NBER WORKING PAPER SERIES

THE EFFECTS OF INCOME TRANSPARENCY ON WELL-BEING: EVIDENCE FROM A NATURAL EXPERIMENT

Ricardo Perez-Truglia

Working Paper 25622
http://www.nber.org/papers/w25622

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2019, Revised March 2019

Thanks to Ipsos MMI for providing access to the survey data. Thanks to Ottar Hellevik for all his help with this study – without him, this project would have not been possible. Thanks to Erzo F.P. Luttmer for invaluable feedback. I benefited from the feedback from several colleagues as well as seminar discussants at Harvard, Princeton, Yale, Universidad de San Andres, ITAM, CIDE, Universidad Catolica de Chile, Universidad de la Republica, RAND, the AEA Annual Meetings, the USC Happiness Conference and the NBER Summer Institute. Diego Marino Fages provided outstanding research assistance. The first draft of this paper circulated in August 2015 under the title "Measuring the Value of Self and Social-Image." The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

The Effects of Income Transparency on Well-Being: Evidence from a Natural Experiment
Ricardo Perez-Truglia
NBER Working Paper No. 25622
February 2019, Revised March 2019
JEL No. D03,D31,D60,D83,I31,Z10

ABSTRACT

In 2001, Norwegian tax records became easily accessible online, allowing everyone in the country to observe the incomes of everyone else. According to the income comparisons model, this change in transparency can widen the gap in well-being between richer and poorer individuals. We test this hypothesis using survey data from 1985–2013. Using multiple identification strategies, we show that the higher transparency increased the gap in happiness between richer and poorer individuals by 29%, and it increased the life satisfaction gap by 21%. We provide suggestive evidence that some, although probably not all, of this effect relates to changes in self-perceptions of relative income. We provide back-of-the-envelope estimates of the importance of income comparisons, and discuss implications for the ongoing debate on transparency policies.

Ricardo Perez-Truglia
Anderson School of Management
University of California, Los Angeles
110 Westwood Plaza
Los Angeles, CA 90095
and NBER
ricardo.truglia@anderson.ucla.edu

1 Introduction

The income comparisons model proposes that individual well-being largely depends on how that individual's income compares to the incomes of others (Luttmer, 2005). This is a fundamental aspect of individual preferences, yet no consensus has been reached about the economic significance of income comparisons.[1] In this study, we offer novel evidence based on a unique natural experiment: in 2001, Norwegian tax records became easily accessible online, allowing everyone in the country to observe the incomes of everyone else quickly and easily. We test the hypothesis that, consistent with the model of income comparisons, increased income transparency widens the gap in well-being between richer and poorer individuals.

Tax records have been public in Norway since the nineteenth century, but they have not always been easily accessible. Before 2001, one had to make a formal request in person at the tax agency to see someone else's income. In the fall of 2001, the Norwegian media digitized tax records and created websites that allowed any individual with Internet access to search anyone's tax records. Every Norwegian was one click away from finding out the incomes of everyone else in the country.

We use various data sources to show the massive popularity of these online tax lists. During the busiest week of the year, these websites were more popular than YouTube. We also show that, rather than using the tax lists for legitimate goals (e.g., uncovering corruption or tax evasion), most used the websites to snoop on friends, relatives, and social contacts. For example, users could create leaderboards showing the highest and lowest earners among their Facebook friends or maps showing the incomes of everyone living around a specific location. This behavior became so pervasive that the Norwegian media dubbed it "tax porn."

Because income transparency facilitates income comparisons, it can widen the gap in well-being between richer and poorer individuals. Poorer individuals often lose this game of income comparisons. For example, if they learn that they are poorer than they thought (Cruces, Perez-Truglia, and Tetaz, 2013), it can lower their self-esteem. If their social contacts learn how poor they are, it can reduce their social-esteem. In contrast, richer individuals often benefit from this game. Learning that they are richer than they thought can boost their self-esteem. And being looked up by their social contacts can boost their social-esteem.[2]

To test this hypothesis, we measure the effect of the increase in transparency on the gradient between subjective well-being and individual income rank (hereinafter referred to as the

[1] This question is important to understand preferences more deeply, and also due to its implications for income taxation and other policies. For example, income comparisons can create positional externalities that reduce social welfare and could be corrected with taxes (e.g., Boskin and Sheshinski, 1978; Frank, 1985).

[2] Different individuals may react differently to the increased transparency. For example, while some rich individuals may feel happy that their neighbors caught a glimpse of their income, others may feel uneasy about the same situation (e.g., if they do not feel deserving of their high income). In this study, we can only measure which of these different mechanisms dominates on average.

happiness-income gradient). We use survey data from Norway from 1985–2013 that includes the two most widely used measures of subjective well-being: happiness and life satisfaction. Despite some limitations of these subjective measures, evidence suggests that they contain useful information about well-being. For example, life satisfaction and happiness have been shown to be significantly correlated with objective measures of well-being and with decision utility (Di Tella, MacCulloch, and Oswald, 2003; Benjamin, Heffetz, Kimball, and Rees-Jones, 2012).

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

We use multiple strategies to identify the causal effect of the transparency change of 2001. First, we conduct an event-study analysis and find that the happiness-income gradient stayed constant in the years before the change, increased in 2001, and persisted at the higher level during the subsequent twelve years of higher transparency.

Second, we identify individuals who were most likely to be exposed to the effects of online tax lists, based on observable characteristics that predict Internet access. We show that, between 1985 and 2000, the happiness-income gradient remained stable for individuals with low and high Internet access. After 2001, the happiness-income gradient 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.

Our third identification strategy reproduces the analysis using similar survey data from Germany, a country that was not affected by the Norwegian change in income transparency. Similar results for Germany would indicate that another factor, such as the dot-com bubble, caused the change in the happiness-income gradient in Norway. In sharp contrast to the Norwegian findings, however, the life satisfaction-income gradient in Germany did not change around 2001. The event-study analysis shows that this gradient 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.

Anecdotal evidence supports our finding that higher income transparency increased the well-being gap between richer and poorer households. For example, the media reported that the online tax lists led to bullying of kids from poorer households and that adults from poorer households felt that they disappointed themselves and others (Aftenposten, 2008; New York Times, 2009). Our findings also align with survey data indicating that, relative to richer households, poorer households were more likely to oppose the income transparency policy (Aftenposten, 2011).

The effects of income transparency may operate through multiple mechanisms. We provide suggestive evidence for one specific mechanism: self-perceptions. According to this channel, richer individuals may be happier because they learn that they are richer than they thought, and poorer individuals may be unhappier because they learn that they are poorer than they thought. We show that, indeed, transparency increased the gradient between perceived income rank and actual income rank by 8.5% (p-value<0.001) and the gradient between the perceived adequacy of one's income and income rank by 4.7% (p-value=0.083). This evidence cannot prove or rule out the self-perceptions channel, but it does serve as suggestive evidence. Moreover, the perceived rank and income adequacy gradients (8.5% and 4.7%) are smaller than the changes in the happiness and life satisfaction gradients (29% and 21%), which suggests the presence of other mediating factors besides self-perceptions.

We use the estimated effects of transparency to assess the economic significance of income comparisons. Our back-of-the-envelope calculations suggest that, as a conservative lower bound, income comparisons accounted for 22% of the happiness that individuals in Norway derived from their incomes during the period of higher transparency. Moreover, we show that this lower bound is consistent with the effect of relative income on happiness, as reported in related studies.

Our evidence also relates to the ongoing debate on transparency. Technological advances have made it possible for everyone to know potentially everything about everyone else, sparking debates on whether the government should disclose its data, such as tax records. Some arguments that favor or oppose transparency are rooted in philosophical grounds.[3] However, most arguments seem to be based on the potential effects of transparency. In particular, detractors of income transparency argued that the tax lists were used in despicable ways to harm the well-being of poorer individuals.[4] This argument, however, was based on qualitative and anecdotal evidence. This study provides the first quantitative evidence on this matter.

Beyond the Norwegian experience, information disclosure may directly affect well-being in other contexts. In the 2000s, Sweden, Iceland, and Finland had to decide whether to make their tax records as easily accessible as in Norway. Outside of Scandinavia, governments disclose all sorts of sensitive information, such as the salaries of public employees (Card, Mas, Moretti, and Saez, 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). Our findings suggest that it is important to measure the well-being effects of disclosing sensitive data and to account for them in the cost-benefit analysis.

This paper relates to various strands of literature. Most important, it relates to a literature on the effect of relative income on well-being. In a seminal contribution, Easterlin (1974) showed evidence that happiness and income are positively correlated across individuals within a country

[3]For example, some of the supporters of income transparency in Norway see it as a fundamental principle of democracy, while some opponents see it as a violation of privacy rights.

[4]This was by no means the only negative consequence from income transparency that was debated. For example, some detractors of open disclosure argued that the tax records could be used by criminals to target rich individuals. However, in a letter to the Ministry of Justice, the Norwegian police noted that their investigations ruled this out as a significant source for concern (Dagens Næringsliv, 2010).

but that average happiness in a country does not seem to rise over time as average income rises. One standard explanation for the paradox is that happiness depends on relative income. Within a given country, richer individuals have higher relative income, so they are happier. However, as every individual in the country becomes richer, the average relative income stays constant, and thus average happiness also remains constant. Consistent with this interpretation, several studies have shown that, holding own income constant, subjective well-being decreases with the mean income of neighbors (Luttmer, 2004; Ferrer-i-Carbonell, 2005).[5]

However, this evidence is subject to concerns about causal identification. For example, the Balassa–Samuelson model (Balassa, 1964; Samuelson, 1964) predicts that consumer prices should be higher in areas where nominal incomes are higher. Thus, even in the absence of income comparisons, happiness should be negatively correlated to the average income of neighbors, reflecting a higher cost of living. More generally, the average income in an area could be correlated with other unobservable attributes of the location that also affect well-being, thus generating omitted-variable biases. We contribute to this literature by presenting novel evidence on the effects of relative income on happiness that relies on a new identification approach, based on quasi-experimental variation in income transparency.

This study also relates to Bø, Slemrod, and Thoresen (2015), who measured the effect of the Norwegian disclosure of tax records on tax evasion. Disclosing tax records may deter tax evasion by encouraging others with relevant information about true tax liability to come forward and by threatening evaders with social sanctions (see also Perez-Truglia and Troiano, 2018). Bø, Slemrod, and Thoresen (2015) found that the change in income disclosure increased reported income among business owners by 2.7%, resulting in a total gain of 0.2% in income tax revenues. This evidence confirms a benefit of disclosure, as alleged by its supporters. We present evidence on an unintended effect, income comparisons, as alleged by detractors of income transparency.

This work also relates to the study by Card, Mas, Moretti, and Saez (2012) of the effects of income transparency on job satisfaction. The researchers sent emails to a random sample of university employees with information on how to access a website that listed the wages of all employees working in the same university. In a follow-up survey, they found that, for workers with below-median salaries within their unit and position, having access to the website decreased satisfaction with their wages and their jobs. Consistent with this finding, Rege and Solli (2015) show evidence that the disclosure of tax records in Norway increased the probability of quitting among workers with lower salaries. Their findings suggest that some poor individuals may benefit from income transparency, because they can find out if they are under-paid and look for a better job. On the contrary, our evidence suggests that income transparency increased the

[5]There are some conflicting accounts about the evidence. See for instance Hagerty and Veenhoven (2003), Stevenson and Wolfers (2008), Easterlin et al. (2010) and Easterlin (2017) on the effect of income growth on happiness, and Senik (2004), Clark, Westergård-Nielsen and Kristensen (2009) and Deaton and Stone (2013) on the effect of relative income on happiness.

well-being of richer individuals at the expense of the well-being of poorer individuals.

Last, this study relates to a literature documenting how individuals misperceive their positions in the income distribution and how providing objective information can correct these misperceptions (Cruces, Perez-Truglia, and Tetaz, 2013; Karadja, Mollerstrom, and Seim, 2017). These studies are based on artificial contexts in which researchers provide information through a survey. We contribute to this literature by exploiting the variation in information access in a natural, large-scale setting and by showing that correcting these misperceptions may affect well-being.

The rest of the paper proceeds as follows. Section 2 describes relevant details about the disclosure policy. Section 3 presents the econometric specification and the survey data. Section 4 presents the results. The last section concludes.

2 Relevant Institutional Details

2.1 Origin of the Online Tax Lists

Although tax records have been publicly available in Norway since the middle of the nineteenth century, they were not easily accessible before 2001. Individuals who wanted to learn about someone else's income had to visit the local tax office or city hall during a three-week period and search through a book with records for thousands of taxpayers from the same municipality.[6] In the fall of 2001, a Norwegian newspaper made these tax records searchable online for the first time so that any Norwegian with Internet access could view them easily and at any time (see Figure 1.a for a screenshot of this website). All major newspapers soon created their own websites, which remained popular in the country for the following decade.[7] These websites listed full names and net incomes (see Figure 1.b for a sample search result), and could also list additional information such as taxes, net worth, birth years, cities and postal codes. These websites allowed visitors to search by multiple fields. For example, visitors could search for their own last name to find relatives. Or they could search by postal code to find neighbors.

Although all other Scandinavian countries (except Denmark) make tax records publicly available, the Norwegian tax disclosure during 2001–2013 was exceptional because of its accessibility (Aftenposten, 2011).[8] In Finland, accessibility of tax records is similar to that in Norway prior

[6] In selected municipalities and only shortly before 2001, some local organizations sold books with information from the local tax rolls. Bø, Slemrod and Thoresen (2013) exploit variation across these municipalities to identify the effects of transparency. We cannot use this same identification strategy because we lack sufficient data (we would need a survey sample orders of magnitude higher and with a higher frequency).

[7] The following is a sample of these websites: www.skattelister.no, www.nrk.no/skatt, www.tu.no/skattelister and skatt.na24.no.

[8] In other countries, information about incomes can be easily accessible online for a subset of the population (e.g., public employees in some U.S. states).

to 2001, requiring individuals to visit the tax agency in person (New York Times, 2018).[9] In Sweden, the requests for tax returns are not anonymous and must be done by phone, making this practice much less popular than it is in Norway.[10] In Iceland, tax records are also difficult to access and available during two weeks of the year only.

2.2 Evolution of the Online Tax Lists

Between 2001 and 2013 (the last year of our survey data), several factors may have contributed to increased or decreased use of online tax lists, though none of these changes in visibility was remotely comparable in size to the change of 2001.

Some factors may have contributed to a gradual increase in income visibility. For example, the media added convenient and engaging ways to browse tax records. One newspaper released an app that connected to Facebook and automatically created leaderboards showing the highest and lowest earners among Facebook friends, as shown in Figure 1.c.[11] Another application allowed users to tap on a map to see the incomes of everyone living near that position.[12] Just like the websites, these smartphone apps were incredibly popular (Digi, 2009; Teknologiradet, 2010).[13] There was also a modest increase in Internet access during the 2000s, which may have contributed to higher income visibility: according to Statistics Norway, the share of Internet users increased from 72.8% in 2002 to 95.1% in 2013.

On the other hand, some government regulations may have decreased the degree of income transparency. From 2004 to 2006, regulators introduced restrictions to the use of the tax lists: visitors had to use an official search tool conduct searches, which was only available during three weeks of the year (Teknologiradet, 2010). The official search tool was easy to use, and the newspapers seamlessly embedded it in their own websites. The three-week restriction may not

[9] In the day that the tax records are released, a few dozen Finish journalists line up in the tax agency to look up the incomes of some news-worthy individuals. However, that number pales in comparison to the millions of searches conducted in Norway every year. One exception of the Finish law is that the tax records are searchable online for the top 10,000 richest individuals. Also, there was a period in which requests for tax records could be done over the phone – however, to the best of our knowledge, this option was not nearly as widespread as the online searches were in Norway.

[10] In 2006, a credit reporting company called Ratsit published a website with a search tool for tax records similar to the ones offered in Norway (The Local, 2015). However, it was taken down by the Sweden's Chancellor of Justice shortly thereafter. The website was later allowed, but the searches were non-anonymous and subject to a fee. In 2015, this same company began selling physical copies of the tax records at the municipality level, just like in some Norwegian municipalities prior to 2001.

[11] Also, the top-right corner of Figure 1.a shows an advertisement for one of these apps.

[12] Appendix Figure A.1.b shows an screenshot of one of these apps.

[13] In addition to showing individual records, some of the websites and apps offered tools to navigate aggregate data. For example, Appendix Figure A.1.c shows the screenshot of one of the websites offering an interactive tool to figure out the user's position in the income distribution of the country or a city (Dagens Næringsliv, 2014). Also, the data published on the tax lists were eventually indexed by all the popular search engines. As a consequence, searching for the name of a Norwegian citizen in Google would show the individual's tax record at the top of the search results (Teknologiradet, 2010).

have had significant effects either, as individuals could conduct the same number of searches, just in a concentrated period. Indeed, most searches occurred during that same three-week period even when the restriction was not in place, because the timing coincided with tax record updates. For example, about 60% of data searches published in October of 2013 were conducted during the first three weeks after the tax lists were posted (E24, 2014), even though individuals were allowed to search all year long.

The 2004 restrictions were removed in 2007, with no new restrictions added until 2011. From 2011 to 2013, the government required individuals who wanted to search tax records to log in to the official website of the tax agency using a pin-code and a password.[14] Most individuals presumably already had accounts for filing their taxes and other online services.[15] Search volume probably declined due to this added hassle, but it remained substantial (Aftenposten, 2013).

The final and most significant restriction was introduced in 2014, when the searches became non-anonymous. This change is relevant for our empirical analysis, because it occurred after the last year of the survey data. However, we discuss this regulatory change below, because it provides evidence about how individuals had been using the search lists.

2.3 Popularity of the Online Tax Lists

We use three sources of data to assess the popularity of the tax lists. The most direct evidence comes from a 2007 survey conducted by Synovate, which was representative of the population of taxpayers. Around 40% of respondents reported to have used the online search tools (Skatte Betaleren, 2008). This behavior may be under-reported in surveys because of social desirability bias. Thus, the true fraction of Norwegians using these websites may have been even larger than 40%.

Web traffic data confirm media claims about the massive popularity of the online tax lists, with one website reporting 29.4 million searches in the year after the publication of the tax records for 2007 (VG, 2008). This figure implies 7.47 searches per capita among 3,935,000 Internet users in Norway in 2007. Even if these statistics are inflated due to self-reporting by the website owners, this figure excludes traffic from other websites and smartphone apps offering access to the tax records, making the likely number of searches even higher.

There is also publicly available data from to the period when tax records were accessible only from the tax agency's official website. According to Norway's Ministry of Finance (2014),

[14] The 2011 legislation also introduced a limit on the maximum number of searches per month (500), although it seems that such restriction would not be binding for the vast majority of individuals. Also, the government still allowed the media to disseminate some information from the tax lists, such as the lists of the top-100 richest individuals or break downs of average income by county – see for example the following website, which is still functional: www.vg.no/spesial/skattelister/. For reference, Appendix Figure A.1.a shows a screenshot of the search tool from the official website of the tax agency as of 2015.

[15] However, individuals who were not registered in the tax agency, such as minors or visitors from other countries, could not log into the website.

920,896 unique users conducted slightly more than 17 million searches in 2013.[16] In that year, only adults with a valid account could log in to the official website to conduct searches. Among the 3,797,822 adults in Norway, about 24.25% searched for at least one tax record in 2013, and the average user made 18.46 searches.

The statistics reported for 2007 and 2013 are not directly comparable to each other, because they come from different sources and are probably based on different definitions. With that caveat, the number of individuals conducting searches and the number of searches per capita decreased from 2007 to 2013. This difference is probably due to the 2011 requirement that users log in to the official tax agency website to search tax records.

We also assess the popularity of the income search tool using data from Google Trends, which include the number of times that a keyword is searched in the Google search engine.[17] For the main search category, skattelister, we include searches for the two words used most often to refer to the tax records, "skattelister" and "skattelistene," which both translate literally to "tax list." For instance, one popular website with access to the tax records was www.skattelister.no. As benchmarks, we use data on two keywords that are consistently among the most popular keywords around the world: "weather" and "YouTube." As a proxy for the general interest in information about taxes, we study the number of searches for "tax."

Figure 2.a shows the popularity of selected keywords in 2010 (the last year when users could conduct searches outside of the official website of the tax agency). The left half of Figure 2.a shows the results for Norway. Google Trends does not provide information about the absolute number of searches, so the search totals are normalized as a fraction of YouTube searches.[18] The data suggest a remarkable interest in the tax lists: for every five searches for YouTube, there was about one for skattelister. Norwegians were more likely to search for the tax records than to search for the weather. Searches for the tax lists were roughly three times higher than those for taxes, suggesting that a general interest in taxes does not explain the popularity of the search tool. As a robustness check, the right half of Figure 2.a provides comparable search data for Sweden, where there is no reason for individuals to search for skattelister. The volumes of searches for weather, taxes, and YouTube are roughly similar between Norway and Sweden, but searches for skattelister are virtually nonexistent in Sweden.

