Soon We'll All Be Royalty

Ari Juels

January 12, 2014

Abstract

This paper contemplates a future shaped by two trends. Privacy, in the sense of the ability to keep personal relationships and behaviors secret, will substantially erode. Decision-making, in the sense of organizational allocation of resources, will become increasingly algorithmic and rely on ever richer sources of data.

Rather than seeking to shore up individual privacy, then, this paper posits that it may be more effective to address the *consequences* of lost privacy directly. One way to do this is to audit decision-making algorithms to ensure that they adhere to equitable policies. In other words, rather than trying to conceal information, we might instead seek to enforce its fair use at the point of consumption.

1 Trend 1: Palatial Privacy

The Palace of Versailles, the immensely costly royal chateau that came to symbolize the absolute power of King Louis XIV, showcased many of the greatest architectural luxuries of its day. To the modern visitor, however, one is conspicuously absent. The palace has no corridors. Indeed,

> ...a seventeenth-century palace was totally without privacy. Architects had not yet invented the corridor. To get from one part of the building to another, one simply walked through a succession of other people's rooms, in which literally anything might be going on... The character of the circumambient architecture was such that one could scarcely avoid the spectacle of others being born, dying, relieving nature, and making love.[1]

[1] Huxley, Aldous. *The Devils of Loudun*. Carroll & Graf, 1986, p. 12.

Privacy is defined by the American English Dictionary as “the right to keep one’s personal matters and relationships secret.” In this sense, it is a historical anomaly, a privilege that even the most powerful sovereign of seventeenth-century Europe didn’t enjoy. In the face of pervasive sensing devices and network connections, our own world begins to resemble that of the Palace of Versailles.

The problem of location privacy offers a good example. As wireless microchips, often known as Radio-Frequency IDentification (RFID) tags, starting becoming ubiquitous some ten years ago, privacy advocates warned of their dangers. The use of RFID tags opened up the possibility of automated, clandestine tracking of their bearers by networks of wireless scanning devices.

Today, that concern seems quaint. Mobile devices regularly report their users’ whereabouts to a variety of services, and many consumers carry a multiplicity of uniquely identifiable, beaconing wireless devices (fitness devices, tablets, and so forth). Digital cameras are proliferating, while face-recognition is rapidly improving. Users are even developing the habit of intentionally broadcasting their location through services such as Foursquare. The ability of an ordinary individual to protect her location privacy—and perhaps the desire to do so—is eroding. It seems reasonable to assume that many other forms of privacy will follow suit.

2 Trend 2: Algorithmic Decision Making

Many significant organizational decisions that directly impact the lives of workers and consumers are today made algorithmically. For example, for several years, some Silicon Valley companies have required job candidates to fill out questionnaires (“Have you ever set a regional-, state-, country-, or world-record?”) as part of their application process. These companies apply classification algorithms to the answers to filter applications.[2]

Research in behavioral economics supports the efficiacy of such approaches in what are known as “low-validity environments,” where cultivation of accurate human intuition is challenging. There are many domains in which statistical predictions have been shown to outperform those of human experts. Examples include the diagnosis of cardiac disease, the longevity of cancer patients, the suitability of foster parents, and the future value of Bordeaux wines.[3]

Increasing availability of data, increasing automation of decision making, and increasing evidence of the superiority of algorithms to human judgment in many

[2] S. Hansell, “Google Answer to Filling Jobs Is an Algorithm,” *New York Times*, 3 Jan. 2007.

[3] Kahneman, Daniel. *Thinking, Fast and Slow*. Farrar, Straus and Giroux, 2011, pp. 223-4

domains are all driving the pervasiveness of algorithmic decision making.

3 From Privacy to Accountability

How can we ensure equitable treatment for individuals when personal information is hard to conceal and aggressively consumed in algorithmic decision making? Stated more starkly, if we give up on major forms of privacy, how can we prevent harm?

One idea is to design auditable decision-making systems. Such systems might *prove* two characteristics of their algorithms: (1) They consume certifiably correct data and (2) They comply with a set of published policies governing data use.

For example, using cryptographic or related techniques, a health-insurance company might prove a (mathematical) statement expressed as follows in prose:

> The annual premium for this health-insurance policy is the output of an algorithm P whose only health-related input X originates with an accredited hospital, and contains no genetic data. (X might be digitally signed by the hospital.) The influence of X on the output of P is at most $50. (I.e., $P(X) - P(X') \leq \$50$ for any input pair (X, X').)

Such an attestation can even be constructed so as to reveal no further information about P.

Given this approach, even if a patient's health record is revealed or her fitness gadgets betray information about her health habits—even, at an extreme, if she enjoys no privacy—she may still obtain assurance at an institutional level against the abusive consequences of her personal information being revealed.

Challenges remain, though, such as ensuring against the social stigma or personal implications of information disclosure. Enforcing fair-use policies in augmented reality systems or other mediators of interaction with the physical world might offer a similar approach to limiting the consequences of personal data loss.[4]

[4] Louis XIV took such an approach. Abundant hair was regarded in seventeenth-century Europe as a particularly important indicator of health and vitality. Louis began balding at the age of 17, a fact he couldn't reasonably keep *secret*. Instead, he suppressed the fact when it mattered by augmenting appearances; he created a new court protocol. Louis began to wear a wig (a perruque). Courtiers copied him, as soon did gentlemen across Europe, setting a trend that lasted for well over a century.