What happens when we find out how much everyone else is making?

Technological advances have made it possible for everyone to know potentially everything about everyone else. Sci-fi shows such as Netflix's Black Mirror imagine dystopian scenarios that could result from these new technologies. In the real world, social media is already allowing individuals to disclose details about their personal lives with strangers. This technological change has sparked a policy debate too: should governments disclose data such as tax records?

Tax transparency in Norway

To understand how far reaching the societal consequences of transparency could be, it is useful to look at the experience of one of the pioneers of transparency, namely, Norway. Tax records have been public in Norway since the 19th century, but they have not always been easily accessible. Before 2001, an individual had to make a formal request in person at a tax agency to see someone else's income. In the fall of 2001, the Norwegian media digitised tax records and created websites that allowed any individual with internet access to search the entire country's tax records easily and effortlessly. Every Norwegian could find out the incomes of anyone else in the country in a matter of seconds.

In the following decade, Norwegians engaged in a heated debate about whether tax records should be easily accessible online. A similar debate took place in neighbouring Sweden, Iceland, and Finland. These other Scandinavian countries also had laws making tax records public, and thus had to decide whether to make them easily available online as in Norway.

There was no consensus among politicians or the general public on whether Norway's transparency was good or bad. At its core, the disagreement was about what the effects of transparency actually are. Some supporters argued that public records could serve to deter corrupt politicians and tax evaders (Bø et al. 2015). Meanwhile, detractors claimed that the tax records would be used in objectionable ways, to snoop on the incomes of friends, for example. Thus far, only qualitative and anecdotal evidence has substantiated this contention. In a recent paper (Perez-Truglia 2019), I provide the first quantitative evidence on the matter.

Figure 1 shows a screenshot of one of the websites that allowed Norwegians to browse the tax records. These websites were easy to use and became incredibly popular in the country for the following decade.

Figure 1. Search tool from skattelister.no

Source: Snapshot of the website as of 16 June 2010, downloaded from web.archive.org

The data indicate that, rather than using tax records to serve the public interest (e.g. uncovering corruption or tax evasion), most visitors used the online tax lists to snoop on friends, relatives, and other social contacts. Snooping became so pervasive that the Norwegian media coined a new term for it – 'tax porn'. This is perhaps best illustrated by some of the tools used to browse the tax records. One such tool is a smartphone app that allows users to create leaderboards with the highest and lowest earners among their Facebook friends. Another shows a map with the incomes of everyone living around a specific location.

Survey data indicate that 40 per cent of Norwegian adults reported using the online tax lists in 2007. To put this number in perspective, a decade later, no social media companies other Facebook have reached this penetration rate (in 2019, 35 per cent of US adults reported using Instagram, 25 per cent LinkedIn, and 24 per cent Twitter).

Poorer individuals stand to lose in this game of income comparisons. For example, they may learn that they are poorer than they thought, lowering their self-esteem. If their social contacts learn such information, they might experience shame. Conversely, richer Norwegians stand to gain from income comparisons. Learning that they are richer than they thought can boost their self-esteem, as can being looked up by their social contacts.

To test this hypothesis, we measure the effect of income transparency on the happiness-income relationship, i.e. the gradient between individuals' subjective wellbeing and their position in the income distribution. We use survey data from Norway from 1985–2013 that include the two most widely used measures of subjective wellbeing: happiness and life satisfaction.

Consistent with our hypothesis, we show that the 2001 income transparency change led to a 29 per cent increase in the happiness-income relationship (p-value=0.005) and a 21 per cent increase in the life satisfaction-income relationship (p-value=0.026).

We use multiple strategies to argue that the change in the happiness–income relationship was caused by the increase in income transparency of 2001. First, we conduct an event-study analysis and find that the happiness-income relationship stayed constant in the years before the change, increased in 2001, and persisted at the higher level during the subsequent 12 years of higher transparency.

Second, we identify individuals who were most likely to be exposed to the effects of online tax information, based on observable characteristics that predict internet access. As shown in Figure 2, the happiness–income relationship remained at the pre-2001 level for individuals with lower internet access but increased substantially and persisted at the higher level for individuals with higher internet access.