Figure 2.b shows the distribution of Google searches over the course of each week of 2010. Search volumes are normalized so that searches in all categories sum up to 1 in the first week of 2010. During most of the year, searches for the tax lists remained stable at roughly twice

[16] The statistics that have been reported for 2011 and 2012 are similar in magnitude to those for 2013: Bergens Tidende (2014a) reports over 700,000 unique visitors making over 13 million searches in 2011, and over 900,000 unique visitors making over 16.5 million searches in 2012.

[17] The data can be accessed using the following URL: trends.google.com. For a discussion of the advantages and limitations of this type of data, see Stephens-Davidowitz (2014).

[18] As additional benchmark, the number of searches for Youtube are slightly higher than the combined searches for "porn" and its Norwegian translation, "porno."

the volume of searches for taxes and at about the same as weather-related searches. In the third week of October, when data from the previous tax calendar year were released, searches for the tax lists increased sharply.[19] During that week, the number of searches for the tax lists exceeded the number of searches for YouTube, suggesting that Norwegians were more interested in learning about others' incomes than in watching videos on YouTube.

2.4 Uses of the Online Tax Lists

This section presents some evidence that the online tax lists were being used primarily to snoop on social contacts.

Perhaps the best piece of evidence comes from the regulatory change that took place in 2014, when searches for tax records stopped being anonymous. Specifically, any individual could use the same website to identify who searched for their tax records. This non-anonymity should have discouraged individuals from unsavory uses of the tax records, such as snooping, due to the threat of social sanctions. Consistent with this hypothesis, the tax agency reported that the number of searches dropped by 88% after the removal of anonymity. Furthermore, the number of users logging in to the system did not decrease much; however, instead of searching for others' incomes, most users logged in to find out who searched for them.[20]

The aforementioned 2007 Synovate survey offers more direct evidence about the uses of the tax lists (Skatte Betaleren, 2008). The survey asked whether respondents searched for specific types of individuals: 61% reported searching for close relatives, 53% for themselves, 42% for friends, 26% for work colleagues, 25% for other relatives, 23% for neighbors, 18% for celebrities, and 6% for politicians. This pattern is more consistent with snooping on social contacts than investigating corruption. Indeed, around 77% of respondents who used the tax records reported using them for curiosity or fun and only 2% for monitoring, such as uncovering tax evasion (Digi, 2008). And in another survey conducted by Synovate in 2011, only 15% of respondents believed that the tax lists provided useful information (Sunnmørsposten, 2011).

We also present evidence based on Internet browsing behavior from panel data covering a significant share of Internet users in Norway in 2010. We focus on visitors to a popular website that provided access to the tax records. The data span 200,000 unique browser sessions with at least one visit to this website. Figure 3.a shows the distribution of total visits by the number of profile visits. The results suggest that most traffic is not directed to famous people, such as athletes and politicians, because visits to popular profiles (i.e., visited at least 100 times) account for less than 3% of total traffic. Figure 3.b provides a histogram of the number of

[19] We find a consistent peak in Google Trends data for other years. This peak is also consistent with the Internet browsing data discussed below.

[20] Some individuals started selling a search service under their names to allow users make anonymous searches, although the service has not met popular demand (Bergens Tidende, 2014b).

profiles visited per user session on the day of the release of the 2009 tax calendar data.[21] The data suggest that large-volume users, such as mass marketers, did not contribute heavy traffic to these websites. For example, users visiting more than 100 profiles per session account for only 0.27% of total visits. Figure 3.b also shows that individuals did not search only for their own incomes, as the typical session involved searching for several individuals. Even under the conservative assumption that all sessions with a single profile visit corresponded to individuals searching for their own incomes, this type of searches comprise just 2.62% of total traffic.

Last, we discuss the possibility that individuals used the online tax lists to learn information about salaries for salary negotiations and career choices (Cullen and Pakzad-Hurson, 2018). In the previously mentioned survey data, only 26% of individuals searched for work colleagues in the tax lists, and they may have been snooping rather than researching. Moreover, due to the nature of the data, the Norwegian tax lists have been described as "completely useless" for salary comparisons (NRK, 2008). As a benchmark, the website of state employees studied in Card, Mas, Moretti, and Saez (2012) publishes information about salaries, with breakdowns by base salary and other forms of compensation. In contrast, the Norwegian tax records reveal the net income of the individual, which aggregates all salaried income, including bonuses and commissions, and non-salaried income, such as capital gains, self-employed income, and social benefits. Thus, if you found out that a coworker was listed in the tax records with a higher net income than yours, you would not be able to tell whether that coworker has a higher salary or whether he or she has additional sources of income.

3 Econometric Specification and Survey Data

3.1 Econometric Specification

The baseline specification is the following:

$$SWB_{i,t} = \alpha_1 \cdot IncomeRank_{i,t} + \alpha_2 \cdot IncomeRank_{i,t} \cdot I_t^{01-13} + X_{i,t}\beta + \delta_t + \epsilon_{i,t} \tag{1}$$

$SWB_{i,t}$ is a measure of subjective well-being of individual i in year t, with a higher value denoting higher well-being. $IncomeRank_{i,t}$ is the position of individual i in the national distribution of household income in year t, from 0 (the poorest household) to 1 (the richest). I_t^{01-13} is a dummy variable indicating the period of higher income transparency, equal to 1 if $t \in [2001, 2013]$ and 0 otherwise. δ_t denotes the year effects, $X_{i,t}$ is a vector with additional control variables, and $\epsilon_{i,t}$ denotes the error term.

The coefficient α_1 corresponds to the average gradient between $SWB_{i,t}$ and $IncomeRank_{i,t}$

[21] This is a lower bound on the number of profiles visited per user that day, for instance, because one may have visited the website from multiple devices.

between 1985 and 2000. We expect this coefficient to be positive, meaning that being richer is associated to higher subjective well-being. This association may arise purely from intrinsic utility from consumption (e.g., richer individuals can afford nicer houses, food and entertainment), or from a combination of intrinsic utility and income comparisons (e.g., richer individuals get higher self-esteem and social-esteem). The coefficient α_2 measures the change in the happiness-income gradient from 1985–2000 to 2001–2013. Our main hypothesis is that α_2 is positive: i.e., by facilitating income comparisons, the higher transparency increased the happiness-income gradient.

This regression has a differences-in-differences interpretation in which I_t^{01-13} corresponds to the indicator of post-treatment period and $IncomeRank_{i,t}$ corresponds to the intensity of treatment (from 0 to 1). An important concern with this specification, as in every other differences-in-differences design, is the possibility of differential pre-trends. In other words, it is possible that the happiness-income gradient had been gradually increasing even before 2001, yielding $\alpha_2 > 0$, even if there was not a discontinuous change in this gradient around 2001. The following specification is a traditional way of addressing this concern, by allowing for differential trends:

$$\begin{aligned} SWB_{i,t} = \alpha_1 \cdot IncomeRank_{i,t} + \alpha_2 \cdot IncomeRank_{i,t} \cdot I_t^{01-13} + \\ + \gamma \cdot IncomeRank_{i,t} \cdot (t - 1985) + X_{i,t}\beta + \delta_t + \epsilon_{i,t} \end{aligned} \tag{2}$$

In this specification, the coefficient α_1 corresponds to the happiness-income gradient in 1985. The coefficient γ corresponds to the linear trend for this gradient from 1985 to 2013. And the coefficient α_2 corresponds to the change in the happiness-income gradient around 2001, above and beyond the linear trend.

Another standard method to assess differential pre-trends is based on the following specification:

$$\begin{aligned} SWB_{i,t} = \alpha_1 \cdot IncomeRank_{i,t} + \alpha_2 \cdot IncomeRank_{i,t} \cdot I_t^{01-13} + \alpha_3 \cdot IncomeRank_{i,t} \cdot I_t^{97-00} + \\ + X_{i,t}\beta + \delta_t + \epsilon_{i,t} \end{aligned} \tag{3}$$

Where I_t^{97-00} is a "fake" treatment indicator that occurs just before the actual change in disclosure: i.e., a dummy variable that equals 1 if $t \in [1997, 2000]$ and 0 otherwise. In this specification, α_1 corresponds to the happiness-income gradient from 1985–1996, whereas α_2 measures the change in that gradient from 1985–1996 to 2001–2013, and α_3 measures the change in the happiness-income gradient from 1985–1996 to 1997–2000. If the happiness-income gradient changed sharply around 2001, we would expect $\alpha_2 > 0$ and $\alpha_3 = 0$. We also present event-study graphs, which extend this specification by including further interactions with I_t^{89-92}, I_t^{93-96} and so on.

If the happiness-income gradient increased in 2001, it is possible that this increase was caused by another significant change besides transparency that occurred in 2001 and persisted over the following twelve years. To the best of our knowledge, there was no major event around that time that could have had such a large effect on the happiness-income gradient. We use two additional identification strategies to address this concern.

The first strategy consists of a "placebo" analysis that reproduces the regressions for another country (i.e., Germany), for which there is similar survey data but no change of disclosure around 2001. If the effects in Norway are due to an event that also happened in Germany, such as the growth of information technology or the dot-com burst of 2001, then the results in Germany should be similar to the results in Norway.

The second strategy consists of a triple-differences specification. Ideally, we would construct a variable indicating the type of individuals who would be most exposed to the effects of online tax lists. This exposure variable would identify individuals who are likely to search for themselves, to be searched for by their social contacts, to be aware that their social contacts are searching for them, and so on. Then, we could test if the well-being effects are stronger for these individuals. Unfortunately, we cannot construct this ideal exposure variable, because our survey data do not contain information such as whether a respondent visited the tax list websites. Instead, we construct our exposure variable based on Internet access data.

Let the dummy variable $HigherInternet_{i,t}$ take the value 1 if individual i's observable characteristics in year t, such as the age and education, predict above-median Internet access at home.[22] Consider the following triple-differences specification:

$$\begin{aligned} SWB_{i,t} = {} & \alpha_1 \cdot IncomeRank_{i,t} + \alpha_2 \cdot IncomeRank_{i,t} \cdot I_t^{01-13} + \alpha_3 \cdot HigherInternet_{i,t} + \\ & + \alpha_4 \cdot HigherInternet_{i,t} \cdot I_t^{01-13} + \alpha_5 \cdot IncomeRank_{i,t} \cdot HigherInternet_{i,t} + \\ & + \alpha_6 \cdot IncomeRank_{i,t} \cdot HigherInternet_{i,t} \cdot I_t^{01-13} + X_{i,t}\beta + \delta_t + \epsilon_{i,t} \end{aligned} \tag{4}$$

The coefficient α_2 is interpreted as the effect of the policy on individuals with lower Internet access, which we expect to be small or even zero. On the other hand, the parameter α_6 measures the differential effect of transparency for individuals with higher Internet access, relative to individuals with lower Internet access. Our main hypothesis is that $\alpha_6 > 0$: i.e., the change in disclosure had a greater effect on individuals with higher Internet access.[23]

[22] We cannot base the triple-differences strategy on a dummy variable for whether the respondent has Internet access or not. First of all, the question about Internet access was not added until 1999. Most important, the the share of individuals with Internet access has increased dramatically in the sample period. For example, a small share of the population had Internet access before 1996: according to the World Telecommunication Development Report, only 6.4% of Norwegians had Internet access in 1995. As a result, even if we had data on Internet access for that period, it would make little sense to estimate the happiness-income gradient for individuals with Internet access.

[23] Moreover, we can also use these two coefficients to predict the effects on the happiness-income gradient for individuals with higher Internet access ($\alpha_2 + \alpha_6$) and for the population at large ($\alpha_2 + \frac{1}{2}\alpha_6$).

3.2 Survey Data

We use data from the Norwegian Monitor Survey, which was a repeated cross-sectional survey conducted by the market research institute Ipsos MMI. The data were collected every other year in 1985–2013 through a self-completion questionnaire sent by mail to a representative sample of Norwegians. This dataset has been used to explore the relationship between well-being and age (Hellevik, 2002), between well-being and values (Hellevik, 2003), and between well-being and sustainability (Hellevik, 2015).

The final sample used in our regression analysis comprises 48,570 observations collected in 15 different years, implying an average of 3,238 observations per survey year. This sample seems to be representative of the general population in some observable characteristics. For example, in the year 2011, 53.0% of respondents were women, the median age was 37, and the mean gross household income was $129,684 (in 2011 U.S. dollars). In comparison, administrative data from Norway for that same year suggest a share of women of 50.5%, a median age of 39.1, and a mean gross household income of $152,890.[24]

The survey team did not collect information about the date when each survey was completed or mailed back, but they believe that the questionnaires were completed between late September and early December.[25] Recall that the tax agency releases the income data for the previous fiscal year in mid-October. In the weeks following the data release, traffic to the online tax lists is highest. Thus, a substantial share of the respondents may have completed the survey during a time when income transparency was most salient. Our estimated effects of income transparency thus may overestimate the effects of income transparency on an average day of the year. On the other hand, a significant share of survey responses for 2001 may have been collected before the change in disclosure took place, thus leading to an under-estimation of the effects of disclosure during the year 2001.

We discuss below the definitions of the main variables:

Subjective Well-Being. The main outcome of interest is subjective well-being. The Norwegian Monitor Survey includes questions about happiness and life satisfaction, which are the two most widely used measures of subjective well-being (Easterlin, 2004; Kahneman and Deaton, 2010). The happiness question is: "Will you mostly describe yourself as: Very happy; Quite happy; Not particularly happy; Not at all happy." The life satisfaction question is, "How satisfied are you with your life? Very satisfied; Somewhat Satisfied; Neither satisfied nor dissatisfied; Slightly dissatisfied; Very dissatisfied." Happiness and life satisfaction are known as evaluative measures of well-being, because answering them requires respondents to think about

[24] The data sources are: Central Intelligence Agency's World Factbook for the female share and median age, and the Euromonitor's World Consumer Income and Expenditure Patterns for the mean household income.

[25] According to private communications with the administrators of the survey, the questionnaires were typically sent to the respondents in the third week of September (following a national or local election), and the vast majority of surveys are sent back before the second week of December.

their lives in general.[26] It is well established that evaluative measures do not vary over the days of the week, are significantly correlated with income, and remain correlated with income even at high levels of income (Kahneman and Deaton, 2010). We use the happiness question in our baseline regressions because it was asked in all survey waves from 1985 to 2013, whereas life satisfaction was asked starting in 1999.

In the baseline specification, instead of arbitrarily assigning values 1, 2, 3, or 4 to the four possible answers to the happiness question, we employ the Probit-OLS method to assign these values (van Praag and Ferrer-i-Carbonell, 2008). By construction, a higher value denotes higher happiness. We use this method with all subjective questions, including life satisfaction. Moreover, to facilitate the interpretation of the regression coefficients, we standardize all subjective outcomes to a mean of 0 and standard deviation of 1. Table 1 summarizes the definitions for happiness, life satisfaction, and all other main variables in the analysis. Table 2 provides the corresponding descriptive statistics.

Although subjective well-being measures have some well-documented limitations, a growing body of evidence indicates that they contain significant information about the individual's true well-being. Subjective well-being is positively correlated to objective measures of well-being, such as emotional expressions (Sandvik et al., 1993), aggregate suicide rates (Di Tella, MacCulloch and Oswald, 2003), and activity in the pleasure centers of the brain (Urry et al., 2004). Subjective well-being also positively correlates with decision utility. For instance, Benjamin, Heffetz, Kimball, and Rees-Jones (2012) conducted a survey in which subjects were shown pairs of hypothetical scenarios with tradeoffs between two aspects (e.g., higher income versus longer workdays). They showed that, despite some deviations, most respondents choose a scenario that maximizes life satisfaction (see also Benjamin, Heffetz, Kimball, and Rees-Jones, 2014). Similarly, Perez-Truglia (2015) showed that the expenditure choices predicted by life satisfaction data are largely consistent with the actual expenditure behavior of the same individuals.

Income Rank. The variable *Income Rank* is the position of the respondent in the distribution of household income for the current year.[27] As is typical in household surveys, respondents

[26]The alternative to evaluative measures are hedonic measures, which are assessed by asking about the presence of various emotions in the experience of yesterday (e.g., happiness, sadness, worry), and they often have different correlates than evaluative measures (Deaton and Stone, 2013). Unfortunately, our survey data do not include hedonic measures.

[27]While this measure is based on the rank in the national income distribution, individuals probably care the most about the comparison to narrower reference groups such as their relatives, friends, neighbors and coworkers (Clark and Senik, 2010). We cannot construct measures of *Income Rank* based on these more specific reference groups because we do not know who the respondent's relatives, friends or other social contact are. Due to this source of measurement error, our results may under-estimate the importance of income comparisons. Similarly, since the tax records disclosed the wealth of individuals, it is likely that individuals also engaged in wealth comparisons. This is another source of measurement error that may lead us to under-estimate the overall importance of (income and wealth) comparisons.

were asked about their annual gross household income using bins.[28] This question provides no information to rank households within a particular income bin and year. To ameliorate this measurement error, we follow the standard imputation method from the literature (e.g., Stevenson and Wolfers, 2008; Kahneman and Deaton, 2010), using information on other household characteristics that correlate with income (e.g., education, age, county) to break ties within income bins.[29]

Higher Internet. The goal of this variable is to split individuals by whether their observable characteristics are associated with higher Internet access (once the Internet becomes available). We base this exercise on the dummy variable *Internet Access*, which equals 1 if the individual has Internet access at home and 0 otherwise. Using survey responses for 2001, we estimate an OLS regression of *Internet Access* on a series of observable characteristics: age, age squared, and dummy variables for gender, education, marital status, household size, and number of working household members. Appendix Table A.8 reports the results from this auxiliary regression. The coefficients suggest that individuals with higher Internet access are, on average, more likely to be male, educated, and young, and their households are likely to be larger with more working members. These correlations are largely consistent with the correlations reported in other studies of Internet access and Internet use in developed countries (File and Ryan, 2013). We use the estimated coefficients to predict *Internet Access* for the entire survey sample. The dummy variable *I{Higher Internet}* equals 1 if the individual's own predicted Internet access exceeds the median for the current year.[30]

Perceived Income Rank and Income Adequacy. Starting in 1993, the survey included a subjective question about self-perceived income rank: "In comparison to other Norwegians, would you say that your economic situation is...? Much worse than average; Slightly worse than average; Average; Slightly better than average; Much better than average." We construct the variable *Perceived Rank* using responses to this question, which are coded with the Probit-OLS method and then standardized. By definition, higher values of this variable denote a higher perceived rank. For an additional test of the self-perceptions channel, we use data on another question that was added to the survey in 1993: "How do you feel about your economic situation? Do you really need more money than you have to be able to live a satisfying life, do you manage with your current income, or would you be able to cope with less if you had to?" The possible answers are "I need more money," "I manage with what I have," and "I could cope with less."

[28] We drop 6% of the sample corresponding to individuals who did not respond the income question.

[29] The first step of this procedure consists in estimating, for each year, an interval regression of the logarithm of income on dummies for gender, education, marital status, age, number of household members and county. The second step consists of using the estimated parameters to predict the logarithm of income for each individual, conditional on belonging to the reported income bracket. We can then construct *Income Rank* by ranking individuals based on their predicted household income.

[30] This definition guarantees that the distribution of *I{Higher Internet}* will be stable over time: i.e., in any given year, half of the sample has *I{Higher Internet}*=1 and the other half has *I{Higher Internet}*=0.

We construct the variable *Income Adequacy* using the Probit-OLS method and then standardize it so that its mean is 0 and standard deviation is 1. Higher values of this variable indicate that one's income is more adequate.

Control Variables. We include a standard set of control variables used in studies of subjective well-being: age, age squared, and dummies for gender, education, marital status, total number of household members, and number of working household members.

German Data. For the placebo test, we employ data from the German Socio-Economic Panel survey, collected every year from 1985 to 2013.[31] The data do not include a question on happiness but do include a question on life satisfaction: "How satisfied are you with your life, all things considered?" Responses are measured on a 11-point scale ranging from "Completely dissatisfied" (0) to "Completely satisfied" (10). We code and standardize this outcome using the same method as for the Norwegian data. Moreover, we reproduce the same regression specification for Germany, including all the same control variables and the same procedure to create *I{Higher Internet}*.[32] The final number of observations in Germany (107,906) is more than twice that of the Norwegian Monitor Survey (48,570).

4 Results

4.1 Effects on the Happiness-Income Gradient

Table 3 explores the effects of the change of disclosure on the happiness-income gradient. The dependent variable in column (1) is *Happiness*. This column uses the simplest specification from equation (1). The estimated coefficient on *Income Rank* (0.311) is positive, precisely estimated and statistically significant (p-value<0.001). This coefficient implies that during 1985–2000, going from the lowest to the highest income rank in Norway was associated with an increase in happiness of 0.311 standard deviations. This happiness-income gradient is in the same order of magnitude as the corresponding gradients reported in other studies.[33]

[31] To maximize power, we use all the years available in the German data – Appendix A.1 shows that the results are robust if we focus on responses on odd years, like in the Norwegian survey.

[32] There are a couple of differences between the German and Norwegian surveys. In Germany, we have to include dummies for years of education instead of the dummies for levels of educational attainment used in Norway. And while in Norway we use *Internet Access* for the year 2001, that question was not included in Germany in 2001 so we have to use the responses for the year 2002 instead. Last, we restrict the German data data to household heads in West Germany.

