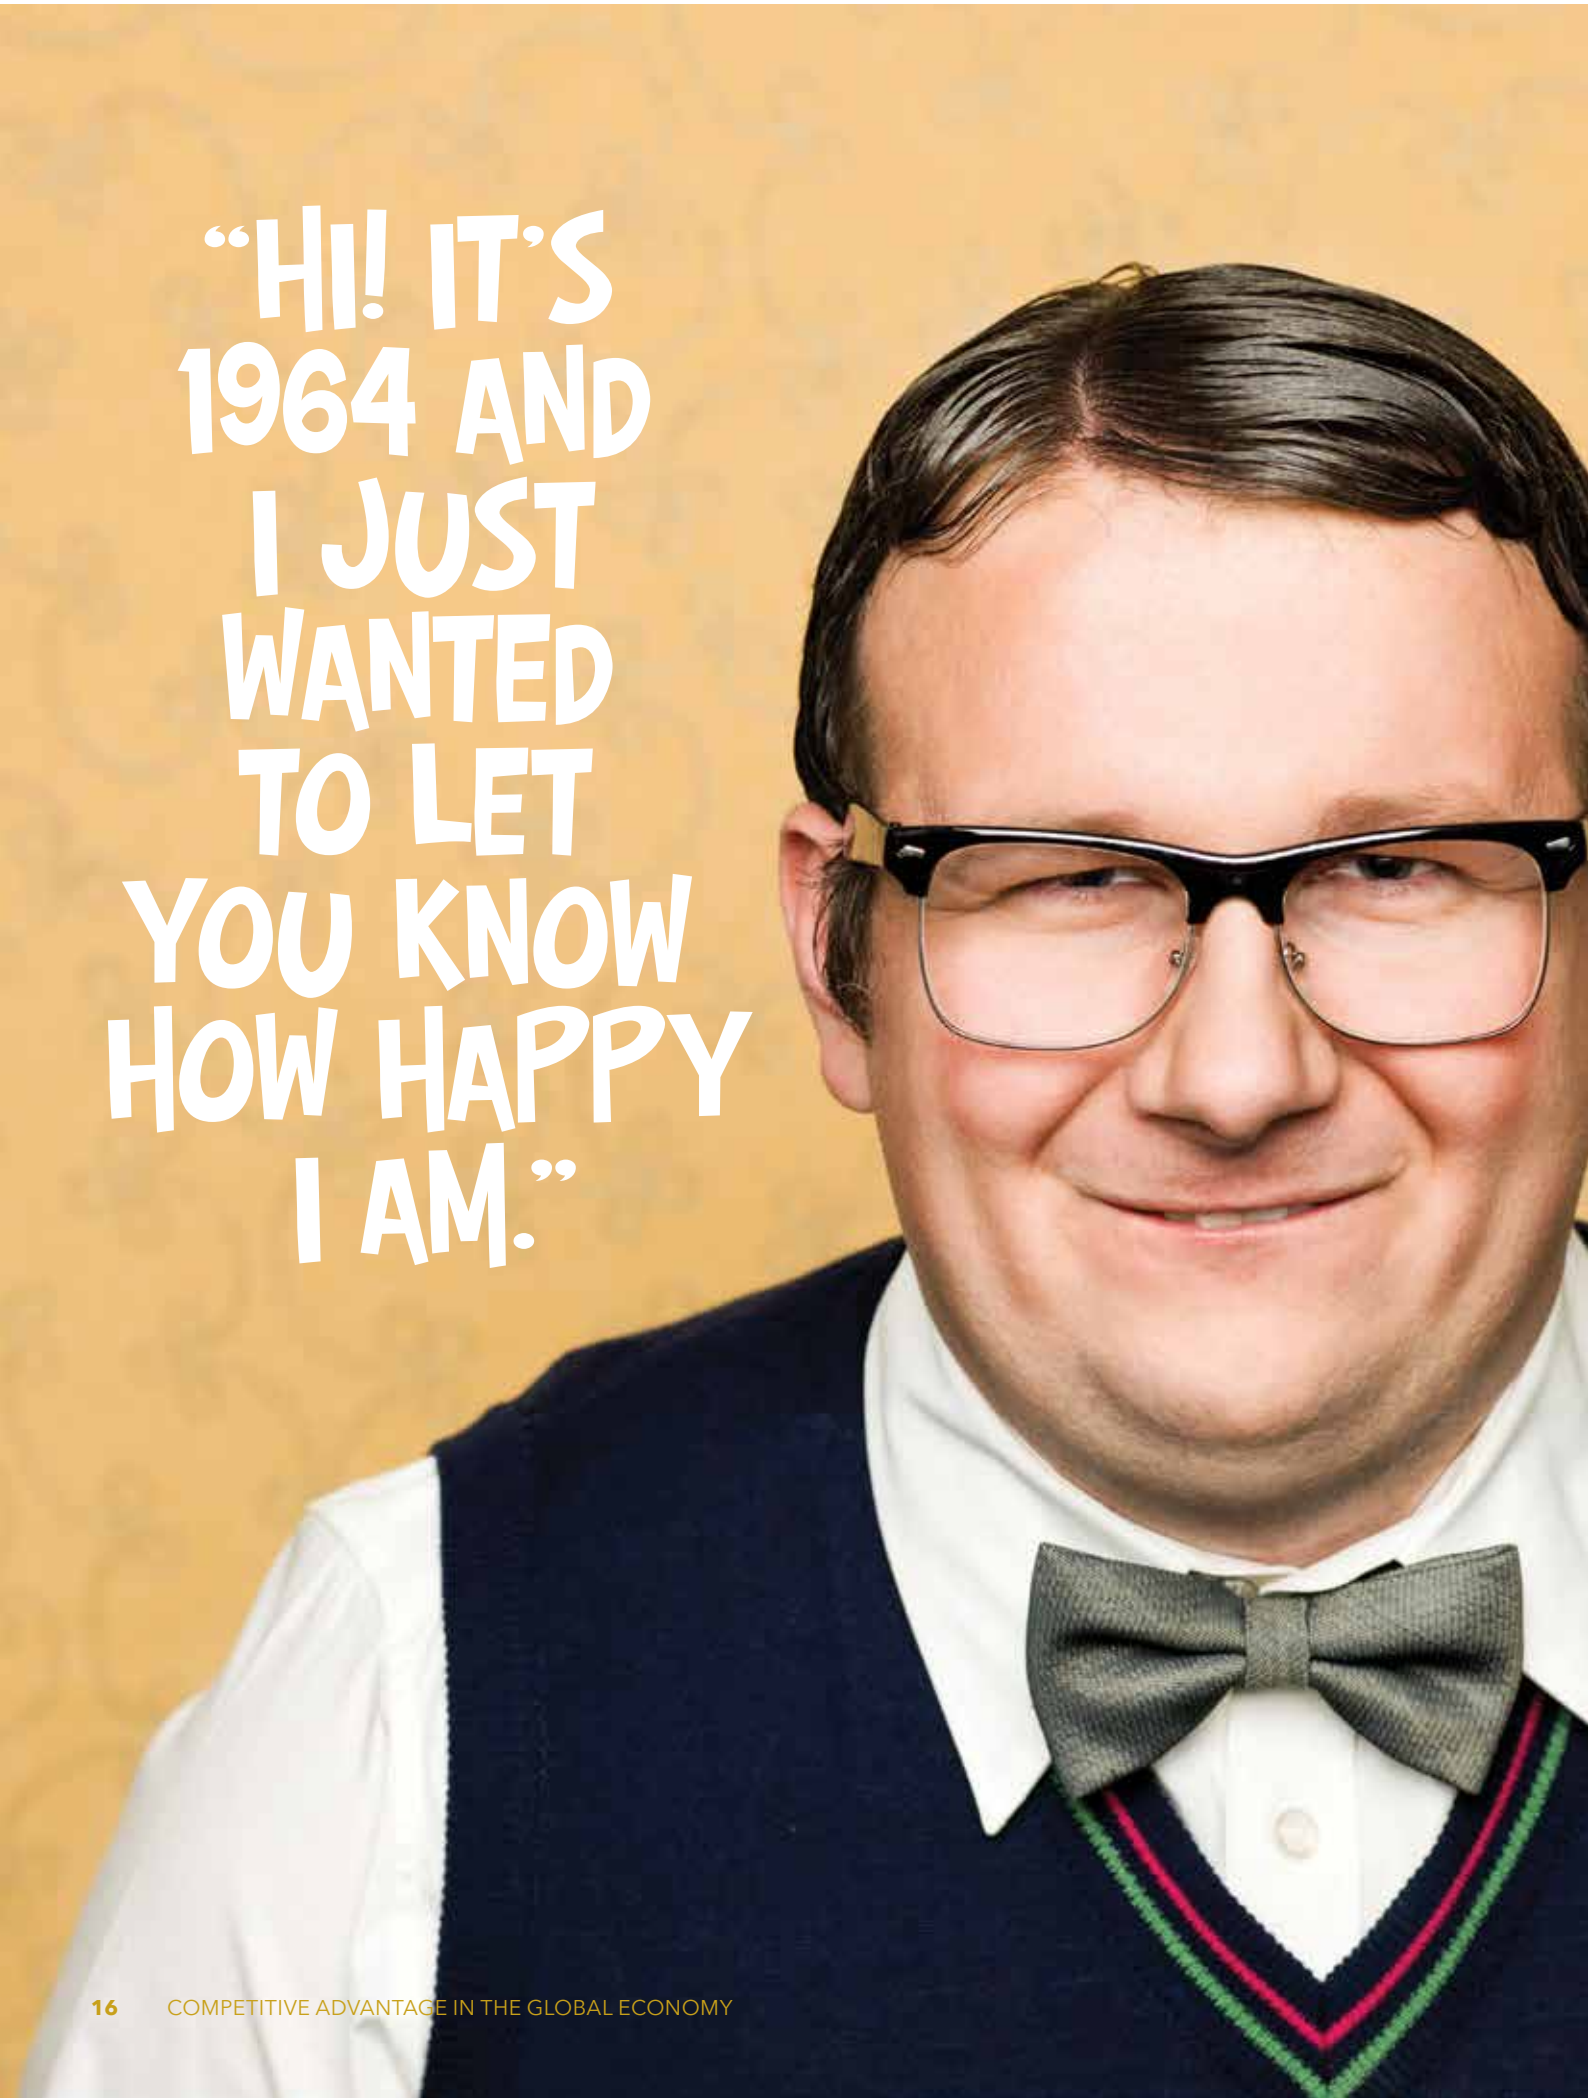
"HI! IT'S 1964 AND I JUST WANTED TO LET YOU KNOW HOW HAPPY I AM."

