economist, encapsulated the crux of the matter in **a very powerful statement**, when he pointed out that it is economically rational for consumers to want to share some personal data and to protect other data, which is itself the very essence of privacy. For instance, in the case of telemarketing, you may want telemarketers to know what you’re interested in so at least when they call you at night and disturb your dinner, you may actually be interested in their offers. However, simultaneously, you may not want them to know *how much* you’re interested in those products – because if they do, they will know exactly how much you’re willing to pay and will therefore ask for exactly that amount of money. The economic surplus of that transaction will go entirely to the seller via first degree price discrimination.

Privacy is not about blocking all information flows, it’s about *managing* information flows. **Recent research** shows that when there is no privacy protection, consumers may be worse off, especially if they make decisions based on short term horizons (which is the reality of much consumer behaviour). In the absence of external protection, over the long term, the consumer surplus may be appropriated by sellers through price discrimination. The point, here, is that there is a **very clear** and obvious economic rationale for privacy.



Privacy does not just affect the allocation of surplus, it also affects the interaction between data holders and data subjects, and the relative balance of powers – the relative surpluses of different economic data subjects. Let me explain how this works in practice. In a **recent study** conducted in the US (the results of which are forthcoming), we sent out CVs for fictional job candidates to employers in the US. We also created LinkedIn profiles for the candidates, as well as Facebook profile pages. On the Facebook profile pages, we included information which was not included in the CVs nor on the LinkedIn pages. Specifically, we included personal information (such as religious beliefs) that is a protected characteristic under US employment law. The aim of the experiment was to determine whether employers sought out information from Facebook, and then whether they were more likely to call one candidate for an interview than another. We found that the more conservative the state was, the more likely the employer was to call the Christian candidate to discuss the job opening as opposed to the Muslim candidate. In some states, the Christian candidate was 6 times more likely to be called for an interview than the Muslim candidate. This suggests that the disclosure of personal information does affect the hiring opportunities of individuals, and therefore the allocation of surplus. Privacy, clearly, makes sense for individuals from a microeconomic perspective.

A second claim in the current debate over privacy is that in order to benefit from technological progress, a loss of privacy is necessary; it is the price we pay for the benefits of big data. There are certainly cases in which obstacles to data sharing create economic inefficiencies, which in turn reduce aggregate welfare and economic surplus. However, there are also cases where, conversely, the absence of privacy and data protection creates economic inefficiencies and reduces economic surplus.

Another claim I mentioned is that sharing personal data is an economic win-win – for instance, in the context of targeted advertising. However, our recent research suggests that the surplus different stakeholders obtain from targeted ads is a function of how much information about consumers is being collected and shared; and – importantly – sharing more information is not always better for consumers.

An example of this is provided by real time bidding in the online advertising world. Real time bidding happens every time you go online. When using sites such as CNN.com (in this instance, the ‘publisher’), there will be advertisers who want to sell you products (for example, trainers, or clothes). Ad exchanges run by companies such as Facebook or Google determine which companies can show adverts on a consumer’s screen. When you, the Internet visitor, reach CNN, the website sends information about you to the ad exchange. The ad exchange recognises you, infers some additional information about you, and passes it on to the advertisers with a request to know if they want to show you an ad. The companies selling the products then decide how much they want to bid to show you an ad, depending upon how valuable you, as a potential customer, look to them. The companies submit bids to the ad exchange, and the ad exchange uses an auction to determine the winning bid and therefore which ad will be shown to the consumer. This occurs in a matter of milliseconds; the consumer doesn’t even realise this is happening.

Now, most economic theory papers that study online ads typically focus on two-player models (usually just the consumer and the advertiser, or the platform and the advertiser). However, the online ad eco-system can be modelled with three players: the consumer, the advertiser and the ad exchange. What we found in our research is that the amount and type of information the ad exchange captures about consumers and shares with the advertisers results in huge differences in how much surplus each of the three parties receive.

When no information about consumers is gathered, most of the surplus goes to the ad exchange, but the consumers get a fair share. When ‘horizontal’ information is shared (i.e. what a customer’s product preferences are), there is a somewhat equal allocation of surplus across the three players. When only ‘vertical’ information is collected and shared (i.e., roughly, whether a consumer is wealthy or not), the consumer gets something. When complete information is shared, consumers get close to zero surplus (because of perfect price discrimination). The bottom line, here, is that



sharing information is not always an economic win-win for consumers and producers. Data trails will determine winners and losers, and the actual outcome is dependent on the context.

The question about the relationship between data and economics is whether we are going to use data in a way which simply changes the allocation of economic surplus, or whether we are going to use data in a way that will increase the total amount of surplus for everyone. I believe that privacy has a role to play in determining which of these two scenarios will occur.

*This blog gives the views of the author and does not represent the position of the LSE Media Policy Project blog, nor of the London School of Economics and Political Science.*

*This post was published to coincide with a workshop held in April 2016 by the Media Policy Project, 'Automation, Prediction and Digital Inequalities'. This was the third of a series of workshops organised throughout 2015 and 2016 by the Media Policy Project as part of a grant from the LSE's Higher Education Innovation Fund (HEIF5). To read a summary of the workshop, please click [here](#).*

---

July 27th, 2016 | [Digital Inequalities](#), [Featured](#) | [0 Comments](#)

---

⌵