[33] For example, results reported in Table 2 from Stevenson and Wolfers (2008) suggest that, using data for a number of countries from the World Values Survey, the ordered probit regression of happiness on the logarithm of household income yields a coefficient of 0.244 (SE 0.008). We can provide a direct comparison by estimating the same regression with our Norwegian data, which yields a coefficient in the same order of magnitude (0.307; SE 0.008). This gradient for Norway is higher than, and statistically different from, the corresponding gradient from Stevenson and Wolfers (2008). However, we would not expect them to be exactly equal: there is no reason to believe that Norway should be representative of the world average; additionally, these differences may be due to differences in how income and subjective well-being are measured in the two datasets.

Column (1) of Table 3 also reports the coefficient on the interaction between *Income Rank* and *I{2001–2013}*. This estimated coefficient (0.090) is positive, large and statistically significant (p-value=0.005). These findings suggest that the happiness-income gradient increased by 29% (i.e., from 0.311 to 0.401) from 1985–2000 to 2001–2013, indicating an economically significant increase.

The first concern is that the coefficient on the interaction of *Income Rank* with *I{2001–2013}* may not correspond to the 2001 change in disclosure, but instead results from a gradual change in this gradient that started years before 2001. To address this concern, column (2) of Table 3 shows results for the specification corresponding to equation (2), which includes the interaction between *Income Rank* and the time trend. The coefficient on the interaction between *Income Rank* and the time trend (-0.001) is close to zero and statistically insignificant (p-value=0.877), whereas the coefficient on the interaction of *Income Rank* with *I{2001–2013}* (0.098) remains positive and statistically significant (p-value=0.099). Indeed, we cannot reject the null hypothesis that the latter coefficient (0.098) equals the corresponding coefficient of 0.090 from column (1) (p-value=0.877).[34]

In turn, column (3) of Table 3 presents results from the specification corresponding to equation (3), which introduces the interactions of *Income Rank* with *I{2001–2013}* and *I{1997–2000}* simultaneously. The coefficient on the interaction of *Income Rank* with *I{2001–2013}* (0.090) reported in column (3) is statistically significant (p-value=0.015) and identical in magnitude to the corresponding coefficient from column (1) (0.090). On the contrary, the coefficient on the interaction of *Income Rank* with *I{1997–2000}* is close to zero (0.001) and statistically insignificant (p-value=0.975). Furthermore, in column (3) the coefficients on the interactions with *I{1997–2000}* (0.001) and *I{2001–2013}* (0.090) are statistically different from each other (p-value=0.043).

Figure 4.a takes the last specification a step further by means of an event-study analysis. This figure shows the evolution of the happiness-income gradient over the entire 1985–2013 period. Each coefficient denotes the change in the happiness-income gradient relative to 1997–2000. Thus, the coefficient on 1997–2000 is normalized to zero.

To attribute the coefficient on the interaction between *Income Rank* and *I{2001–2013}* to the effect of the disclosure policy, the happiness-income gradient should be stable during 1985–2000 and then increase after 2001. This is exactly the pattern shown in Figure 4.a. All pre-treatment coefficients are close to zero and statistically insignificant (p-values of 0.749, 0.737, and 0.612), suggesting that the happiness-income gradient remained constant during 1985–2000. And the three coefficients on the post-treatment period are positive and mostly statistically significant

[34] This is an equality test between two coefficients based on the same data but different regressions. To allow for a non-zero covariance between these two coefficients, we estimate a system of Seemingly Unrelated Regressions. In the remainder of the paper, when comparing coefficients from the same data but different regressions, we always use this method.

(p-values of 0.057, 0.024, and 0.311), suggesting that the happiness-income gradient increased after 2001 and remained high.

One challenge in interpreting these findings is that the effects may result from something other than the transparency change in 2001. To the best of our knowledge, there were no events that could explain these patterns, such as major changes to the income tax schedule or welfare benefits. To address this concern more directly, we use the triple-differences design corresponding to equation (4).

The dummy variable *I{Higher Internet}* equals 1 if the individual's characteristics, such as being younger and more educated, predict higher Internet access. For short, we refer to individuals with *I{Higher Internet}*=1 as individuals with higher Internet access. Our triple-difference strategy is then based on the assumption that the individuals with higher Internet access, such as younger and more educated individuals, are the same types of individuals who are more exposed to the effects of the online tax lists.

Several arguments support this assumption. The websites require Internet access, thus individuals with higher Internet access and Internet use are more likely to use the online tax lists, more aware that they exist, and more aware that their social contacts may be searching for them. Due to homophily, individuals with higher Internet access have social contacts who also have higher Internet access, and thus are more likely to be searched in the tax lists by their social contacts. Beyond Internet access itself, individuals with higher Internet access may be the type of individuals who care more about income comparisons. For example, Clark and Senik (2010) show that individuals with higher Internet access report that income comparisons are more important to them. As final evidence consistent with our assumption, the individual characteristics associated with higher Internet access in our data (disproportionately younger, more educated and male) also are associated with self-reported use of the online tax lists (Skatte Betaleren, 2008).

Column (4) of Table 3 reports the results from the triple-differences specification. The evidence indicates that, consistent with the hypothesis that the change in transparency caused the effects, changes in the happiness-income gradient were concentrated entirely among individuals with higher Internet access. The coefficient on the interaction between *Income Rank* and *I{2001–2013}* is close to zero (-0.004), statistically insignificant, and precisely estimated. This coefficient suggests that the happiness-income gradient did not change in 2001 for individuals with lower Internet access. The coefficient on the triple interaction between *Income Rank*, *I{2001–2013}* and *I{Higher Internet}* (0.217) is positive, large, and statistically significant (p-value=0.003). This coefficient indicates that the happiness-income gradient increased substantially for individuals with higher Internet access.[35]

[35] According to column (4) of Table 3, the average effect of the income disclosure on the happiness-income gradient is 0.104 ($= -0.004 + 0.5 \cdot 0.217$) and statistically significant (p-value=0.005). As expected, this average

Figure 4.b provides the event-study equivalent of the triple-differences identification strategy. The results show that, for individuals with lower Internet access, the happiness-income gradient was stable prior to 2001 and remained the same after 2001. For individuals with higher Internet access, the happiness-income gradient was stable prior to 2001, increased after 2001, and remained at the higher for the subsequent twelve years.

As discussed in section 2.2, during the twelve years following the change of disclosure, some factors may have increased or decreased the degree of income transparency. Due to the precision of the estimates, we cannot rule out ups and downs in the effects of the policy during the twelve years. However, based on Figure 4.b, the best guess is that the effects of transparency were stable: we cannot reject the null hypothesis that the three post-treatment coefficients (0.286, 0.260, and 0.221) are equal (p-value=0.757).[36]

As an additional robustness check, we compare the effects on happiness with the effects on life satisfaction. Both outcomes are evaluative measures of well-being, and they are normally found to be significantly correlated to income. Furthermore, despite conceptual differences, happiness and life satisfaction often are treated as interchangeable in the literature on subjective well-being (Easterlin, 2004). We thus expect similar effects of transparency across these two outcomes.

Columns (5) and (6) of Table 3 show results using *Life Satisfaction* as the dependent variable instead of *Happiness*. We are cautious when interpreting this evidence, however. Whereas *Happiness* has been available since 1983, the *Life Satisfaction* question was not added to the survey until 1999 and thus has only one year of pre-treatment data. Consequently, the standard errors for the *Life Satisfaction* regressions are almost twice as large as those for the *Happiness* regressions.

Column (5) of Table 3 presents the results for *Life Satisfaction* under the baseline specification from equation (1). The coefficient on *Income Rank* (0.585) is large and statistically significant (p-value<0.001). Indeed, this gradient is larger than the corresponding gradient for *Happiness* (0.311, from column (1)). This difference in gradients is not unreasonable, given that the two questions are supposed to measure somewhat different aspects of well-being and even use different scales. Most important, column (5) of Table 3 shows that the coefficient on the interaction between *Income Rank* and *I{2001–2013}* (0.122) is positive and statistically significant (p-value=0.026). These estimates imply that higher income transparency increased the life satisfaction–income gradient by 21%, from 0.585 to 0.707. Moreover, we cannot reject the null hypothesis that the 21% increase in the life satisfaction–income gradient equals the corresponding 29% increase in the happiness–income gradient (p-value=0.645).

For the sake of completeness, column (6) reports the triple-differences specification for the *Life Satisfaction* outcome. The key coefficient on the triple interaction between *Income Rank*,

effect is close to the average effect reported in the baseline specification (0.090, from column (1) of Table 3).

[36] See Appendix A.1 a more detailed discussion, including a more disaggregated event-study graph.

I{2001–2013}, and *I{Higher Internet}* (0.168) is positive, large, and statistically indistinguishable from the corresponding *Happiness* coefficient (0.217, from column (4)). However, due to the lower precision of the *Life Satisfaction* results, the coefficient is statistically insignificant (p-value=0.195).

The next robustness check involves estimating the same regressions in a placebo country, Germany, which was not exposed to the Norwegian change in income transparency of 2001. A similar change in the happiness-income gradient in Germany around 2001 would imply that the effects for Norway must be explained by a factor different than the change in disclosure. In turn, finding no effects in Germany would rule out any shocks that began in 2001 and were common between Norway and Germany, such as the dot-com burst or the growth of information technology.

Columns (7) and (8) of Table 3 report the results for Germany, using *Life Satisfaction* as the dependent variable. Column (7) reports the results under the basic specification from equation (1). The results indicate that, prior to 2001, the gradient between income and life satisfaction was similar between Germany and Norway: the coefficient on *Income Rank* (0.496, from column (7)) is statistically significant (p-value<0.001) and in the same order of magnitude as the corresponding coefficient for Norway (0.585, from column (5)).[37] Most important, column (7) indicates that, unlike in Norway, there was no significant change in the life satisfaction-income gradient in Germany around 2001. The coefficient on the interaction of *Income Rank* and *I{2001–2013}* (0.024) is close to zero, statistically insignificant, and precisely estimated. We can reject the null hypothesis that this coefficient for Germany (0.024, from column (7)) equals the corresponding coefficient estimated in Norway (0.090, from column (1)), with a p-value of 0.030.[38] These findings are consistent with Figure 4.c, which reproduces the event-study analysis for Germany (equivalent to Figure 4.a for Norway). In Norway, the happiness-income gradient was stable prior to 2001, then increased, and remained at the higher level. On the contrary, the life satisfaction-income gradient in Germany was stable prior to 2001 and remained at the same level after 2001.

Column (8) of Table 3 reports the results for Germany under the triple-differences specification. The coefficient on the triple interaction between *Income Rank*, *I{2001–2013}*, and *I{Higher Internet}* (-0.001) is close to zero, statistically insignificant, and precisely estimated. This coefficient indicates that in Germany, there was no differential change in the life satisfaction-income gradient between individuals with lower versus higher Internet access. Indeed, we can confidently reject the null hypothesis that the coefficient for Germany (-0.001, from column (8))

[37] The difference between these coefficients is statistically significant (p-value=0.042) – however, modest differences should be expected because these are two different countries and also because the life satisfaction and income questions are elicited in different ways.

[38] We can also reject the null hypothesis that the coefficient from Germany (0.024, from column (7)) is equal to the corresponding coefficient reported in column (5) (0.122) for Norway (p-value=0.090).

equals the corresponding coefficient for Norway (0.217, from column (4)), with a p-value of 0.009. These results are consistent with those in Figure 4.d, which reproduces the event-study analysis for Germany by Internet access (equivalent to Figure 4.b for Norway). As in Norway, the gradient in Germany between well-being and *Income Rank* was stable prior to 2001 for individuals with higher and lower Internet access. In Norway, these two gradients diverged after 2001, whereas in Germany they continued to be similar after 2001.

Last, these results do not address the effect of higher transparency on the average level of well-being. The happiness-income gradient may have increased because richer individuals became happier, because poorer individuals became unhappier, or a combination of the two. Appendix A.2 presents some estimates of the average effects on well-being with a differences-in-differences estimator using the exposure indicator based on Internet access. The results suggest that the change in disclosure did not have a significant effect on average happiness and life satisfaction. These findings suggest that the disclosure policy resulted in a transference of well-being from poorer to richer individuals. Also, this idea that the change of disclosure created winners and losers is consistent with data from the 2007 Synovate survey indicating that about half of the Norwegian population (46%) opposed the income transparency policy (Aftenposten, 2011).[39]

4.2 The Self-Perceptions Channel

We cannot measure all possible channels that could explain the effects of transparency on well-being, but we provide some suggestive evidence of the role of self-perceptions.

There is abundant evidence that individuals perceive themselves to be closer than they are to the middle of income distribution (Cruces, Perez-Truglia and Tetaz, 2013).[40] This middle-class bias is believed to arise due to assortativity neglect: rich people look around and see other rich people, so they incorrectly conclude that they are middle class; likewise, poor individuals see other poor people around them and believe that they are middle class.[41] Unfortunately, the question on perceived income rank included in the Norwegian Monitor Survey uses a subjective scale, so we cannot measure the middle class bias directly, as in Cruces, Perez-Truglia, and Tetaz (2013). However, we observe suggestive signs of this bias: most respondents (89.7%) believe that their incomes are slightly above, slightly below, or about average, and this tendency is true even for individuals in the tails of the income distribution.[42]

[39]This share was still similar (49%) when measured again in 2011 (Sunnmørsposten, 2011).

[40]Cruces, Perez-Truglia and Tetaz (2013) first documented this bias using data from Argentina. Since then, other studies have documented this same middle-class bias in other countries: see for example Poppitz (2016) and Bublitz (2017).

[41]Indeed, Frick, Iijima and Ishii (2018) show that, under some assumptions about payoffs and the information structure, this type of "assortativity neglect" is not only one possible equilibrium, but the unique equilibrium.

[42]Only 16% of households from the top income decile report to be much better than average, and only 25% of households in the bottom income decile report to be much worse than average. However, there are several

Under a middle class bias, a higher income transparency would increase the gradient between perceived income rank and actual income rank. That is, poorer individuals would realize that they are poorer than they thought, and richer individuals would realize that they are richer than they thought. We test this hypothesis using the same regression specification from before but with *Perceived Rank* as the dependent variable instead of *Happiness*. Columns (1) through (4) of Table 4 present the regressions with this dependent variable. These results are based on smaller time frames and sample sizes, because *Happiness* has been collected since 1985, but *Perceived Rank* has been measured only since 1993.

Column (1) of Table 4 presents the results for the simplest specification from equation (1). The coefficient on *Income Rank* (2.172) is positive, large, and statistically significant (p-value<0.001). This coefficient implies that, from 1993–2001, moving from the poorest to the richest household was associated with an increase of 2.172 standard deviations in perceived income rank. This strong gradient between perceived and actual income ranks suggests that self-perceptions of income rank were at least somewhat accurate. Most important, column (1) shows that the coefficient on the interaction between *Income Rank* and *I{2001–2013}* (0.185) is positive, large, and statistically significant (p-value<0.001). These results imply that higher income transparency increased the gradient between perceived income rank and actual income rank by 8.5% (from 2.172 to 2.357). That is, perceptions of income rank became 8.5% more accurate as a result of increased income transparency.

Column (2)–(4) of Table 4 assesses the robustness of the results with the other specifications. Column (2) adds the interaction between the time trend and *Income Rank*. In this alternative specification, the coefficient on the interaction between *Income Rank* and *I{2001–2013}* (0.135) is still positive, large, and statistically significant (p-value=0.015). Column (3), which adds the fake treatment interaction, also suggests that the results are robust: the coefficient on the interaction between *Income Rank* and *I{1997–2000}* is statistically insignificant and statistically different from the coefficient on the interaction between *Income Rank* and *I{2001–2013}* (p-value<0.001). The results reported in column (4), corresponding to the triple-differences specification, are less robust. The triple interaction between *Income Rank*, *I{2001–2013}* and *I{Higher Internet}* (0.092) is statistically insignificant (p-value=0.228). However, because the coefficient is not precisely estimated, we cannot rule out that this coefficient equals the 0.185 coefficient from the baseline specification reported in column (1) (p-value=0.267).[43]

To assess the self-perceptions mechanism further, we measure the effects on *Income Ad-*

reasons why this misalignment may not be attributed to misperceptions. In the extreme case, if individuals interpret the range from "slightly below average" to "slightly above average" as between -90% and 1000% of the average, then almost everyone would be right to pick those categories. Additionally, part of the misalignment may reflect measurement error in the actual income rank.

[43]For the sake of completeness, Appendix A.1 presents the event-study graphs for *Perceived Rank* and *Income Adequacy*.

equacy. Individuals may form their income aspirations by looking at the incomes of others. Richer individuals, who found out that they were richer than they thought, may have felt that their incomes were more adequate; poorer individuals, who learned that they were poorer than they thought, may have felt that their incomes were less adequate. To test this hypothesis, columns (5) through (8) present results with *Income Adequacy* as the dependent variable. Column (5) corresponds to the simplest regression specification. The coefficient on *Income Rank* (1.290) is positive, large, and statistically significant (p-value<0.001). This coefficient suggests that moving from the poorest to the richest household is associated with an increase of 1.290 standard deviations in the adequacy of own income. Column (5) also suggests that the higher transparency increased the gradient between *Income Adequacy* and income rank: the coefficient on the interaction between *Income Rank* and *I{2001–2013}* (0.061) is positive and statistically significant (p-value=0.083). These estimates suggest that the higher transparency increased the gradient between income adequacy and income rank by 4.7% (from 1.290 to 1.351). Indeed, this 4.7% effect is statistically indistinguishable from the 8.5% increase in the gradient between perceived and actual income ranks (p-value=0.149).

Columns (6)–(9) of Table 4 assess the robustness of the results with the other specifications. In column (6), which introduces the interaction with the linear trend, the coefficient on the interaction between *Income Rank* and *I{2001–2013}* is even larger (0.094) than in the baseline specification of column (5). However, due to the loss in precision, this coefficient is statistically insignificant (p-value=0.120). The results from column (7), which includes the fake interaction term, are mixed. On the one hand, the coefficient on the interaction between *Income Rank* and *I{2001–2013}* becomes larger (0.101) and more statistically significant (p-value=0.044). On the other hand, we cannot reject the null hypothesis that this coefficient equals the coefficient on the fake interaction (p-value=0.396). Last, the results from column (8) also are mixed: the coefficient on the interaction between *Income Rank*, *I{2001–2013}* and *I{Higher Internet}* (0.101) suggests effects that are larger than in the baseline specification (0.061, from column (5)), but due to the lack of precision, this coefficient is statistically insignificant (p-value=0.225).

The effects on well-being coincided with the effects on perceived relative income and adequacy of own income. However, this finding does not constitute definitive proof of the self-perceptions channel. Indeed, even if had found no effects on perceived relative income, the self-perceptions channel would have not been ruled out. For example, transparency may affect well-being by making self-perceptions more salient, or it may operate through other self-perceptions such as whether individuals believe that others perceive them as rich. However, the evidence from this section suggests that some effects of transparency on well-being operate through changes in self-perceptions. The change in the perceived rank and income adequacy gradients (8.5% and 4.7%) are smaller in magnitude than the changes in the happiness and life satisfaction gradients (29% and 21%). This finding suggests that the self-perceptions channel cannot fully account for effects on well-being and that additional mechanisms are at play.

4.3 Additional Robustness Checks

In this section, we provide a brief discussion of the robustness checks that, due to space constrains, are reported in Appendix A.

In the baseline specification, we code the happiness and life satisfaction questions using the Probit-OLS method. Appendix A.3 shows nearly identical results, both in terms of magnitude and statistical significance, under two alternative specifications: coding these variables with consecutive integers and using an ordered Probit model instead of OLS.

The baseline specification also assumes a linear relationship between subjective well-being and *Income Rank*. We use this linear specification because it fits the data well and, presumably for that reason, it is widely used in the literature (Ferrer-i-Carbonell, 2004). In Appendix A.4, we present results under a more flexible specification based on binned scatterplots. The findings confirm that the linear specification provides a fair approximation and that the results are not driven by outliers or non-linearities. Relatedly, Appendix A.5 compares the results to an alternative definition of *Income Rank* based on the rank within the county of residency instead of the national rank. The results are qualitatively and quantitatively consistent across the two definitions. If anything, the estimated effects of transparency are slightly larger under the local definition of *Income Rank*.

A potential concern is that *Income Rank* is measured with a survey question, which may introduce measurement error.[44] This concern is probably minor, because survey measures of income correlate highly with their counterparts from the administrative records (Karadja, Mollerstrom, and Seim, 2017). Moreover, measurement error in *Income Rank* does not necessarily challenge the validity of our findings. To explain our findings, the measurement error must have experienced a large, sudden, and permanent reduction in 2001. Moreover, this reduction must be present for individuals with higher Internet access but not for individuals with lower Internet access. Given that individuals with higher and lower Internet access answered the same question about income, this confounding factor seems unlikely.