Measuring historical happiness using millions of digitised books

By Daniel Sgroi

New research demonstrates that through careful analysis of the words people were using in the past, applied to new sources of Big Data, we can start to build a long-run measure of subjective wellbeing.

HAPPINESS HAS LONG SINCE moved from being considered a fringe concern for economists to being a major policy objective. 2011 saw the launch of the UN World Happiness Report and the OECD's Better Life Index and at the level of national governments, the UK is leading the way in thinking about how to take national happiness seriously as a policy objective. However, despite the best of intentions national happiness suffers when compared to national income in one major regard: we have very little historical data. Without historical data we will always struggle to understand what truly drives happiness, and how major shocks or government policies affect happiness at the level of the nation-state. But how can we ever access historical data on happiness? The standard method to measure happiness is to survey opinion: surely surveying opinion from past generations is impossible?

The key insight in our work is that language conveys sentiment, and that the growing availability of digitised text provides unprecedented resources to construct a quantitative history of wellbeing based on historical language use. In particular, the foundation of our work involves combining multiple large corpora of natural language going back two centuries with state-of-the-art methods for deriving public mood (i.e., sentiment) from language. ▶

Without historical data we will always struggle to understand what truly drives happiness, and how major shocks or government policies affect happiness at the level of the nation-state.

The recent digitisation of books, newspapers, and other sources of natural language – such as the Google Books Ngram database – represent historically unprecedented amounts of data on what people thought and wrote over the past few centuries. These databases have already proved fruitful in detecting large-scale changes in language, which in turn correlate with social and demographic change.

These data offer the capacity to infer public mood using *sentiment analysis*. Deriving sentiment from large collections of written text represents a growing scientific endeavour. Examples include recovering large-scale opinions about political candidates, predicting stock market trends, understanding diurnal and seasonal mood variation, detecting the social spread of collective emotions, and understanding the impact of events with the potential for large-scale societal impact such as celebrity deaths, earthquakes, and economic bailouts. Applying the same methods to historical text we can begin to produce more quantitative accounts of national happiness.

In the approach we took, sentiment measures were based on valence norms for thousands of words. These already exist in the literature and are collected from a large group of individuals who are asked to rate a list of words on how those words make them feel. In the present case, valence norms based on the Affective Norms for English Words have already been collected for five languages: been collected for six languages: French, Spanish, Italian, German and separately for (British) English and American English.

We applied these norms to the Google Books corpus for each of these languages, allowing us to derive proxies for subjective wellbeing going back to 1776.