Figure 2. Event-study analysis
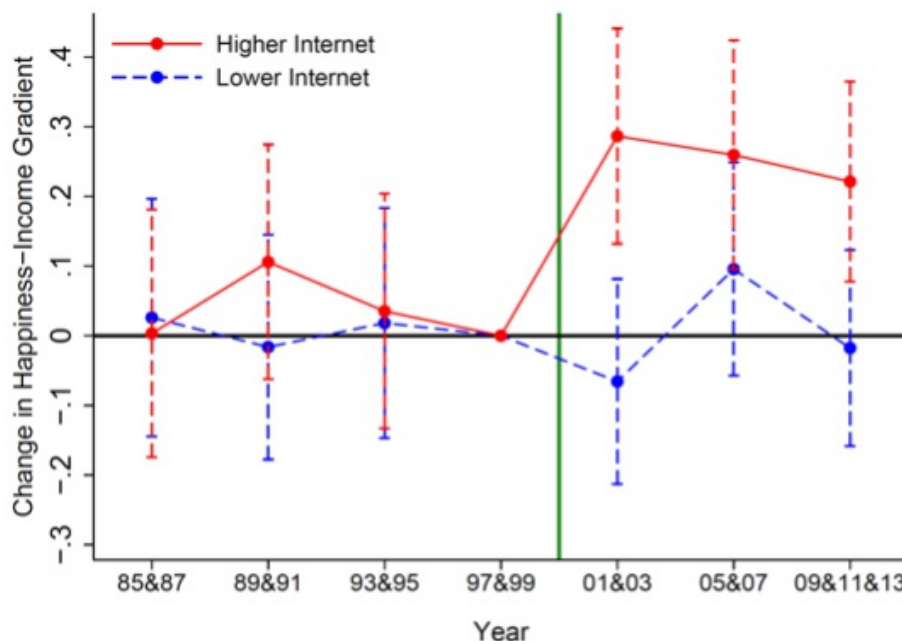
Source: Perez-Truglia (2019)

We also offer a falsification test by reproducing the analysis using similar survey data from Germany, a country that was not affected by a change in income transparency. The event-study analysis shows that the happiness–income relationship remained stable in Germany from 1985 to 2013, both in the population at large and in the sub-populations of individuals with higher and lower internet access.

Our evidence is consistent with related studies suggesting that income comparisons play a significant role in the relationship between income and subjective wellbeing (Luttmer 2005). Our evidence is also consistent with anecdotal evidence reported in the Norwegian media. For example, low-income adults felt shamed by the tax lists and children from poorer households were allegedly bullied at school because of the tax lists. Our findings are also in line with survey data indicating that, relative to richer households, poorer households were more likely to oppose the income transparency policy.

The debate on transparency goes far beyond the disclosure of tax records. Governments around the world disclose all sorts of sensitive information, such as the salaries of public employees (Card et al. 2012, Mas 2017), individual contributions to political campaigns (Perez-Truglia and Cruces 2018), and identities of criminals and tax delinquents (Linden and Rockoff 2008, Perez-Truglia and Troiano 2018). Private organisations can also disclose information that can have societal effects, for example companies can disclose the salaries of their employees (Cullen and Perez-Truglia 2018, Cullen and Pakzad-Hurson 2018).

Policy implications

Our findings suggest that it is important to measure the wellbeing effects of disclosing sensitive data. When choosing transparency policies, these effects must be accounted for in the cost-benefit analysis. However, even if transparency is found to affect the wellbeing of the individuals whose information is being divulged, that does not imply that transparency should be avoided. Rather, we should find alternative ways of disclosing information that mitigates adverse consequences while preserving some of its desirable effects.

For example, the Norwegian government made searches of tax records non-anonymous in 2014. Specifically, any individual could use a website to identify who had looked at his or her tax records. This removal of anonymity was intended to discourage individuals from unsavoury uses of the tax records, such as snooping, due to the threat of social sanctions. The policy was effective – the number of searches on the tax website dropped by 88 per cent after the removal of anonymity. The number of users logging in to the system did not decrease much, but instead of searching for others' incomes, most users logged in to find out who searched for them.

However, non-anonymous searches can still discourage legitimate uses of data. Governments may want to complement the non-anonymous search tools by offering anonymous access to de-identified datasets. For example, some US states list the salaries of all public employees including identifiable information such as full names. Instead, they could offer aggregate data such as average salaries or salary ranges by position and unit. This aggregate data can provide most of the information that individuals need for their decision making, while avoiding harmful effects on the wellbeing of the employees who earn less.

<div align="center">♣♣♣</div>

Notes:

- *This blog post appeared originally on VoxEu, and is published in agreement with the author. It is based on the author's paper "The effects of income transparency on wellbeing: Evidence from a natural experiment", NBER Working Paper No. 25622 (2019)*
- *The post gives the views of its author(s), not the position of LSE Business Review or the London School of Economics and Political Science.*
- *Featured image by derneuemann, under a Pixabay licence*
- *When you leave a comment, you're agreeing to our Comment Policy*

Ricardo Perez-Truglia is an assistant professor of economics at UCLA Anderson School of Management. His research interests include behavioural economics, political economy and public economics. Dr Perez-Truglia received his Ph.D. in economics at Harvard University in 2014. Before joining UCLA Anderson in 2016, he was a postdoctoral researcher at Microsoft Research New England.