To address these concerns more directly, Appendix A.6 presents results under alternative definitions of *Income Rank*. As the income question is elicited in bins, our baseline specification uses the standard method from the happiness literature to impute the values of *Income Rank* within each bin (Stevenson and Wolfers, 2008; Kahneman and Deaton, 2010). In Appendix A.6, we show that the results are robust under the non-imputed version of *Income Rank*. This result should not be surprising, because the imputed and non-imputed versions correlate highly with each other (correlation coefficient of 0.984). A related potential concern is that the ninth

[44] Ideally, we would like to merge the survey responses to the income data from the administrative records. Unfortunately, this is not possible because the Norwegian Monitor Survey does not collect individual identifiers.

bin in the income question was added in 1999, which could contaminate the comparison of the happiness-income gradient around 2001. Appendix A.6 mitigates this concern: the results are almost identical if we pool the ninth and eight bins, as if the ninth bin was never introduced. This result should not be surprising either, because only 1.55% of individuals fell in the ninth bin in 1999.

Appendix A.7 assesses the robustness of the results under alternative definitions of *I{Higher Internet}*: using responses to *Internet Access* for 1999 or 1999–2013 (instead of 2001, as in the baseline specification), dividing the two groups by the median for the whole sample (instead of year-specific medians), and using a Probit model (instead of OLS). The results are quantitatively and qualitatively robust for all definitions. Moreover, we use the same method to construct *I{Higher Internet}* in Germany as in Norway. As a result, if the definition of *I{Higher Internet}* generated spurious effects in Norway, it should have introduced the same spurious effects in Germany.

Another potential concern is that the change in the happiness-income gradient is mechanically driven by an increase in income inequality. This possibility seems highly unlikely, given that it requires a large, sudden, and persistent increase in inequality that would be unprecedented in a developed country. Appendix A.8 addresses this concern more directly. Using survey data from the Norwegian Monitor Survey and administrative data from Statistics Norway, it shows that all measures of income inequality remain remarkably stable in Norway, not only around 2001, but during the entire 1985–2013 period. Relatedly, Appendix A.8 shows that the composition of the sample of survey respondents did not change significantly around 2001 and, consistent with that fact, the regression results are not sensitive to the introduction of sampling weights.

4.4 Interpretation of the Findings

Our evidence suggests an increase in the gradient between subjective well-being and income in 2001 that is probably due to the increase in income transparency. Although we cannot cover every mechanism that may have mediated this effect, we can briefly discuss some plausible mechanisms.

Individuals may be affected by the online tax lists because, as a result of them, they are treated better or worse by others. In other words, rich individuals may have benefited from having their incomes made public, because others recognize them as rich and treat them better (e.g., maybe agreeing to favors or dating them). In turn, poorer individuals may have been treated worse by others. This interpretation aligns with evidence showing that individuals are treated better when they wear expensive clothing (Fennis, 2008) and drive expensive cars (Doob and Gross, 1968). It also aligns with evidence suggesting that individuals will pay more for highly visible goods, such as clothing and cars, to signal their income to others (Charles, Hurst, and Roussanov, 2009).

Another possible mechanism states that individuals care directly about whether they are richer than others, because they get a psychological utility from holding that belief.[45] Richer individuals may be happier because they find out that they are richer than they thought – for example, they may form income aspirations by looking at the incomes of others. Alternatively, richer individuals may be happier by merely thinking that a social contact will find out how rich they are. Some evidence supports this interpretation: even in contexts of anonymity and privacy, individuals seem to care about how their own payoffs in the laboratory compare to the payoffs of other subjects. For example, individuals become less risk averse to avoid being ranked last (Kuziemko et al., 2014). According to brain imaging data, individuals seem to be displeased to learn that other subjects in the lab earned higher rewards (Fliessbach et al., 2007).

Although other channels unrelated to income comparisons may affect the happiness-income gradient, we do not believe that they play a significant role in explaining our findings. For example, although the higher transparency reduced tax evasion (Bø, Slemrod, and Thoresen, 2015), the magnitude of these effects (US$33 per household in 2001) are tiny relative to the magnitude of the change in the happiness-income gradient. Also, if individuals with lower pay used the tax lists to get a raise or a job with higher pay (Rege and Solli, 2015), the resulting effects on the happiness-income gradient should be small and point in the opposite direction.[46]

Regarding the external validity of the findings, Norway is different from other countries in many dimensions, and thus income transparency effects may be more or less pronounced. To assess whether Norway is an exceptional context, we exploit data from the 2006–2007 wave of the European Social Survey. Following Clark and Senik (2010), we use the question, "How important is it for you to compare your income with other people's incomes?" The possible answers ranged from 0 ("Not at all important") to 6 ("Very important"). This question was asked in Norway as well as in other 21 European countries.[47] The importance of income comparisons seems to be reasonably homogeneous across these countries, ranging from an average score of 1.95 in the Netherlands to an average score of 2.82 in Slovakia. Most important, the evidence indicates that Norway is not special in terms of income comparisons, as the average score in Norway (2.25) is close to the average score across the other 21 countries (2.30).

[45] For a discussion on the difference between external and internal benchmarks in income comparisons, see Senik (2009).

[46] As discussed in section 2.4, the tax lists are far from ideal for salary comparisons, and thus this change in behavior may only affect a small fraction of individuals. And, if anything, this channel would predict effects in the opposite direction: if lower-paid individuals are moving to higher-paying jobs, then that should reduce the gap in well-being between richer and poorer individuals.

[47] In addition to Norway, the question was asked in the following countries: Austria, Belgium, Bulgaria, Denmark, Estonia, Finland, France, Germany, Hungary, Ireland, Netherlands, Poland, Portugal, Russian Federation, Slovakia, Slovenia, Spain, Sweden, Switzerland, Ukraine and the United Kingdom.

4.5 Back-of-the-Envelope Calculations and Comparison to Related Studies

We start from the framework that income can affect happiness through two channels, intrinsic utility and income comparisons. We use the previously discussed estimates to assess the relative contribution of income comparisons, ranging from 0% (only intrinsic utility matters) to 100% (only income comparisons matter). To estimate a lower bound, we start from the worst-case scenario that the happiness-income gradient was entirely due to the intrinsic utility channel prior to 2001. Additionally, we assume that the 29% increase in the happiness-income gradient (column (1) of Table 3) was due to the income comparisons channel. Under those two assumptions, it follows that 22% ($=\frac{0.29}{1+0.29}$) of the happiness-income gradient after 2001 is due to income comparisons. As this is based on a worst-case scenario, the 22% provides a lower bound to the importance of income comparisons.[48]

We provide a less conservative lower bound by focusing on individuals with higher Internet access. We assume that the pre-2001 happiness-income gradient was entirely due to the intrinsic utility channel and that the 56% increase in this gradient (column (4) of Table 3) was entirely due to income comparisons. Under these two assumptions, it follows that among individuals with higher Internet access and after 2001, at least 36% ($=\frac{0.56}{1+0.56}$) of the happiness-income gradient can be attributed to income comparisons.

We also benchmark our results with estimates obtained in other studies. Countless studies can be used for this comparison, but we focus on a set of five studies that are used as benchmarks in Bottan and Perez-Truglia (2017). Three are based on happiness regressions using data from the United States, Germany, and Japan, respectively: Luttmer (2005), Ferrer-i-Carbonell (2005), and Clark, Senik and Yamada (2016). The other two studies are based on hypothetical trade-offs between absolute and relative incomes: Johansson-Stenman, Carlsson, and Daruvala (2002) and Yamada and Sato (2013), based on data from Sweden and Japan, respectively. Based on the key estimates reported in these studies, the role of income comparisons is estimated at 82.0% in Luttmer (2005), 49.6% in Ferrer-i-Carbonell (2005), 52.8% in Clark, Senik, and Yamada (2016), 35.0% in Johansson-Stenman, Carlsson, and Daruvala (2002), and 45.8% in Yamada and Sato (2013).[49] These estimates are in the range 35%–82% and are then consistent

[48]We are assuming the worst case scenario that income comparisons did not matter at all prior to 2001. If, instead, we were to assume that income comparisons explained 50% of the happiness-income gradient prior to 2001, then we would have concluded that 61% ($\frac{0.5+0.29}{1+0.29}$) of the post-2001 gradient was due to the income comparisons channel. Appendix B provides a simple model to formalize these back-of-the-envelope calculations.

[49]Let y be an individual's own income and $\bar{y}$ the average income in the individual's reference group. These studies are based on utility functions of the following form: $U = a \cdot log(y) - b \cdot log(\bar{y})$, with $a > 0$ and $b \in [0, a]$. We can use the ratio between the two coefficients, $\frac{b}{a}$, to measure the fraction of the utility from income that is due to income comparisons. Luttmer (2005) reports $a = 0.361$ and $b = 0.296$ in column (3) of Table 1; Ferrer-i-Carbonell (2005) reports $a = 0.456$ and $b = 0.226$ in column (1) of Table 2; Clark, Senik and Yamada (2016) report $a = 0.290$ and $b = 0.153$ in column (1) of Table 3; Johansson-Stenman, Carlsson and Daruvala

with the lower bound of 22% reported in our study.

5 Conclusions

In 2001, Norwegian tax records became easily accessible online, allowing everyone in the country to observe the incomes of everyone else. We propose that, because of income comparisons, higher income transparency can increase the differences in well-being between richer and poorer individuals. Using survey data and multiple identification strategies, we present evidence that higher income transparency caused an increase of 29% in the happiness-income gradient and an increase of 21% in the life satisfaction-income gradient. We provide evidence that some, although probably not all, of these effects operated through changes in self-perceived income rank. We also provide back-of-the-envelope calculations suggesting that income comparisons play a significant role in the relationship between well-being and income.

We conclude by discussing some implications for designing disclosure policies. We provide unique evidence that, as argued by those who opposed it, income transparency had a negative effect on the well-being of individuals with lower incomes. However, this result does not imply that transparency is bad. Alternative ways of publicizing and disseminating information can reduce these adverse consequences while preserving the desirable effects of transparency. For example, in 2014, Norway made searches of the tax records non-anonymous, which seems to have successfully leveraged social norms to discourage unintended uses of the data, such as snooping on friends.

However, in presence of strong privacy norms (Cullen and Perez-Truglia, 2018), a policy of non-anonymous searches can discourage legitimate uses of the data such as for salary negotiations and career planning. Governments may want to complement the non-anonymous search tools by offering anonymous access to de-identified datasets. For example, some U.S. 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 organization, occupation, and unit. This aggregate data can provide most of the information that individuals need while avoiding harmful effects on the well-being of the lowest earners.[50]

(2002) report $\frac{b}{a}$=0.35 in page 373; and Yamada and Sato (2013) report $a = 0.048$ and $b = 0.022$ in column (1) of Table 4.

[50]A similar recommendation is provided in Cullen and Perez-Truglia (2018). They conducted a survey of employees in a large firm with standard pay secrecy. They show that a vast majority of employees would prefer the firm to disclose anonymous salary information, such as average salaries per position. However, most employees would prefer the status quo of pay secrecy over the disclosure of salary records with personally identifiable information.

References

[1] Aftenposten (2008), "Skattelistene fører til mobbing," May 26, 2008.

[2] Aftenposten (2011), "Særnorsk praksis rundt skattelistene," October 12, 2011.

[3] Aftenposten (2013), "Åpne skattelister gir mindre juks," October 18, 2013.

[4] Balassa, B. (1964), "The Purchasing Power Parity Doctrine: A Reappraisal," Journal of Political Economy, Vol. 72 (6), pp. 584–596.

[5] Benjamin, D.J.; Heffetz, O.; Kimball, M.S. and Rees-Jones, A. (2012), "What Do You Think Would Make You Happier? What Do You Think You Would Choose?" American Economic Review, Vol. 102 (5), pp. 2083–2110.

[6] Benjamin, D.J.; Heffetz, O.; Kimball, M.S. and Rees-Jones, A. (2014), "Can Marginal Rates of Substitution Be Inferred From Happiness Data? Evidence from Residency Choices." American Economic Review, Vol. 104 (11), pp. 3498–3528.

[7] Bergens Tidende (2014a), "Nå er det slutt på usynlig snoking," October 16, 2014.

[8] Bergens Tidende (2014b), "Christina sjekker skattelistene for deg," October 17, 2014.

[9] Bø, E.E.; Slemrod, J. and Thoresen, T.O. (2015), "Taxes on the Internet: Deterrence Effects of Public Disclosure," American Economic Journal: Economic Policy, Vol. 7 (1), pp. 36–62.

[10] Boskin, M. and Sheshinski, E. (1978), "Optimal Redistributive Taxation When Individual Welfare Depends on Relative Income," Quarterly Journal of Economics, Vol. 92 (4), pp. 589–601.

[11] Bottan, N. and Perez-Truglia, R. (2017), "Choosing Your Pond: Location Choices and Relative Income," NBER Working Paper No. 23615.

[12] Bublitz, E. (2017), "Misperceptions of income distributions: Cross-country evidence from a randomized survey experiment," LIS Working Paper No. 694.

[13] Card, D.; Mas, A.; Moretti, E. and Saez, E. (2012), "Inequality at Work: The Effect of Peer Salaries on Job Satisfaction," American Economic Review, Vol. 102 (6), pp. 2981–3003.

[14] Charles, K.K.; Hurst, E. and Roussanov, N.L. (2009), "Conspicuous Consumption and Race," Quarterly Journal of Economics, Vol. 124 (2), pp. 425-467.

[15] Clark, A.E.; Frijters, P. and Shields, M.A. (2008), "Relative Income, Happiness, and Utility: An Explanation for the Easterlin Paradox and Other Puzzles," Journal of Economic Literature, Vol. 46 (1), pp. 95–144.

[16] Clark, A.E. and Senik, C. (2010), "Who compares to whom? The anatomy of income comparisons in Europe," Economic Journal, Vol. 120, pp. 573–594.

[17] Clark, A.E.; Kristensen, N. and Westergård-Nielsen, N. (2009), "Economic Satisfaction and Income Rank in Small Neighbourhoods," Journal of the European Economic Association, Vol. 7, pp. 519–527.

[18] Clark, A.; Senik, C. and Yamada, K. (2017), "When experienced and decision utility concur: the case of income comparisons," Journal of Behavioral and Experimental Economics, Vol. 70, pp. 1–9.

[19] Cruces, G.; Perez-Truglia, R. and Tetaz, Martin (2013), "Biased Perceptions of Income Distribution and Preferences for Redistribution: Evidence from a Survey Experiment," Journal of Public Economics, Vol. 98, pp. 100–112.

[20] Cullen, Z. and Perez-Truglia, R. (2018), "The Salary Taboo: Privacy Norms and the Diffusion of Information," NBER Working Paper No. 25145.

[21] Cullen, Z. and Pakzad-Hurson, B. (2018), "Equilibrium Effects of Pay Transparency in a Simple

Labor Market," Mimeo.

[22] Dagbladet (2009), "Nå får du skattesøk på iPhone," October 15, 2009.

[23] Dagens Næringsliv (2010), "Skattelister gir ikke innbruddsbølger," May 22, 2010.

[24] Dagens Næringsliv (2014), "Sjekk skattelister 2014," October 16, 2014.

[25] Deaton, A. and Stone, A. (2013), "Two Happiness Puzzles," American Economic Review: Papers & Proceedings, Vol. 103 (3), pp. 591–597.

[26] Digi (2008), "Vil fjerne skattelistene på nettet," October 16, 2008. 2008.

[27] Digi (2009), "Dette tjener de på iPhone-skattesøk," October 21, 2009.

[28] Di Tella, R.; MacCulloch, R. and Oswald, A. (2003), "The Macroeconomics of Happiness," Review of Economics and Statistics, Vol. 85 (4), pp. 809–27.

[29] Doob, A. and Gross, A. (1968), "Status of Frustrator as an Inhibitor of Horn-Honking Responses," Journal of Social Psychology, Vol. 76, pp. 213-218.

[30] E24 (2014), "Nå strammes innsyns-reglene inn," May 14, 2014.

[31] Easterlin, R.A. (1974), "Does economic growth improve the human lot? Some empirical evidence," Nations and Households in Economic Growth, Vol. 89, pp. 89–125.

[32] Easterlin, R. (2004), "The economics of happiness" Daedalus, Vol. 133 (2), pp. 26–33.

[33] Easterlin, R.; Angelescu, L.; McVey, M.; Switek, O.; Sawangfa, J. and Smith, Z. (2010), "The happiness–income paradox revisited," Proceedings of the National Academy of Sciences, Vol. 107 (52), pp. 22463-22468.

[34] Easterlin, R.A. (2017), "Paradox Lost?" Review of Behavioral Economics, Vol. 4 (4), pp. 311–339.

[35] Fennis, B.M. (2008), "Branded into submission: Brand attributes and hierarchization behavior in same-sex and mixed-sex dyads," Journal of Applied Social Psychology, Vol. 38, pp. 1993–2009.

[36] Ferrer-i-Carbonell, A. (2005), "Income and Well-Being: An Empirical Analysis of the Comparison Income Effect," Journal of Public Economics, Vol. 89, pp. 997–1019.

[37] File, T. and Ryan, C. (2014), "Computer and Internet Use in the United States: 2013," American Community Survey Report.

[38] Fliessbach, K.; Weber, B.; Trautner, P.; Dohmen, T.; Sunde, U.; Elger, C.E. and Falk, A. (2007), "Social Comparison Affects Reward-Related Brain Activity in the Human Ventral Striatum," Science, Vol. 318 (5854), pp. 1305-1308.

[39] Frank, R.H. (1985), "The Demand for Unobservable and Other Positional Goods," American Economic Review, Vol. 75, pp. 101–116.

[40] Hagerty, M.R. and Veenhoven, R. (2003), "Wealth and Happiness Revisited – Growing National Income Does Go with Greater Happiness," Social Indicators Research, Vol. 64, pp. 1–27.

[41] Hellevik, O. (2002), "Age differences in value orientation – life cycle or cohort effect?" International Journal of Public Opinion Research, Vol. 14 (3), pp. 286–302.

[42] Hellevik, O. (2003), "Economy, values and happiness in Norway," Journal of Happiness Studies, Vol. 4 (3), pp. 243–283.

[43] Hellevik, O. (2015), "Is the good life sustainable? A three decade study of values, happiness and sustainability in Norway," in Mueller, M.L. and Syse, K.V.L. (eds.), Sustainable Consumption and the Good Life. London & New York: Routledge.

[44] Johansson-Stenman, O.; Carlsson, F. and Daruvala, D. (2002), "Measuring future grandparents' preferences for equality and relative standing," Economic Journal, Vol. 112, pp. 362–383.

[45] Kahneman, D. and Deaton, A. (2010), "High income improves evaluation of life but not emotional

well-being," Proceedings of the National Academy of Sciences, Vol. 107 (38), pp. 16489–16493.

[46] Karadja, M.; Mollerstrom, J. and Seim, D. (2017), "Richer (and Holier) than Thou? The Effect of Relative Income Improvements on Demand for Redistribution," Review of Economics and Statistics, Vol. 99 (2), pp. 201–212.

[47] Kuziemko, I.; Buell, R.W.; Reich, T. and Norton, M. (2014), "Last-place Aversion: Evidence and Redistributive Implications," Quarterly Journal of Economics, Vol. 129 (1), pp. 105–149.

[48] Linden, L. and Rockoff, J.E. (2008), "Estimates of the Impact of Crime Risk on Property Values from Megan's Laws," American Economic Review, Vol. 98 (3), pp. 1103–27.

[49] Luttmer, E.F.P. (2005), "Neighbors as Negatives: Relative Earnings and Well-Being," Quarterly Journal of Economics, Vol. 120 (3), pp. 963–1002.

[50] Mas, A. (2017), "Does Transparency Lead to Pay Compression?" Journal of Political Economy, Vol. 125 (5), pp. 1683–1721.

[51] New York Times (2009), "Norway Divided by Citizen Wealth Tables," October 23, 2009.

[52] New York Times (2018), "Happy National Jealousy Day!" November 1, 2018.

[53] Norway Ministry of Finance (2014), "Høringsnotat om innstramminger i reglene om offentlige skattelister," January 24, 2014.

[54] NRK (2008), "Dette er økonomisk pornografi!" October 14, 2008.

[55] Origo (2010), "Debatten om skatten," October 20, 2010.

[56] Perez-Truglia, R. (2015), "A Samuelsonian Validation Test for Happiness Data," Journal of Economic Psychology, Vol. 49, pp. 74–83.

[57] Perez-Truglia, R. and Troiano, U. (2018), "Shaming Tax Delinquents," Journal of Public Economics, Vol. 167, pp. 120–137.

[58] Perez-Truglia, R. and Cruces, G. (2017), "Partisan Interactions: Evidence from a Field Experiment in the United States," Journal of Political Economy, Vol. 125 (4), pp. 1208–1243.

[59] Poppitz, P. (2016), "Does self-perceptions and income inequality match? The case of subjective social status," Working Paper.

[60] Rege, M. and Solli, I.F. (2015), "Lagging Behind the Joneses: The Impact of Relative Earnings on Job Separation," Mimeo.

[61] Samuelson, P.A. (1964), "Theoretical Notes on Trade Problems," Review of Economics and Statistics, Vol. 46 (2), pp. 145–154.

[62] Sandvik, E.; Diener, E. and Seidlitz, L. (1993), "Subjective Well-Being: The Convergence and Stability of Self-Report and Non-self-Report Measures," Journal of Personality, Vol. 61 (3), pp. 317–42.

[63] Senik, C. (2004), "When Information Dominates Comparison. Learning from Russian Subjective Panel Data," Journal of Public Economics, Vol. 88, pp. 2099–2133.

[64] Senik, C. (2009), "Direct Evidence on Income Comparison and their Welfare Effects," Journal of Economic Behavior and Organization, Vol. 72, pp. 408–424.

[65] Skatte Betaleren (2008), "Til Kamp Mot Skattelistene," Issue 2008–5.

[66] Stephens-Davidowitz, S. (2014), "The cost of racial animus on a black candidate: Evidence using Google search data," Journal of Public Economics, Vol. 118, pp. 26–40.

[67] Stevenson, B. and Wolfers, J. (2008), "Economic Growth and Subjective Well-Being: Reassessing the Easterlin Paradox," NBER Working Paper No. 14282.

[68] Sunnmørsposten (2011), "Flere sier nei til offentliggjøring av skattelister," May 13, 2011.

[69] Teknologiradet (2010), "Innspill til komitébehandling om offentliggjøring av skattelister," Mimeo.

[70] The Local (2015), "Salary snooping steps up in Sweden," June 9, 2015.

[71] Urry, H.L.; Nitschke, J.B.; Dolski, I.; Jackson, D.C.; Dalton, K.M.; Mueller, C.J.; Rosenkranz, M.A.; Ryff, C.D.; Singer, B.H. and Davidson, R.J. (2004), "Making a Life Worth Living: Neural Correlates of Well Being," Psychological Science, Vol. 15 (6), pp. 367–372.

[72] van Praag, B.M.S. and Ferrer-i-Carbonell, A. (2008), "Happiness Quantified: A Satisfaction Calculus Approach." Oxford: Oxford University Press.

[73] VG (2008), "I natt kommer de nye skattetallene," October 14, 2008.

[74] Yamada, K. and Sato, M. (2013), "Another avenue for anatomy of income comparisons: Evidence from hypothetical choice experiments," Journal of Economic Behavior & Organization, Vol. 89, pp. 35–57.

Figure 1: Screenshots of Websites and Smartphone Apps Designed to Browse the Tax Records

a. Search Tool from skattelister.no as of June 16, 2010

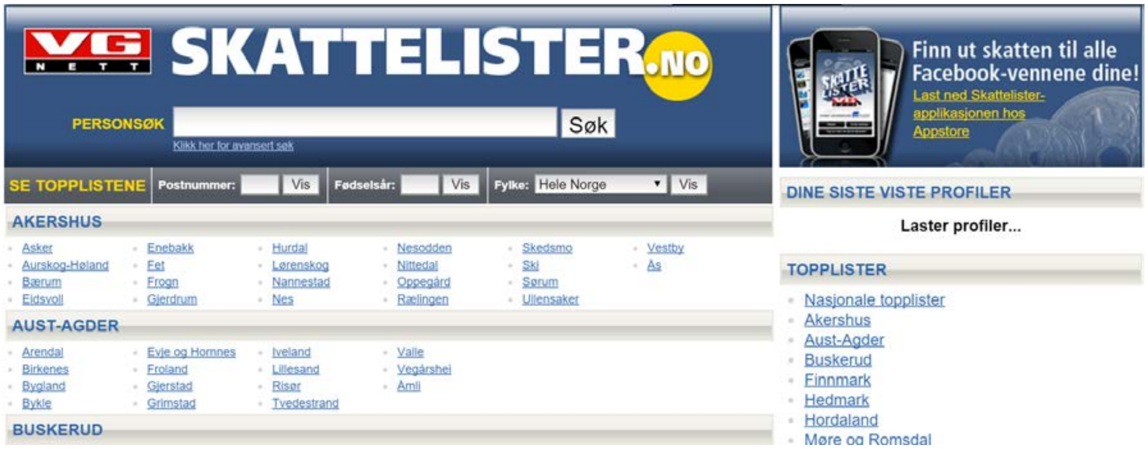

b. Search Result from skatt.na24.no as of August 1, 2015

c. Facebook Leaderboard from an Iphone Application

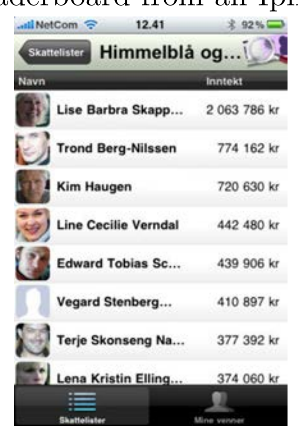

Source: (a) and (b): web.archive.org. (c) Dagbladet (2009).

Figure 2: Popularity of Online Tax Lists Measured by Google Search Data

a. Annual Search Volumes, Norway vs. Sweden

b. Weekly Search Volumes, Norway

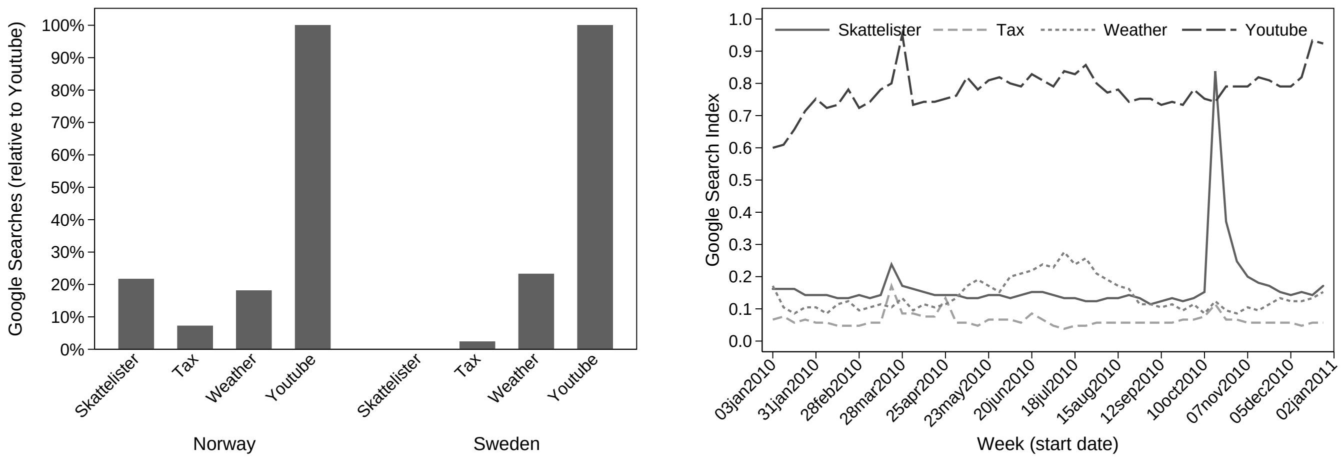

Notes: Google Trends data for 2010. Panel (a) shows the annual number of Google searches for each category of keywords, relative to the searches for "youtube." The Skattelister category is comprised by "skattelister+skattelistene", Tax is comprised by "skatt+skatter", Weather is comprised by "yr+ver" in Norway and "väder" in Sweden. Panel (b) shows the weekly number of searches for each keyword category in Norway, normalized so that total searches sum up to 1 in the first week of 2010.

Figure 3: Internet Browsing Data on the Uses of the Online Tax Lists

a. Traffic by Profile Popularity

b. Traffic by Session Length

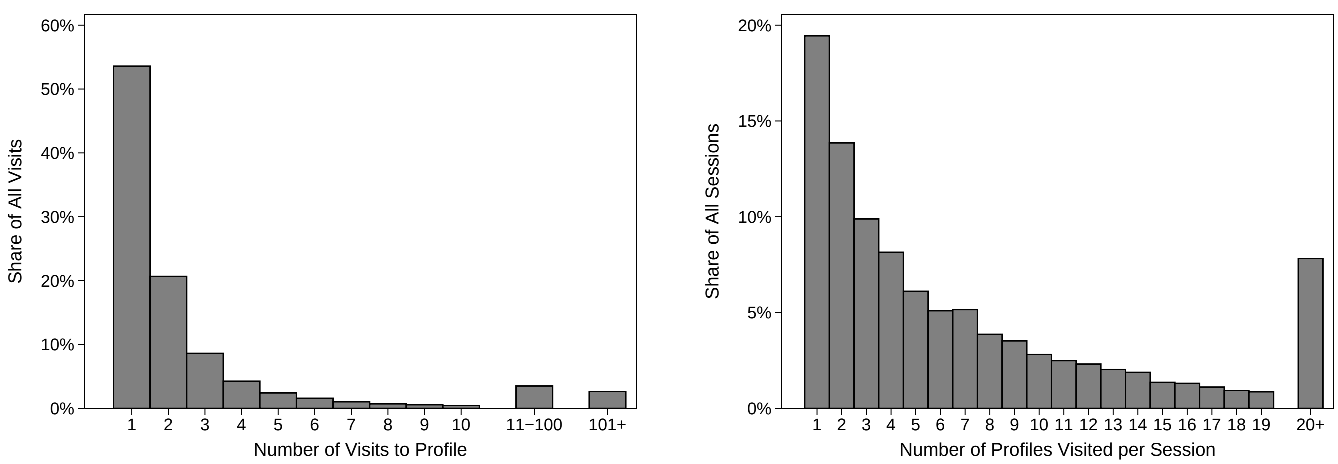

Notes: Data from a proprietary panel of Internet users, on the visits to a popular website that offered a tool to search the Norwegian tax records during 2010. Panel (a) shows the number of repeated visits to profiles during 2010 (for profiles that have been visited at least once). Panel (b) shows the number of profiles visited per session, for sessions with at least one profile visited, corresponding to October 20, 2010 (the release date for the data on incomes corresponding to the 2009 tax calendar). A session begins when the user opens the Internet browser and ends when the user closes the browser.

Figure 4: Event-Study Analysis of the Gradient between Subjective Well-Being and Income Rank

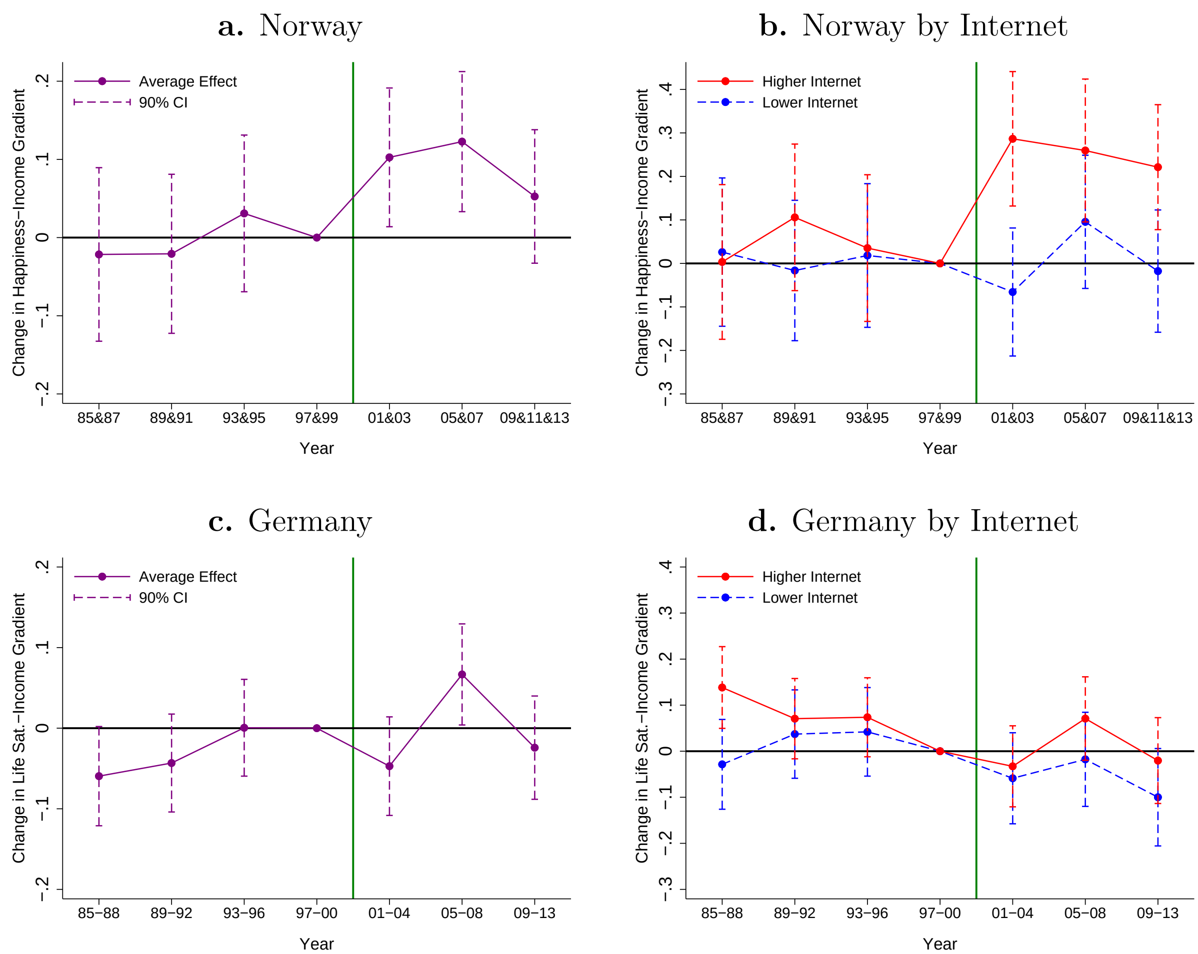

Notes: Each coefficient corresponds to the change in the gradient between subjective well-being and *Income Rank* relative to the period 1997–2000 (by construction, the coefficient for these years is normalized to zero). The green vertical line represents the change in disclosure that took place in Norway in 2001. The measure of subjective well-being in Norway is *Happiness* (panels (a) and (b) and *Life Satisfaction* in Germany ((c) and (d)). Both of these outcomes are normalized to have mean 0 and standard deviation of 1, with higher values denoting higher happiness/satisfaction. *Income Rank* denotes the respondent's household rank for that year, from 0 to 1. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size and three dummies for number of working household members. Panels (a) and (b) are based on 48,570 observations from the Norwegian Monitor Survey, collected every other year in 1985–2013 – see Table 1 for more detailed data definitions and Table 2 for descriptive statistics. Panels (c) and (d) are based on 107,906 observations from the German Socio-Economic Panel Survey, collected every year in 1985–2013. Panel (b) is similar to panel (a), only that two regressions are estimated separately for individuals with *I{Higher Internet}=1* and *I{Higher Internet}=0*. Likewise, panel (d) is similar to panel (c), only that based on two separate regressions by *I{Higher Internet}*.

Table 1: Summary of Data Definitions for the Main Variables from the Norwegian Monitor Survey

Variable Name	**Definition**
Happiness	Based on the following question: "Will you mostly describe yourself as: Very happy; Quite happy; Not particularly happy; Not at all happy." These 4 categories were assigned values using the Probit-OLS method, and then the variable was standardized to have mean 0 and standard deviation 1. Higher values denote higher happiness.
Income Rank	Estimated position in the national distribution of household gross income in a given year, from 0 (lowest income in the country) to 1 (highest). This rank is based on the following question: "What would you estimate the household's total gross income? That is, all total income before taxes and deductions: Less than NKR100,000; NKR100,000-199,000; NKR200,000-299,000; NKR300,000-399,000; NKR400,000-499,000; NKR500,000-599,000; NKR600,000-799,000; NKR800,000-999,000; Above NKR1,000,000." Before 1995, the bottom bin was split in two bins (0K–59K and 60K–100K). In 1995, they merged those two bins and added the seventh and eight bins. They added the ninth bin in 1999. The within-bin ranks are imputed using interval regressions as in Stevenson and Wolfers (2008) and Kahneman and Deaton (2010).
Life Satisfaction	Based on the following question: "How satisfied are you with your life? Very satisfied; Somewhat Satisfied; Neither satisfied nor dissatisfied; Slightly dissatisfied; Very dissatisfied." These 5 categories were assigned values using the Probit-OLS method, and then the variable was standardized to have mean 0 and standard deviation 1. Higher values denote higher satisfaction.
Perceived Rank	Based on question: "In comparison to other Norwegians, would you say that your economic situation is...? Much worse than average; Slightly worse than average; As average; Slightly better than average; Much better than average." These 5 categories were assigned values using the Probit-OLS method, and then the variable was standardized to have mean 0 and standard deviation 1. Higher values denote higher rank.
Income Adequacy	Based on the following question: "How do you feel about your economic situation? Do you really need more money than you have to be able to live a satisfying life, do you manage with your current income, or would you be able to cope with less if you had to?" The possible answers were "I need more money," "I manage with what I have" and "I could cope with less," which were assigned values using the Probit-OLS method, and then the variable was standardized to have mean 0 and standard deviation 1. Higher values denote higher income adequacy.
Internet Access	Dummy variable that takes the value 1 if the individual responds affirmatively to the following question: "Do you have Internet access at home?"
I{Higher Internet}	Dummy variable that takes the value 1 if the respondent has above-median predicted *Internet Access* in the current year. We estimated an OLS regression of *Internet Access* on age, age squared and dummy variables for gender, marital status, education, household size and number of workers in the household using the responses from 2001 (the results from this regression are presented in Appendix Table A.8). Then, for each sample year, we used the estimated coefficients to generate predicted *Internet Access* and then split the observations for that year by the corresponding median value.

Table 2: Descriptive Statistics for the Main Variables from the Norwegian Monitor Survey

Variable Name	**Availability**	**Observations**	**Mean (Std.)**
Happiness	1985–2013	48,570	0.00 (1.00)
Income Rank	1985–2013	48,570	0.50 (0.29)
Life Satisfaction	1999–2013	29,655	0.00 (1.00)
Perceived Rank	1993–2013	38,938	0.00 (1.00)
Income Adequacy	1993–2013	38,950	0.00 (1.00)
Internet Access	1999–2013	29,875	0.76 (0.43)
I{Higher Internet}	1985–2013	48,570	0.50 (0.50)

Notes: See Table 1 for a summary of data definitions for all of the variables listed above. Data from the Norwegian Monitor Survey.

Table 3: Effects on the Gradient between Subjective Well-Being and Income Rank

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
					Life	Life	Life	Life
	Happiness	Happiness	Happiness	Happiness	Satisf.	Satisf.	Satisf.	Satisf.
Income Rank	0.311***	0.315***	0.310***	0.331***	0.585***	0.526***	0.496***	0.561***
	(0.028)	(0.040)	(0.032)	(0.040)	(0.056)	(0.085)	(0.017)	(0.024)
Income Rank * I{2001-2013}$^{(i)}$	0.090***	0.098*	0.090**	-0.004	0.122**	0.050	0.024	-0.052
	(0.032)	(0.059)	(0.037)	(0.051)	(0.055)	(0.088)	(0.021)	(0.034)
Income Rank * I{2001-2013} * I{Higher Internet}				0.217***		0.169		-0.001
				(0.073)		(0.131)		(0.045)
Income Rank * (Year-1985)		-0.001						
		(0.004)						
Income Rank * I{1997-2000}$^{(ii)}$			0.001					
			(0.048)					
P-value (i)=(ii)			0.043					
Country	Norway	Norway	Norway	Norway	Norway	Norway	Germany	Germany
Period	85-13	85-13	85-13	85-13	99-13	99-13	85-13	85-13
Observations	48,570	48,570	48,570	48,570	29,655	29,655	107,906	107,906

Notes: *** p<0.01, ** p<0.05, * p<0.1. Heteroskedasticity-robust standard errors in parenthesis. Each column corresponds to a separate OLS regression. *Happiness* and *Life Satisfaction* are subjective well-being measures normalized to have mean 0 and standard deviation of 1, with higher values denoting higher happiness/satisfaction. *Income Rank* denotes the position of the respondent's household relative to all the other respondents for that year, from 0 to 1. *I{2001–2013}* takes the value 1 for 2001–2013. *I{1999–2000}* takes the value 1 for 1997–2000. *I{Higher Internet}* is a dummy variable that takes the value 1 if the respondent's predicted *Internet Access* is above the median value for a given year. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size and three dummies for number of working household members. Column (4) also controls for all the additional interactive terms listed in equation (4). Columns (1) through (6) are based on data from the Norwegian Monitor Survey, collected every other year in 1985—2013 – see Table 1 for a summary of data definitions and Table 2 for descriptive statistics. Columns (7) and (8) are based on the same specifications used in columns (1) and (3), only that using data from the German Socio-Economic Panel, collected every year in 1985—2013.

Table 4: Effects on the Additional Outcomes: Perceived Income Rank and Adequacy of Own Income

	(1) Perc. Rank	(2) Perc. Rank	(3) Perc. Rank	(4) Perc. Rank	(5) Income Satisf.	(6) Income Satisf.	(7) Income Satisf.	(8) Income Satisf.
Income Rank	2.172***	2.117***	2.130***	2.275***	1.290***	1.326***	1.249***	1.300***
	(0.032)	(0.060)	(0.047)	(0.048)	(0.035)	(0.065)	(0.050)	(0.049)
Income Rank * I{2001-2013}$^{(i)}$	0.185***	0.135**	0.228***	0.138***	0.061*	0.094	0.101**	-0.013
	(0.033)	(0.056)	(0.047)	(0.053)	(0.035)	(0.060)	(0.050)	(0.054)
Income Rank * I{2001-2013} * I{Higher Internet}				0.092				0.101
				(0.077)				(0.083)
Income Rank * (Year-1985)		0.005				-0.003		
		(0.004)				(0.005)		
Income Rank * I{1997-2000}$^{(ii)}$			0.069				0.066	
			(0.055)				(0.059)	
P-value (i)=(ii)			<0.001				0.396	
Country	Norway	Norway	Norway	Norway	Norway	Norway	Norway	Norway
Period	93-13	93-13	93-13	93-13	93-13	93-13	93-13	93-13
Observations	38,938	38,938	38,938	38,938	38,950	38,950	38,950	38,950

Notes: *** p<0.01, ** p<0.05, * p<0.1. Heteroskedasticity-robust standard errors in parenthesis. Each column corresponds to a separate OLS regression. *Perceived Rank* and *Income Adequacy* are subjective measures normalized to have mean 0 and standard deviation of 1, with higher values denoting higher rank/adequacy. *Income Rank* denotes the rank of the household income for that year, from 0 to 1. *I{2001–2013}* takes the value 1 for 2001–2013. *I{1999–2000}* takes the value 1 for 1997–2000. *I{Higher Internet}* is a dummy variable that takes the value 1 if the respondent's predicted *Internet Access* is above the median value for a given year. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size and three dummies for number of working household members. Columns (4) and (8) also control for all the additional interactive terms listed in equation (4). All regressions are based on data from the Norwegian Monitor Survey, collected every other year in 1985—2013 – see Table 1 for a summary of data definitions and Table 2 for descriptive statistics.

Online Appendix: For Online Publication Only

A Additional Results and Robustness Checks

A.1 Additional Event-Study Graphs

In this section, we present some additional event-study graphs.

Figure A.2 presents variations of the event study graphs from Figure 4. Figure A.2.a reproduces Figure 4.b, only that breaking down the coefficients at the yearly level instead of using pairs of years. Given that we are estimating twice as many coefficients with the same number of observations, each individual coefficient is less precisely estimated and thus one must be more careful with the interpretation.

As discussed in section 2.2, there were some factors during 2001–2013 that may have increased or decreased the degree of income transparency. The event-study analysis from Figure A.2.a allows for a closer look at this. We cannot reject the null hypothesis that, for individuals with higher Internet access, all the post-2001 coefficients from Figure A.2.a are equal (p-value=0.201). This evidence suggests that the effects of higher transparency were stable over time. One possible interpretation is that the events during 2001–2013 did not affect the degree of income transparency. An equally valid interpretation is that there were factors reducing and increasing the degree of transparency over time, but these positive and negative effects canceled each other out.

We can use Figure A.2.a to get a closer look at what happened in 2001. For individuals with lower Internet access, the coefficient on 2001 is close to zero and statistically insignificant. On the contrary, for individuals with higher Internet access, the coefficient on 2001 is positive, large and statistically significant. The point estimate for 2003 is higher than the point estimate for 2001, suggesting that perhaps the effects of the higher transparency took a bit more time to fully materialize.[51] However, the difference between these two coefficients is not statistically significant (p-value=0.1744). Moreover, this difference could be due to measurement error: due to the timing of the survey collection, a non-trivial share of the survey responses for 2001 may have been collected before the change in disclosure took place, which will bias the coefficient for 2001 downwards.

Figure A.2.a also allows for a closer look at what happened in 2004–2016, when the online tax lists were limited to the three-week period following the release of the data. Due to the

[51] It may take some time after the online tax lists are published for its effects on well-being to fully materialize. Consider for example the anecdote about kids being bullied at school because their classmates found out that their parents were poor through the website. Even though some kids may have been bullied the day after the publication of the online tax lists, most kids would probably not be bullied until months after the tax lists became available.

timing of our survey, this restriction is unlikely to affect our estimates. The survey was collected every two years, and thus only one of the survey years, 2005, falls in this sub-period. Moreover, in 2005 the income search tool was available during the last two weeks of October and the first week of November. Since the survey responses are collected from late September to early December, it is likely that most of the survey responses in 2005 were collected while the income search tool was available. And since the tool was not available during the rest of the year, the income search tool may have been especially salient in this period of 2005. Figure A.2.a shows a coefficient for higher Internet in 2005 that is positive and statistically significant, suggesting that the effects of higher transparency did not decline in 2005.[52]

We can use Figure A.2.a to get a closer look at what happened in 2011–2013, when the government introduced a hassle to use the tax lists: individuals had to log into the official website of the tax agency with a pin-code and a password. The coefficients for 2011 and 2013 suggest that the effects of the transparency policy persisted through 2011–2013. This finding is consistent with the evidence discussed in section 2.3, according to which the volume of searches was still substantial in 2011–2013. Moreover, even if the search volume was reduced, that should not necessarily undo the effects of transparency. Indeed, even if the online tax lists were removed altogether, it may take years for the effects on well-being to vanish. If an individual became unhappy because she found out that she is poorer than she thought, removing her access to the online tax lists will not make her suddenly forget how poor she is. Similarly, if the individual is being bullied by others, those others will not suddenly forget that the individual is poor after losing access to the website.

Figure A.2.b presents a robustness check for the analysis of the data from the German Socio-Economic Panel survey. The German survey was collected every year in 1985–2013, while the data from the Norwegian Monitor Survey was collected every odd year in 1985–2013. In the baseline specification, to maximize power, we use all the years available in the German data. As a robustness check, Figure A.2.b reproduces the German event study from Figure 4.c, only that restricting the German observations to the odd years, to mimic the frequency of the Norwegian Monitor Survey. The coefficients in Figure A.2.b are less precisely estimated because we are discarding roughly half of the data. However, the main result is still robust: Figure A.2.b shows that the life satisfaction-income gradient did not change in Germany post–2001.

Figure A.3 shows the event-study graphs for the other two outcomes used in the analysis: *Perceived Rank* and *Income Adequacy*. Note, however, that we have to be more careful when interpreting these results. First, these two questions were not included in the survey until 1993, and thus the results will be less precisely estimated than for *Happiness*. Second, the magnitude of the effects for these two outcomes (8.5% and 4.7%, respectively) are substantially smaller

[52]On the other hand, if we measured the effect of higher transparency as the gap between the coefficient for lower and higher Internet, that estimate would suggest that the effects of transparency were smaller in 2005.

than the effects for Happiness (29%), and thus these effects would be more difficult to detect even if we could hold constant the number of observations.

Figure A.3 shows that, consistent with the results presented in regression form in section 4.2, the event-study analysis for *Perceived Rank* and *Income Adequacy* are suggestive but not nearly as sharp as the results for well-being. Figure A.3.a reproduces Figure 4.b, only that using *Perceived Rank* as dependent variable instead of *Happiness*. Figure A.3.a shows that the gradient between perceived and actual income rank evolved similarly for individuals with higher and lower Internet access, but then diverged after 2001. However, this divergence is less precisely estimated. Figure A.3.b reproduces Figure 4.b, only that using *Income Adequacy* as dependent variable instead of *Happiness*. Again, the event-study graph suggests a divergence in the gradient after 2001, but this finding is not nearly as precisely estimated as the corresponding finding for *Happiness*.

A.2 Effects on the Average Level of Well-Being

Section 4.1 reports the findings for the effect of transparency on the gradient between subjective well-being and income rank. In this section, we report the findings for the effects of transparency on the average level of subjective well-being.

To estimate these average effects, we follow a differences-in-differences strategy that is based on the exposure variable discussed in section 4:

$$SWB_{i,t} = \alpha_1 \cdot HigherInternet_{i,t} + \alpha_2 \cdot HigherInternet_{i,t} \cdot I_t^{01-13} + X_{i,t}\beta + \delta_t + \epsilon_{i,t} \tag{A.1}$$

$SWB_{i,t}$ denotes subjective well-being of individual i in year t. $HigherInternet_{i,t}$ takes the value 1 if individual i's observable characteristics in year t, such as the age and education, predict above-median Internet access at home. I_t^{01-13} is a dummy variable indicating the period of higher income transparency. $X_{i,t}$ is a vector with a set of control variables. δ_t denotes the year dummies. And $\epsilon_{i,t}$ denotes the error term.

The coefficient α_1 estimates the well-being gap between individuals with higher and lower Internet access during 1985–2000, while α_2 measures the change in that gap from 1985–2000 to 2001–2013. If we assume that individuals with $HigherInternet_{i,t} = 0$ were not affected by the change in transparency while individuals with $HigherInternet_{i,t} = 1$ were affected by it, then α_2 measures the effect of the 2001 change of disclosure on the average level of well-being.[53]

The results from section 4 suggest that the effects on the happiness-income gradient for individuals with higher Internet was 0.217 (column (3) of Table 3). If that effect was purely a redistribution of happiness from poorer to richer individuals, we would expect the effect

[53] If individuals with lower Internet access were affected less than individuals with higher Internet access but still affected to some extent, then the coefficient α_2 would still measure the average effect of transparency but suffer from attenuation bias.

on average happiness to be null (i.e., $\alpha_2 = 0$). However, it is possible that the effects were asymmetric. However, if poorer households lost more happiness than the happiness gained by richer households, then we would expect $\alpha_2 < 0$. On the contrary, if richer households gained more happiness than the happiness lost by poorer households, then we would expect $\alpha_2 > 0$.

Column (1) of Table A.1 reports the regression results using *Happiness* as the dependent variable. The estimated α_2 (0.014) is close to zero and statistically insignificant (p-value=0.460). This evidence is consistent with a pure redistribution of happiness from poorer to richer households ($\alpha_2 = 0$). As a robustness check, column (2) introduces the falsification test to assess the possibility that the parameter α_2 is biased because of differential pre-trends across individuals with high and low Internet access, as in the specification from equation (3) in section 3. The results from column (2) are consistent with column (1): the coefficient α_2 (0.019) is still small and statistically insignificant (p-value=0.362), and the coefficient on the interaction with *I{1997–2000}* (0.015) is also small and statistically insignificant (p-value=0.582). As an additional robustness check, column (3) uses life satisfaction as the dependent variable instead of happiness. Since this outcome was measured starting in 1999, the pre-treatment period consists of just one year of data and thus the findings must be taken with a grain of salt. The estimated α_2 (0.027) is still small and statistically insignificant (p-value=0.390). In sum, the estimates for happiness and life satisfaction suggest that the change in disclosure had a null effect on the average level of well-being.

For the sake of completeness, columns (4) through (7) of Table A.1 reproduce the results from columns (1) and (2), but with the additional outcomes (*Perceived Rank* and *Income Adequacy*) as dependent variables. Column (4) suggests that the effect on the average perceived income rank is positive (0.061) and statistically significant (p-value=0.001). This finding is robust to the alternative specification, reported in column (5): the coefficient on the interaction with *I{1997–2000}* (-0.017) is close to zero, statistically insignificant (p-value=0.570), and statistically different from the interaction with *I{2001–2013}* (p-value=0.002). These results imply that individuals found out through the online tax lists that, on average, their income rank was higher than they thought. This evidence is consistent with the findings from Karadja, Mollerstrom, and Seim (2017) that, on average, households under-estimate their own income rank and thus will update their perceptions upwards when provided with accurate information.

Last, columns (6) and (7) show the results for *Income Adequacy* as the dependent variable. Column (6), which reports the most basic specification, indicates that the coefficient on the interaction with *I{2001–2013}* (0.037) is close to zero, although statistically significant (p-value=0.069). In the alternative specification, reported in column (7), this coefficient is even smaller (0.003) and becomes statistically insignificant (p-value=0.924). These results suggest that were no significant effects on the average level of *Income Adequacy*.

A.3 Probit-OLS versus OLS and Ordered Probit

In the baseline specification, we coded the subjective questions using the Probit-OLS method. In this section, we show that the results are robust under alternative econometric models.

When constructing the happiness outcome, instead of arbitrarily assigning values 1, 2, 3, and 4 to the four possible answers to the happiness question, we employ the Probit-OLS method to assign these values (van Praag and Ferrer-i-Carbonell, 2008). This method consists of assigning values to match the distribution of responses to a normal distribution. For example, if a fraction q reports the lowest category ("not at all satisfied"), the Probit-OLS method assigns the lowest category an score of $E[z|z < q]$, where z is distributed standard normal. The resulting values for the happiness scores are 1.36 ("very happy"), -0.17 ("quite happy"), -1.67 ("not particular happy") and -2.79 ("not at all happy").

Table A.2 explores the robustness of the results to different treatments of the subjective data. Columns (1) and (2) denote the baseline specifications for Happiness and Life Satisfaction, respectively (these results are identical to columns (3) and (5) from Table 3). Columns (3) and (4) from Table A.2 correspond to the specifications from columns (1) and (2), but with the responses to the happiness and life satisfaction questions coded from 1 to 4 and 1 to 5, respectively. The results from columns (3) and (4) are qualitatively consistent with the results from columns (1) and (2). The magnitudes of the coefficients from columns (3) and (4) are not directly comparable to the coefficients from columns (1) and (2), because of the differences in scales of the dependent variables. With that caveat in mind, the findings are quantitatively robust across the two specifications: column (3) suggests that the change in disclosure increased the happiness-income gradient by 27% ($= \frac{0.049}{0.179}$), which is close to (and statistically indistinguishable from) the 29% increase implied by the coefficients from column (1); and column (4) indicates that the change in disclosure increased the life satisfaction-income gradient by 20% ($= \frac{0.089}{0.452}$), which is close to (and statistically indistinguishable from) the 21% increase implied by the coefficients from column (2).

Columns (5) and (6) from Table A.2 estimate the same specifications from columns (1) and (2), except that they use an Ordered Probit model instead of the OLS model from the baseline specification. Columns (5) and (6) report the raw coefficients from the Ordered Probit model, which cannot be compared in magnitude directly to the OLS coefficients from columns (1) and (2). In terms of signs and statistical significance, the results from columns (5) and (6) are consistent with the results from columns (1) and (2). The results are also quantitatively robust: column (5) suggests that the change in disclosure increased the happiness-income gradient by 30% ($= \frac{0.113}{0.380}$), which is very close to (and statistically indistinguishable from) the 29% increase implied by the coefficients from column (1). Similarly, column (6) indicates that the change in disclosure increased the life satisfaction-income gradient by 25% ($= \frac{0.166}{0.667}$), which is close to (and statistically indistinguishable from) the 21% increase implied by the coefficients from column (2).

A.4 Binned Scatterplots

In the baseline specification, we assume a linear relationship between subjective well-being and *Income Rank*. In this section, we present results under a more flexible specification using binned scatterplots.

In addition to exploring the role of outliers and non-linearities, the binned scatterplots can also shed light on the distribution of the effects of transparency. These effects could be unevenly distributed along the income distribution due to an uneven exposure to the online tax lists. Indeed, the evidence from the triple-differences specification already indicates that the effects were concentrated in individuals with higher Internet access. Since the individuals with higher Internet access tend to be richer than the rest of the population, then we would expect effects that are stronger for richer individuals.[54]

We follow the baseline regression specification but, instead of letting *Income Rank* enter linearly in the right hand side of the regression equation, we include this variable as a set of dummies for nine equal-sized income groups: one set of dummies for the post-2001 period, and another one for the 2001–2013 period. The middle groups are set as omitted categories, and thus their coefficients are normalized to zero.

Figure A.4 presents the results for the binned scatterplot analysis. Figure A.4.a shows the results for *Happiness*, and Figures A.4.b through A.4.d show the results for the rest of the outcomes. Most important, these figures suggest that the linear specification used in the baseline specification for *Income Rank* provides a fair approximation. These figures also confirm that the results are not driven by outliers or non-linearities. Last, the results are consistent with the expectation that, due to differences in Internet access, the effects were stronger in the upper part of the income distribution.

A.5 Local vs. National Income Rank

In the baseline specification, *Income Rank* corresponds to the position of the respondent in the national distribution of household income. To the extent that this does not fully capture the income comparisons that people care about, this can be a source of measurement error and thus

[54]Additionally, richer individuals may be more salient in the online tax lists – for example, it was common for the online search tools to provide rankings with the richest individuals in each city. Also, the effects may be stronger for some income groups due to the nature of income comparisons. For example, if individuals are last-place-averse (Kuziemko et al., 2014), then the effects of transparency may be particularly strong at the bottom of the income distribution. Similarly, if well-being is a concave function of relative pay (Fehr and Schmidt, 1999), the effects of transparency may be weaker among richer individuals (Card et al., 2012).

introduce an attenuation bias. In this section we present results under an alternative definition, based on the local instead of the national income rank.

Using the local alternative for *Income Rank* may be a more appropriate specification for a number of reasons. For example, due to the Balassa-Samuelson effect, the local income rank may predict purchasing power better than the national income rank. And regarding income comparisons, individuals may care the most about the comparisons to their social contacts, which are probably drawn disproportionately from the same area of residence.

The most disaggregated geographic identifiers in the Norwegian survey data correspond to the county identifiers. There are 19 counties in Norway, with populations in 2001 ranging from 73,417 in Finnmark to 599,230 in Oslo, with a median of 233,705 in Vestfold (Source: Statistics Norway). We constructed a local *Income Rank*, based on the within-county rank instead of the national rank. The income distributions across these 19 regions are fairly similar, and as a result the national and county ranks are highly correlated to each other (correlation coefficient of 0.9681).

The results are presented in Table A.3. Columns (1) and (2) denote the baseline specifications for happiness and life satisfaction, respectively, which are identical to columns (3) and (5) from Table 3. Columns (5) and (6) from Table A.3 reproduce the specifications from columns (1) and (2) but using the local version of *Income Rank* instead. The results from columns (5) and (6) are similar to (and statistically indistinguishable from) the results from columns (1) and (2) – if anything, and consistent with the argument of attenuation bias, the estimated effects of transparency are slightly larger under the local definition of *Income Rank*.

A.6 Alternative Definitions of Income Rank

This section provides some robustness checks related to the construction of *Income Rank*.

One potential source for concern is that the income question added a bin in 1999, which could contaminate the comparison of the happiness-income gradient around 2001. Since only 1.55% of respondents fell in the ninth bin in 1999, this is probably a minor concern. To address any remaining concerns, Table A.4 presents a sharp robustness check. The results from columns (1) through (3) correspond to the baseline definition of *Income Rank* (identical to columns (1), (3) and (4) of Table 3). The results from columns (4) through (6) are based on the same specifications but with an alternative version of *Income Rank*, which ignores the distinction between the eight and ninth bin. That is, we treat the data as if the ninth bin was never introduced: we start by pooling the responses to the ninth and eight bins in the raw data, and then we replicate the data construction and analysis starting from this revised dataset. The results are robust to this check: the coefficients from columns (1) through (3) are almost identical to the corresponding coefficients from columns (4) through (6). For example, column (1) indicates that transparency increased the happiness-income gradient by 0.090 (p-value=0.005), while the corresponding coefficient in column (4) is 0.090 (p-value=0.005).

Since the income question is elicited in bins, our baseline specification uses the standard method from the happiness literature to impute the values of *Income Rank* within each bin (Stevenson and Wolfers, 2008; Kahneman and Deaton, 2010). Table A.4 also assesses the sensitivity of the results to this imputation method. Again, columns (1) through (3) correspond to the results with the baseline definition of *Income Rank*, which uses the imputation. Columns (7) through (9) show the results under the alternative definition of *Income Rank* using the raw data (i.e., without the within-bin imputation). Most important, the results are qualitatively and quantitatively robust across the imputed and non-imputed versions of *Income Rank*. For example, column (1) indicates that the happiness-income gradient increased by 29% after 2001, while column (7) indicates that it increased by 25% – moreover, we cannot reject the null hypothesis that these two estimates are equal. Additionally, the comparison of the coefficient on *Income Rank* between columns (1)–(3) and (7)–(9) suggests that the imputation method is indeed helping to ameliorate the measurement error: the happiness-income gradients are almost 10% higher when using the imputation.

A.7 Alternative Definitions of Higher Internet

This section shows the robustness of the results under alternative definitions of *I{Higher Internet}*.

The results are presented in Table A.5. Column (1) corresponds to the baseline definition of *I{Higher Internet}* (identical to column (4) of Table 3). Columns (2) through (5) are based on the same specification but with alternative versions of *I{Higher Internet}*. In column (2), *I{Higher Internet}* is identical to the baseline definition except that, instead of using responses to *Internet Access* for 2001, it is based on responses for 1999. In column (3), *I{Higher Internet}* is identical to the baseline definition except that it is based on responses to *Internet Access* for the entire period for which the *Internet Access* is available (1999–2013). In column (4), *I{Higher Internet}* is identical to the baseline definition except that it splits individuals by the median value of predicted Internet access over the entire sample, rather than splitting them by the median value for the respective year. In column (5), *I{Higher Internet}* is identical to the baseline definition except that we predict Internet access using a Probit regression instead of an OLS regression. The coefficients reported in Table A.5 indicate that the results are qualitatively and quantitatively robust between the baseline specification (column (1)) and the alternative specifications (columns (2)–(5)).

A.8 Evolution of Inequality and Sample Composition

Since the survey is a repeated cross-section, we need to check that the composition of the survey respondents has not changed abruptly around 2001. To address this concern, Table A.6 presents the yearly averages of some individual characteristics. Consistent with the gradual changes in the composition of the universe of Norwegians, over the 28 years there were gradual changes in the composition of of the survey respondents: a gradual increase in age and education, and a gradual decrease in marriage rates. Most important, the results confirm that the composition of the survey respondents has not changed abruptly around 2001.

Relatedly, Table A.3 explore the robustness of the results to the use of sampling weights. Columns (1) and (2) denote the baseline specifications for happiness and life satisfaction, respectively, which are identical to columns (3) and (5) of Table 3. Columns (3) and (4) of Table A.3 reproduce the same regressions from columns (1) and (2), but using individual-specific sampling weights computed by the team in charge of collecting the survey data. As expected, using sample weights does not change the results: the coefficients in columns (3) and (4) are very similar to (and statistically indistinguishable from) the coefficients from columns (1) and (2).

Another potential source for concern is that the change in the happiness-income gradient is mechanically driven by an increase in income inequality. In other words, richer individuals may become happier because they are able to afford more stuff, while poorer individuals become less happy because they can afford less stuff. This possibility seems highly unlikely, because it would require a large, sudden and persistent increase in inequality that would be unprecedented in a developed country. To address this concern more directly, Table A.7 presents data on the evolution of income inequality in Norway during the sample period.

Table A.7 shows the evolution of income inequality according to the data from the Norwegian Monitor Survey. The coefficient of variation in incomes did not change abruptly in 2001 – furthermore, it was stable during the entire 1985–2001 period. For example, we can mimic the comparison from the event-study analysis: the coefficient of variation went from an average of 0.521 in 1997/99 to an average of 0.520 in 2001/03, amounting to a mere reduction of 0.21%. This change in inequality is tiny compared to the estimated 29% increase in the happiness-income gradient around 2001.

Table A.7 also presents results for two measures of inequality constructed with administrative data.[55] The *Gross Gini Index* is an index of inequality based on the gross household income, while the *Net Gini Index* is an index of inequality based on the net household income. Just like the survey data, the administrative data indicates that there was no abrupt change in income inequality around 2001, and that the levels of inequality were stable over the entire 1985–2013 period. Moreover, we can mimic the comparison from the event-study analysis with

[55] These measure of inequality were obtained from the Chartbook of Economic Inequality and are publicly available in the following URL: https://www.chartbookofeconomicinequality.com/.

these alternative measures of inequality. The *Gross Gini Index* increased from an average of 44.24 in 1997/99 to an average of 44.79 in 2001/03. This amounts to a mere 1.3% increase in income inequality, which is insignificant relative to the 29% increase in the happiness-income gradient measured around 2001. Similarly, the *Net Gini Index* increased from an average of 24.55 in 1997/99 to an average of 25.15 in 2001/03, amounting to a mere 2.4% increase in income inequality. Again, these changes in income inequality are minuscule compared to the estimated 29% increase in the happiness-income gradient around 2001.

Figure A.1: More Screenshots of Websites and Apps Designed to Search the Tax Records

a. Official Search Tool for the Tax Records of 2014

b. TV2s Skattelisten Iphone App

c. Interactive Tool to Learn about the Income Distribution

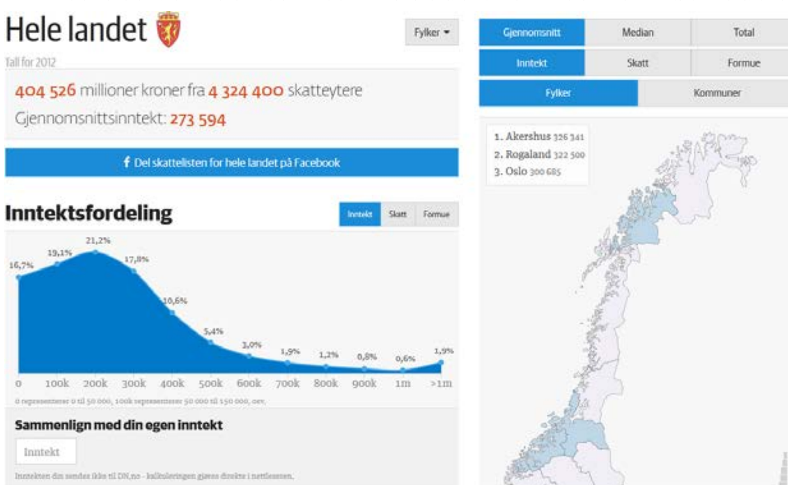

Source: (a) and (b): Origo (2010). (c) web.archive.org.

Figure A.2: Alternative Event-Study Graphs for Subjective Well-Being

a. Norway (yearly) **b.** Germany (odd years)

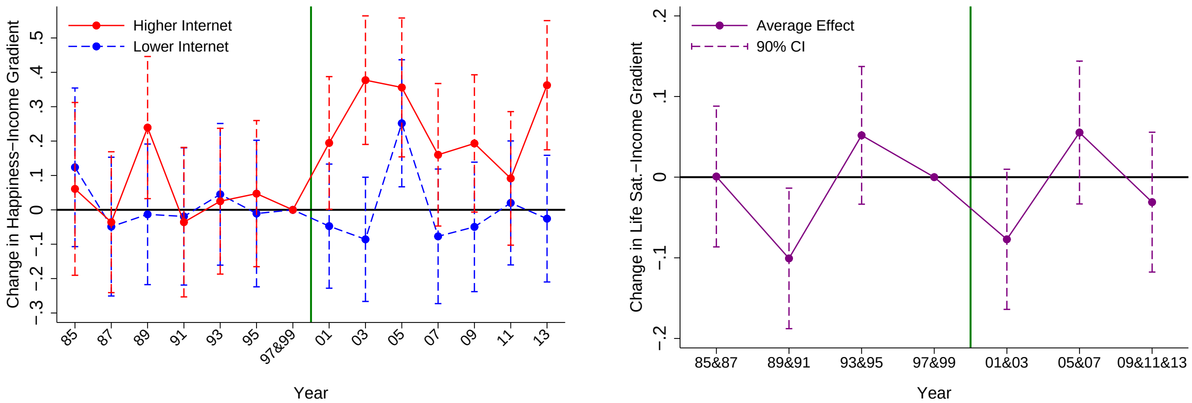

Notes: Variations of the event study graphs presented in Figure 4. Panel (a) reproduces panel (b) from Figure 4, except that it breaks down the coefficients at the year level. Panel (b) reproduces panel (c) from Figure 4, only that it restricts the data to odd years in 1985–2013. See notes to Figure 4 for more details about the specification and the data.

Figure A.3: Event-Study Graphs for Perceived Income Rank and Income Adequacy

a. *Perceived Rank* **b.** *Income Adequacy*

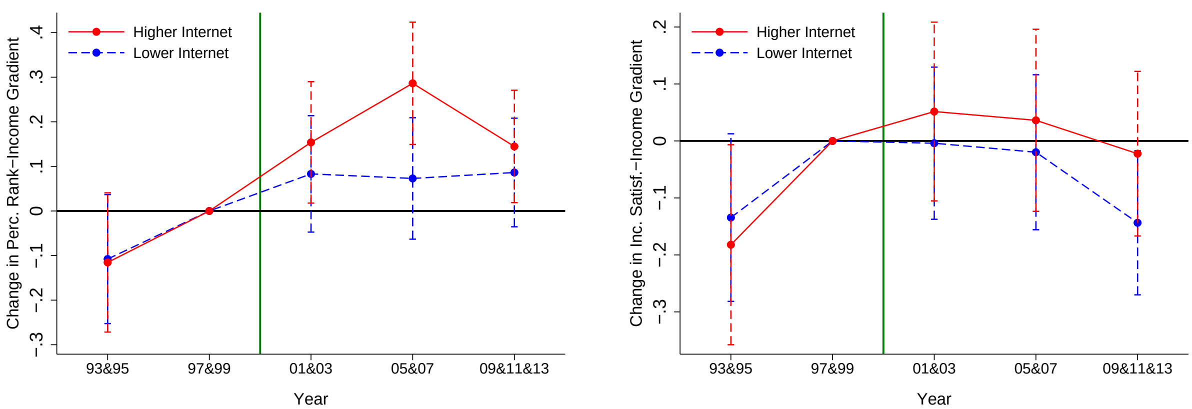

Notes: Panel (a) reproduces panel (b) from Figure 4, only that it uses *Perceived Rank* as the dependent variable instead of *Happiness*. Panel (b) reproduces panel (b) from Figure 4, only that it uses *Income Adequacy* as dependent variable instead of *Happiness*. All the dependent variables have been normalized to have mean 0 and standard deviation of 1. See notes to Figure 4 for more details about the specification and the data.

Figure A.4: Binned Scatterplot Showing the Change in Gradient Between Happiness and Income Rank

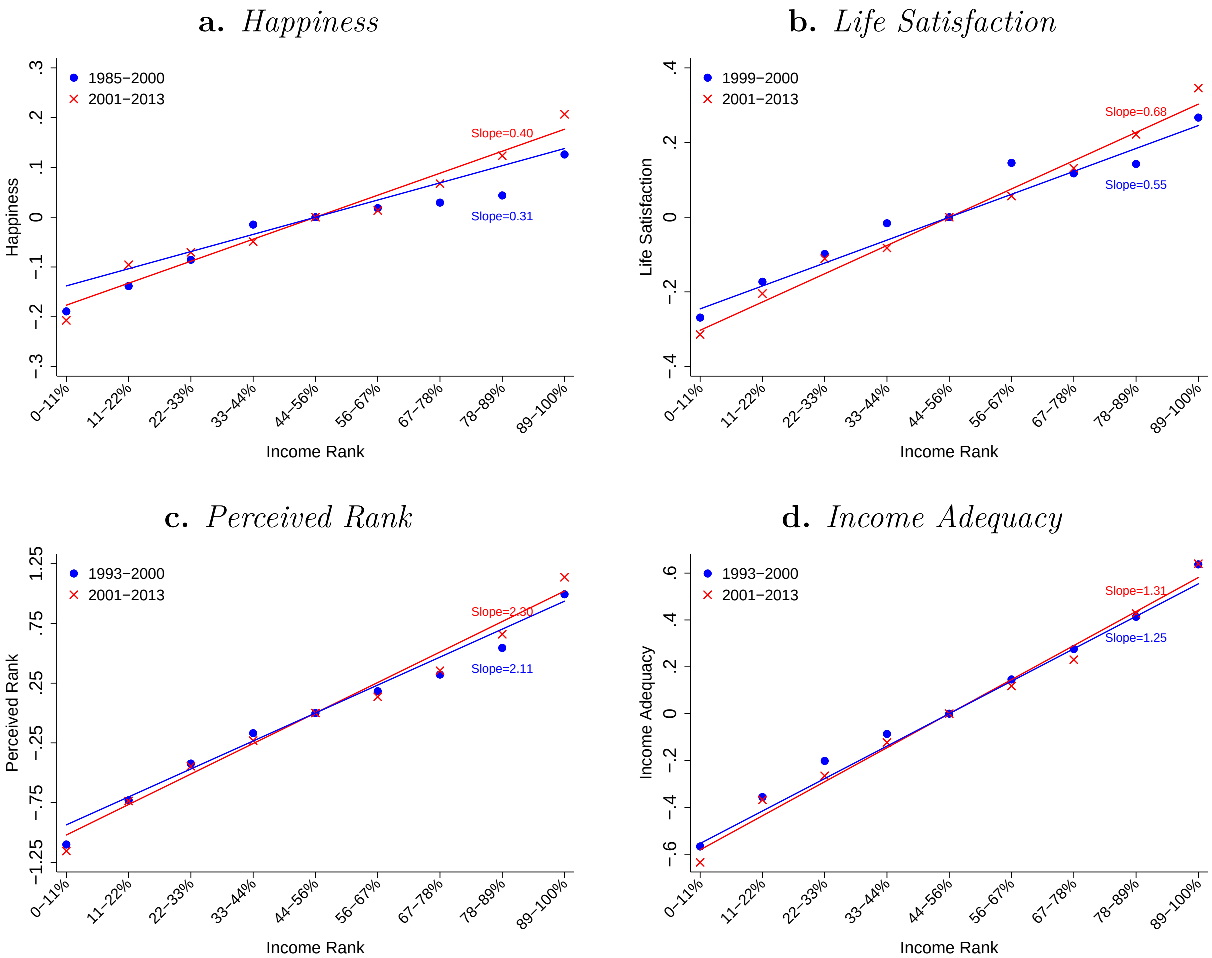

Notes: Panel (a) corresponds to the same regression from column (1) of Table 3, only that *Income Rank* is introduced as two sets of nine equally-sized dummies (one set for the post-2001 period and another one for the pre-2001 period), with the coefficients on the middle categories normalized to zero. Panel (b), (c) and (d) are identical to panel (a), only that instead of *Happiness* they use the dependent variables *Life Satisfaction*, *Perceived Rank* and *Income Adequacy*, respectively. All these dependent variables have been normalized to have mean 0 and standard deviation of 1, with higher values denoting higher happiness/satisfaction/rank/adequacy. See notes to Table 3 for more details about the regression specification and the data.

Table A.1: Effects on the Average Level of Well-Being

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Happiness	Happiness	Life Satisf.	Perc. Rank	Perc. Rank	Income Adequacy	Income Adequacy
I{Higher Internet}	-0.010	-0.016	-0.037	-0.078***	-0.067**	-0.048**	-0.013
	(0.017)	(0.020)	(0.033)	(0.019)	(0.027)	(0.021)	(0.029)
I{Higher Internet} * I{2001-2013}$^{(i)}$	0.014	0.019	0.027	0.061***	0.050*	0.037*	0.003
	(0.018)	(0.021)	(0.032)	(0.018)	(0.026)	(0.020)	(0.029)
I{Higher Internet} * I{1997-2000}$^{(ii)}$		0.015			-0.017		-0.055*
		(0.027)			(0.030)		(0.034)
P-value (i)=(ii)		0.861			0.002		0.016
Period	85-13	85-13	99-13	93-13	93-13	93-13	93-13
Observations	48,570	48,570	29,655	38,938	38,938	38,950	38,950

Notes: *** p<0.01, ** p<0.05, * p<0.1. Heteroskedasticity-robust standard errors in parenthesis. Each column corresponds to a separate OLS regression. *Happiness, Life Satisfaction*, *Perceived Rank* and *Income Adequacy* were normalized to have mean 0 and standard deviation of 1, with higher values denoting higher happiness/satisfaction/rank/adequacy. *I{Higher Internet}* is a dummy variable that takes the value 1 if the respondent's predicted *Internet Access* is above the median value for a given year. *I{2001–2013}* takes the value 1 for 2001–2013. *I{1997–2000}* takes the value 1 for 1997–2000. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size, three dummies for number of working household members and *Income Rank* (i.e., the respondent's position in the national distribution of household income for that year). Data from the Norwegian Monitor Survey, which has been collected every other year in 1985–2013. See Table 1 for a summary of data definitions and Table 2 for descriptive statistics.

Table A.2: Robustness to Alternative Econometric Models

	(1) Happiness	(2) Life Satisf.	(3) Happiness	(4) Life Satisf.	(5) Happiness	(6) Life Satisf.
Income Rank	0.310***	0.585***	0.179***	0.452***	0.380***	0.667***
	(0.032)	(0.056)	(0.018)	(0.043)	(0.039)	(0.064)
Income Rank * I{2001-2013}$^{(i)}$	0.090**	0.122**	0.049**	0.089**	0.113**	0.166***
	(0.037)	(0.055)	(0.021)	(0.042)	(0.045)	(0.062)
Income Rank * I{1997-2000}$^{(ii)}$	0.001		0.001		-0.000	
	(0.048)		(0.028)		(0.059)	
P-value (i)=(ii)	0.043		0.053		0.034	
Model	OLS	OLS	OLS	OLS	O-Probit	O-Probit
POLS Transformation	Yes	Yes	No	No	No	No
Period	85-13	99-13	85-13	99-13	85-13	99-13
Observations	48,570	29,655	48,570	29,655	48,570	29,655

Notes: *** p<0.01, ** p<0.05, * p<0.1. Heteroskedasticity-robust standard errors in parenthesis. Column (1) through (4) report coefficients from OLS regressions, while columns (5) and (6) report raw coefficients from an Ordered Probit model. *Happiness* and *Life Satisfaction* are responses to subjective questions where higher value denotes higher happiness/satisfaction. In columns (1) and (2), responses to these questions were coded using the Probit-OLS method, and then normalized to have mean 0 and standard deviation of 1. In columns (3) and (5), responses to the happiness question are assigned values from 1 (not at all happy) to 4 (very happy). In columns (4) and (6), responses to the life satisfaction question are assigned values from 1 (very dissatisfied) to 4 (very satisfied). *Income Rank* denotes the position of the respondent's household relative to all the other respondents for that year, from 0 to 1. *I{2001–2013}* takes the value 1 for 2001–2013. *I{1997–2000}* takes the value 1 for 1997–2000. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size and three dummies for number of working household members. Data from the Norwegian Monitor Survey, which has been collected every other year in 1985–2013. See Table 1 for a summary of data definitions and Table 2 for descriptive statistics.

Table A.3: Robustness to the Local Definition of Income Rank and the Sample Weights

	(1)	(2)	(3)	(4)	(5)	(6)
		Life		Life		Life
	Happiness	Satisf.	Happiness	Satisf.	Happiness	Satisf.
Income Rank	0.310***	0.585***	0.306***	0.563***	0.268***	0.554***
	(0.032)	(0.056)	(0.036)	(0.063)	(0.032)	(0.056)
Income Rank * I{2001-2013}$^{(i)}$	0.090**	0.122**	0.087**	0.112*	0.098***	0.115**
	(0.037)	(0.055)	(0.043)	(0.061)	(0.037)	(0.055)
Income Rank * I{1997-2000}$^{(ii)}$	0.001		-0.022		0.006	
	(0.048)		(0.056)		(0.048)	
P-value (i)=(ii)	0.043		0.034		0.035	
Weights	No	No	Yes	Yes	No	No
Income Rank	Nation	Nation	Nation	Nation	County	County
Period	85-13	99-13	85-13	99-13	85-13	99-13
Observations	48,570	29,655	48,570	29,655	48,570	29,655

Notes: *** p<0.01, ** p<0.05, * p<0.1. Heteroskedasticity-robust standard errors in parenthesis. Each column corresponds to a separate OLS regression. *Happiness* and *Life Satisfaction* are responses to subjective questions where higher value denotes higher happiness/satisfaction, normalized to have mean 0 and standard deviation of 1. *Income Rank* denotes the position of the respondent's household, from 0 to 1, relative to all the other respondents for that year in the nation (columns (1) through (4)) or county (columns (5) and (6)). *I{2001–2013}* takes the value 1 for 2001—2013. *I{1997–2000}* takes the value 1 for 1997—2000. In columns (3) and (4), the regressions use population weights computed by the group in charge of conducting the survey. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size and three dummies for number of working household members. Data from the Norwegian Monitor Survey, which has been collected every other year in 1985–2013. See Table 1 for a summary of data definitions and Table 2 for descriptive statistics.

Table A.4: Robustness to Alternative Definitions of Income Rank

	(1) Happiness	(2) Happiness	(3) Happiness	(4) Happiness	(5) Happiness	(6) Happiness	(7) Happiness	(8) Happiness	(9) Happiness
Income Rank	0.311***	0.310***	0.331***	0.313***	0.313***	0.333***	0.291***	0.283***	0.305***
	(0.028)	(0.032)	(0.040)	(0.028)	(0.032)	(0.040)	(0.026)	(0.030)	(0.037)
Income Rank * I{2001-2013}$^{(i)}$	0.090***	0.090**	-0.004	0.090***	0.091**	-0.006	0.072**	0.081**	-0.003
	(0.032)	(0.037)	(0.051)	(0.032)	(0.037)	(0.052)	(0.031)	(0.036)	(0.048)
Income Rank * I{2001-2013} * I{Higher Internet}			0.217***			0.223***			0.168**
			(0.073)			(0.073)			(0.069)
Income Rank * I{1997-2000}$^{(ii)}$		0.001			0.001			0.024	
		(0.048)			(0.048)			(0.048)	
P-value (i)=(ii)		0.043			0.042			0.196	
Income Rank Definition	I	I	I	II	II	II	III	III	III
Observations	48,570	48,570	48,570	48,570	48,570	48,570	48,570	48,570	48,570

Notes: *** p<0.01, ** p<0.05, * p<0.1. Heteroskedasticity-robust standard errors in parenthesis. Each column corresponds to a separate OLS regression. *Happiness* is normalized to have mean 0 and standard deviation of 1, with higher values denoting higher happiness. *Income Rank* denotes the position of the respondent's household in the national income distribution, from 0 to 1. In columns (1) through (3), *Income Rank* is defined as in the baseline specification described in Table 1. In columns (4) through (6), the ninth bin (introduced in 1999) is merged with the eight bin before constructing *Income Rank*. In columns (7) through (9), *Income Rank* is constructed without the imputation of the within-bin ranks. *I{2001–2013}* takes the value 1 for 2001—2013. *I{1997–2000}* takes the value 1 for 1997—2000. *I{Higher Internet}* is a dummy variable that takes the value 1 if the respondent's predicted *Internet Access* is above the median value for a given year. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size and three dummies for number of working household members. Data from the Norwegian Monitor Survey, which has been collected every other year in 1985–2013. See Table 1 for a summary of data definitions and Table 2 for descriptive statistics.

Table A.5: Robustness to Alternative Definitions of Higher Internet

	(1) Happiness	(2) Happiness	(3) Happiness	(4) Happiness	(5) Happiness
Income Rank	0.331***	0.336***	0.382***	0.351***	0.334***
	(0.040)	(0.040)	(0.041)	(0.040)	(0.040)
Income Rank * I{2001-2013}$^{(i)}$	-0.004	0.016	-0.023	-0.010	0.007
	(0.051)	(0.050)	(0.053)	(0.053)	(0.052)
Income Rank * I{2001-2013} * I{Higher Internet}	0.217***	0.152**	0.209***	0.211***	0.187**
	(0.073)	(0.071)	(0.075)	(0.074)	(0.074)
I{Higher Internet} Definition	I	II	III	IV	V
Observations	48,570	48,570	48,570	48,570	48,570

Notes: *** $p<0.01$, ** $p<0.05$, * $p<0.1$. Heteroskedasticity-robust standard errors in parenthesis. Each column corresponds to a separate OLS regression. *Happiness* is normalized to have mean 0 and standard deviation of 1, with higher value denoting higher happiness. *Income Rank* denotes the position of the respondent's household in the national income distribution, from 0 to 1. *I{2001–2013}* takes the value 1 for 2001—2013. *I{Higher Internet}* is a dummy variable that takes the value 1 if the respondent's predicted *Internet Access* is above the median value for a given year. In column (1), *I{Higher Internet}* is defined as in the baseline specification described in Table 1. In columns (2) through (5), *I{Higher Internet}* is based on alternative definitions. In column (2), *I{Higher Internet}* is identical to the baseline definition except that it is based on responses to *Internet Access* for 1999 instead of 2001. In column (3), *I{Higher Internet}* is identical to the baseline definition except that it is based on responses to *Internet Access* for 1999–2013 instead of 2001. In column (4), *I{Higher Internet}* is identical to the baseline definition except that instead of splitting the sample within each year, we split the sample using the median value over the entire 1985–2013 period. In column (5), *I{Higher Internet}* is identical to the baseline definition except that it is constructed using a Probit model instead of an OLS model. All regressions control for year dummies, age, age squared, a gender dummy, three education dummies, four dummies for marital status, four dummies for household size and three dummies for number of working household members. Data from the Norwegian Monitor Survey, which has been collected every other year in 1985–2013. See Table 1 for a summary of data definitions and Table 2 for descriptive statistics.

Table A.6: Descriptive Statistics by Year

	Female	Age	Married	College	Oslo	Density
1985	0.499	40.861	0.721	0.231	0.118	185.719
	(0.011)	(0.361)	(0.010)	(0.010)	(0.007)	(9.388)
1987	0.505	40.733	0.731	0.249	0.119	188.673
	(0.010)	(0.325)	(0.009)	(0.009)	(0.006)	(8.338)
1989	0.508	39.759	0.684	0.280	0.108	176.713
	(0.010)	(0.318)	(0.009)	(0.009)	(0.006)	(7.789)
1991	0.501	39.352	0.622	0.270	0.114	182.496
	(0.010)	(0.333)	(0.010)	(0.009)	(0.006)	(8.063)
1993	0.494	39.989	0.527	0.288	0.122	192.982
	(0.010)	(0.304)	(0.010)	(0.009)	(0.006)	(8.182)
1995	0.490	40.906	0.483	0.311	0.119	188.013
	(0.010)	(0.334)	(0.010)	(0.009)	(0.006)	(8.371)
1997	0.517	41.212	0.488	0.360	0.113	181.017
	(0.008)	(0.245)	(0.008)	(0.007)	(0.005)	(6.358)
1999	0.537	42.492	0.486	0.394	0.117	186.894
	(0.008)	(0.253)	(0.008)	(0.008)	(0.005)	(6.597)
2001	0.554	44.571	0.496	0.412	0.114	183.217
	(0.008)	(0.252)	(0.008)	(0.008)	(0.005)	(6.542)
2003	0.542	45.205	0.510	0.421	0.114	185.332
	(0.008)	(0.257)	(0.008)	(0.008)	(0.005)	(6.626)
2005	0.559	47.563	0.539	0.461	0.112	179.919
	(0.008)	(0.261)	(0.008)	(0.008)	(0.005)	(6.674)
2007	0.531	51.840	0.573	0.489	0.117	188.027
	(0.008)	(0.253)	(0.008)	(0.008)	(0.005)	(6.727)
2009	0.561	48.434	0.546	0.521	0.116	187.730
	(0.008)	(0.293)	(0.008)	(0.008)	(0.005)	(7.004)
2011	0.531	50.763	0.579	0.598	0.124	198.645
	(0.008)	(0.249)	(0.008)	(0.008)	(0.005)	(6.850)
2013	0.535	47.821	0.503	0.577	0.137	215.520
	(0.008)	(0.304)	(0.008)	(0.008)	(0.006)	(7.310)

Notes: Data from the Norwegian Monitor Survey (48,570). *Female* takes the value 1 if the respondent is female. *Age* is the age in years. *Married* takes the value 1 if married. *College* takes the value 1 for College graduates. *Oslo* takes the value 1 for Oslo residents. *Density* corresponds to the population per square kilometer in the county of residence (Statistics Norway, 2011). Heteroskedasticity-robust standard errors in parenthesis.

Table A.7: Income Inequality by Year

	Administrative Data		Norwegian Monitor Survey						
	Gini Index		*Absolute Income*			*Income Rank*		*I{Higher Internet}*	
Year	*Gross*	*Net*	Mean	SD	CV	Mean	SD	Mean	SD
1985	40.37	21.00	5.10	2.96	0.581	0.50	0.29	0.50	0.50
1986	40.47	21.00							
1987	40.84	21.00	5.76	3.24	0.563	0.50	0.29	0.50	0.50
1988	40.72	21.10							
1989	40.07	22.80	5.69	3.09	0.543	0.50	0.29	0.50	0.50
1990	40.58	21.70							
1991	41.07	21.90	5.53	3.32	0.600	0.50	0.29	0.50	0.50
1992	42.59	22.30							
1993	43.98	22.90	4.44	2.40	0.540	0.50	0.29	0.50	0.50
1994	44.35	24.10							
1995	44.18	23.60	4.22	2.47	0.586	0.50	0.29	0.50	0.50
1996	44.36	24.50							
1997	44.72	24.90	4.79	2.48	0.518	0.50	0.29	0.50	0.50
1998	43.59	23.80							
1999	43.76	24.20	5.18	2.71	0.524	0.50	0.29	0.50	0.50
2000	45.57	26.20							
2001	43.51	22.90	5.32	2.79	0.524	0.50	0.29	0.50	0.50
2002	45.30	26.40							
2003	46.08	27.40	5.86	3.02	0.516	0.50	0.29	0.50	0.50
2004	46.90	28.30							
2005	50.62	32.70	6.06	3.11	0.513	0.50	0.29	0.50	0.50
2006	44.67	24.30							
2007	45.68	25.20	6.40	3.27	0.511	0.50	0.29	0.50	0.50
2008	44.90	24.80							
2009	44.68	24.10	6.70	3.44	0.513	0.50	0.29	0.50	0.50
2010	45.13	24.50							
2011	45.14	24.70	7.45	3.80	0.510	0.50	0.29	0.50	0.50
2012	45.22	24.90							
2013	45.44	25.00	7.72	4.52	0.585	0.50	0.29	0.50	0.50

Notes: The *Gini Index* are inequality indeces. The *Gross Gini Index* is based on gross household income, while the *Net Gini Index* is based on the net household income. Both of these measures were obtained from the Chartbook of Economic Inequality and are based on administrative data. The rest of the outcomes are based on data from the Norwegian Monitor Survey (N=48,570). *Absolute Income* is the gross household income measured in hundreds of thousands of Kroner, and converted to 2013 prices with the consumer price index from Statistics Norway. *SD* stands for standard deviation, and *CV* stands for coefficient of variation (i.e., the ratio between the standard deviation and the mean). *Income Rank* is respondent's position in the distribution of *Absolute Income* in a given year, and *I{Higher Internet}* is a dummy variable that takes the value 1 if the respondent's predicted *Internet Access* is above–median – see Table 1 for more detailed data definitions.

Table A.8: Auxiliary Regression Results: Predictors of Internet Access

Dependent Variable: *Internet Access * 100*			
Female	-9.294***	Number of HH Members (omitted: 1)	
	(1.391)	2	8.449***
Age	0.479*		(3.029)
	(0.266)	3	13.561***
Age Squared	-0.013***		(3.334)
	(0.003)	4	12.808***
Education (omitted: Primary School)			(3.656)
Middle School	11.688***	5+	16.999***
	(2.920)		(3.945)
High School	18.881***	Number of HH Workers (omitted: 0)	
	(2.732)	1	9.448***
College	30.729***		(2.609)
	(2.731)	2	18.859***
Marital Status (omitted: Married)			(2.794)
Cohabitant	-10.327***	3+	24.909***
	(2.218)		(4.055)
Single	-7.443**	Constant	33.856***
	(3.327)		(8.607)
Separated/Divorced	-4.310		
	(3.363)		
Widowed	2.287		
	(4.063)		

Notes: $N = 3,931$. *** p<0.01, ** p<0.05, * p<0.1. Heteroskedasticity-robust standard errors in parenthesis. Coefficients of an OLS regression. *Internet Access*100* takes the value 100 if the respondent has Internet access at home and 0 otherwise. Data from the Norwegian Monitor Survey for the year 2001. The average of the dependent variable is 60.75 percentage points.

B A Simple Model of the Effects of Income Transparency Through Income Comparisons

This section provides a simple model to illustrate how an increase in income transparency can increase the gradient between utility and income rank. This model is not intended to capture all the different forms that income comparisons can take. Instead, it is based on two specific channels (social-esteem and self-esteem) that have received some attention in the economics literature.

B.1 The Model

There is a continuum of individuals with a non-degenerate income distribution, where r_i^{true} denotes the true relative position in the income distribution of individual i.

Intrinsic Utility. We assume that intrinsic utility from income is a linear function of the true income rank of the individual:

$$U_i^{intrinsic} = \eta_0 \cdot r_i^{true} \tag{B.1}$$

Using this particular functional form for the intrinsic utility function is made just to simplify the notation – the intuitions would still apply under more standard functional forms, such as making intrinsic utility equal to the logarithm of absolute income.

We model income comparisons through two distinct channels: self-image utility and social-image utility. These two sources of utility depend on the social interactions with other individuals from the same population. The incomes of the individuals involved in each interaction are observable with some exogenous probability $\nu \in [0, 1]$. This parameter ν is a reduced-form representation of the degree of income visibility.[56]

Social-Image Utility. Each individual is paired with a random individual from the same population. A third party, the allocator, allocates some social-image utility worth $\eta_1 > 0$. The allocator wants to give the social-image utility to the individual in the pair with the higher income. With probability ν, the allocator can observe the incomes of the two individuals in the pair, in which case the allocator gives the social-esteem to the individual with the higher income. With probability $1-\nu$ the allocator cannot observe incomes, in which case the allocator simply randomizes who gets the social-esteem. As a result, the ex-ante utility from social-image is the following:

[56] We use a unique ν to simplify the notation. In reality, there should be at least two ν's: one relevant for the formation of self-image and another for the formation of social-image. For example, it is probably easier to observe information to infer one's position in the income distribution, for which it suffices to have access to aggregate income statistics, than to observe the income of a particular individual.

$$U_i^{social} = \nu \cdot r_i^{true} \cdot \eta_1 + (1-\nu) \cdot \frac{1}{2} \cdot \eta_1 \tag{B.2}$$

By taking the derivative of (B.2) with respect to r_i^{true}, we obtain the gradient between social-image utility and income. By taking an additional derivative with respect to ν, we show that this gradient increases with income visibility:

$$\frac{\partial^2 U_i^{social}}{\partial r_i^{true} \partial \nu} = \eta_1 > 0 \tag{B.3}$$

The intuition behind this result is simple. An increase in visibility would make an individual with a below-median income worse off, because with a higher probability her peers would observe her income and learn that she is poorer than they would have thought otherwise. On the other hand, a visibility increase would make an individual with above-median income better off, because with a higher probability the peers would observe her income and learn that she is richer than they would have thought otherwise.

Self-Image Utility. Self-image utility is similar to social-image utility, only that the individual is comparing herself to others rather than being compared by the allocator. The individual is paired with another individual randomly chosen from the population. The individual must decide whether she deserves some self-image utility worth $\eta_2 > 0$. Whether the individual feels deserving of or not depends on whether she thinks she is richer than the individual she is paired with. With probability ν, she can observe the actual income of the peer, in which case she gets the self-image utility if and only if her income is higher than the income of the peer. From an ex-ante perspective, this happens with probability r_i^{true} (i.e., equal to the probability of being paired with someone poorer). With probability $1-\nu$, the income of the paired peer is not observable, in which case she will get a fraction of the self-image utility, equal to the perceived probability of being richer than the other individual. Let r_i^{self} be this prior perceived probability. We let the prior beliefs be heterogeneous and (possibly) correlated to the actual income ranks: $r_i^{self,prior} = \theta_0 + \theta_1 \cdot r_i^{true}$.

Let $r_i^{self,post}$ be the posterior belief about the own income rank. If the individual learns rationally, the expectation of the posterior beliefs should be as follows:

$$r_i^{self,post} = \nu \cdot r_i^{true} \cdot \eta_2 + (1-\nu) \cdot \theta \cdot r_i^{true} \tag{B.4}$$

Following the evidence on the middle class bias (Cruces, Perez-Truglia and Tetaz, 2013), we assume $\theta_1 < 1$.[57] As a result, the ex-ante expected utility from self-image is:

[57] To figure out whether θ_1 is equal, above or below 1, we can test an intermediate prediction of the model: $\frac{\partial^2 r_i^{self,post}}{\partial r_i^{true} \partial \nu} = (1-\theta_1)$. Thus, if $\theta < 1$, we would predict that higher income visibility increases the gradient between self-perceived income rank and actual income rank.

$$U_i^{self} = \nu \cdot r_i^{true} \cdot \eta_2 + (1-\nu) \cdot \Big(\theta_0 + \theta_1 \cdot r_i^{true}\Big) \cdot \eta_2 \tag{B.5}$$

By taking the derivative of (B.5) with respect to r_i^{true}, we obtain the gradient between self-image utility and income. By taking an additional derivative with respect to ν, we can show that this gradient is increasing in income visibility:

$$\frac{\partial^2 U_i^{self}}{\partial r_i^{true} \partial \nu} = \eta_2 \cdot (1 - \theta_1) > 0 \tag{B.6}$$

The intuition for this result is also straightforward. When incomes are more easily observable, poor individuals learn that they are actually poorer than they thought, thus losing self-image utility; and rich individuals learn that they are actually richer than they thought, thus gaining self-image utility.

Finally, we can also explore the predictions of this model for the effect of income transparency on average well-being. Regarding social-image, it is straightforward to check that the average effect is zero. Intuitively, increasing visibility transfers social-image from poor to rich individuals, but no social-image utility gets created or destroyed in the process. Regarding self-image, it is straightforward to check that the average effect of higher visibility depends on whether θ_0 is above or below $\frac{1}{2}$. Intuitively, if $\theta_0 > \frac{1}{2}$, it means that on average individuals were over-estimating their own position in the income distribution. Since the higher transparency corrects this systematic bias, there is a net loss in utility from self-image. Similarly, if if $\theta_0 < \frac{1}{2}$, higher visibility would lead to an increase in average happiness; and if the average bias in self-perceived income rank was zero ($\theta_0 = \frac{1}{2}$), then higher income transparency would have no effect on the average utility from self-image.

B.2 Back-of-the-Envelope Calculations

To simplify the notation, we focus on the most obvious case of $\theta_1 = 0$: i.e., if all incomes were completely unobservable ($\nu = 0$), then self-perceptions about income rank would be orthogonal to actual income ranks. To obtain the overall utility, we must add up the three sources of utility: $U_i = U_i^{intrinsic} + U_i^{social} + U_i^{self}$. We add ((B.1), (B.2) and (B.5)) up and then re-arrange as follows:

$$U_i = (\beta_1 + \beta_2 \cdot \nu) \cdot r_i^{true} + \epsilon_i, \tag{B.7}$$

where $\beta_1 = \eta_0$, $\beta_2 = \eta_1 + \eta_2$ and $\epsilon_i = (1-\nu) \cdot \Big(\frac{1}{2} \cdot \eta_1 + \theta_0 \cdot \eta_2\Big)$. Note that β_1 measures the intrinsic utility from income, while $\beta_2 \cdot \nu$ measures the utility from income through income comparisons (in this model, self-image and social-image). As a result, $\frac{\beta_2 \cdot \nu}{\beta_1 + \beta_2 \cdot \nu}$ measures the value of income comparisons relative to intrinsic consumption. Intuitively, in a world where all incomes are unobservable ($\nu = 0$), an increase in one's income cannot increase the utility from self-image or social-image. On the other extreme, increasing one's income has the highest possible effect on self-image and social-image when incomes are perfectly visible ($\nu = 1$).

Let $\bar{\nu}_{t<2001}$ and $\bar{\nu}_{t\geq 2001}$ denote the visibility before 2001 and after 2001, respectively. Note that $\bar{\nu}_{t<2001}$ must be greater than zero, because even when the tax records were private, individuals could use other means to learn about the incomes of social contacts and about the income distribution. For instance, individuals could learn about the income distribution from school, from the media, or by talking with others about wages and consumption. Also, individuals reveal their own income to their social contacts, or signal it through conspicuous consumption. Similarly, $\bar{\nu}_{t\geq 2001}$ must be lower than 1, because even when the tax records were easily accessible online, there was still a small cost in attention, memory, and time to search those records. Thus, individuals did not search for the incomes of everyone with whom they interacted. Also, the online tax lists probably could not make incomes visible among total strangers – to find out someone's income in the tax records, you need to know the name of that person.

We can measure the importance of income comparisons as $s_t = \frac{\beta_2 \cdot \nu_t}{\beta_1 + \beta_2 \cdot \nu_t}$, which is the share of the happiness-income gradient that can be explained by income comparisons. Combining the regression model (1) with (B.7) and re-arranging:

$$s_{t<2001} = \frac{\frac{\alpha_2}{\alpha_1}}{\frac{\bar{\nu}_{t\geq 2001} - \bar{\nu}_{t<2001}}{\bar{\nu}_{t<2001}}}, \quad s_{t\geq 2001} = \frac{1 + \frac{1}{\frac{\bar{\nu}_{t\geq 2001} - \bar{\nu}_{t<2001}}{\bar{\nu}_{t<2001}}}}{1 + \frac{1}{\frac{\alpha_2}{\alpha_1}}} \tag{B.8}$$

The value of income comparisons depends on two parameters: $\frac{\alpha_2}{\alpha_1}$ and $\frac{\bar{\nu}_{t\geq 2001} - \bar{\nu}_{t<2001}}{\bar{\nu}_{t<2001}}$. The first parameter is the proportional growth in the happiness-income gradient as a result of the change in disclosure in 2001. This is the main parameter estimated in this study. The second parameter is the effect of the change in disclosure on income visibility.[58] Because we do not have an estimate of this second parameter, we present results assuming different values for it.

First, we estimate a lower bound for the value of income comparisons. Note that $s_{t\geq 2001}$ is strictly increasing in $\bar{\nu}_{t<2001}$. Thus, by assuming $\bar{\nu}_{t<2001} = 0$, we can estimate a lower bound on $s_{t\geq 2001}$. This is a conservative lower bound because it is highly unlikely that income information was completely private before 2001. Assuming that $\bar{\nu}_{t<2001} = 0$ implies that, after 2001, income comparisons explain at least 22% of the happiness-income gradient (i.e., $s_{t\geq 2001} = \frac{0.090}{0.311+0.090}$) and 17% of the life satisfaction-income gradient (i.e., $s_{t\geq 2001} = \frac{0.122}{0.585+0.122}$). These results suggest that the value of income comparisons is bound to be economically significant.

Second, we estimate an upper bound for the value of income comparisons. As $\frac{\bar{\nu}_{t\geq 2001} - \bar{\nu}_{t<2001}}{\bar{\nu}_{t<2001}}$ approaches $\frac{\alpha_2}{\alpha_1}$ from above, both $s_{t<2001}$ and $s_{t\geq 2001}$ converge to 1. That is, a change in visibility

[58] The formula for $s_{t<2001}$ from (B.8) can be interpreted as a Wald estimate: i.e., the ratio between the effect on the happiness-income gradient (i.e., the reduced form effect) and the effect on visibility (i.e., the first stage effect).

of $\frac{\alpha_2}{\alpha_1}$ would imply that income comparisons explain the entire relationship between income and well-being. In the case of happiness, we would need to assume that visibility increased by 29% as a result of the change in disclosure. In the case of life satisfaction, we would have to assume that visibility increased by 21%. Given all the evidence about the widespread use of the search tool, it seems somewhat unlikely that the publication of tax records increased income visibility by just 21%. In this sense, our results suggest that income comparisons may not be the only factor mediating the effect of income on well-being.