An initial comparison with subjective wellbeing collected with survey data is shown in Figure 1. The data reflect the residuals after controlling for country fixed effects and clearly show a strong

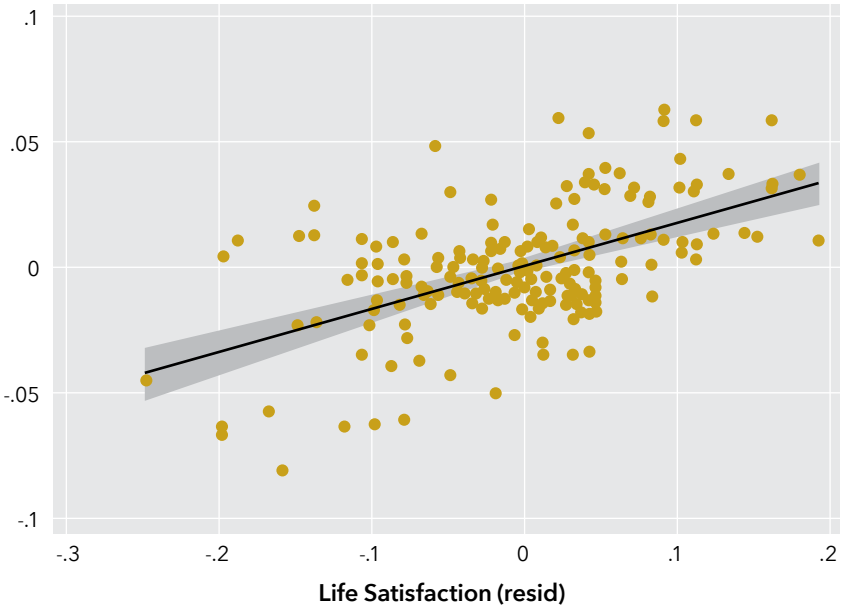

Figure 1. Comparison between survey measures of life satisfaction and residuals (after controlling for country fixed effects) for our measure based on sentiment from historic text. The grey area represents the 95% confidence interval.

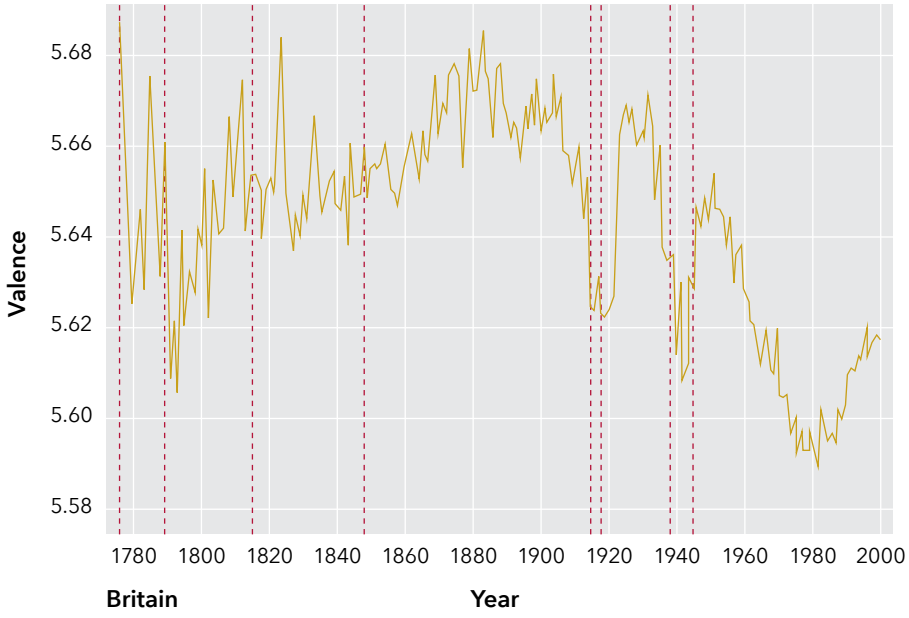

Figure 2. The Average Valences Over the Period 1776-2000 Vertical red lines correspond to 1789, the year of the French Revolution, the Napoleonic Wars (1803-15), the year of the revolutions (1848), World War I (1915-18) and World War II (1939-45).

and significant correlation with our measure based on historic language. This is reinforced in regressions, clearly showing that our measure is very significantly linked to life satisfaction measured from survey data where both are available.

Rolling the text-derived measures of subjective wellbeing back to 1776 reveals a quantitative picture of how public sentiment has changed across the six countries. Glancing at Figure 2 we can see the 1920s, the depression era, and World War I and II show clear and distinguishable influences on subjective wellbeing in the UK (and we can see similar patterns across the other countries we have investigated). We can also see the boost to happiness after World War II (a period of high aspirations) and the fall back (perhaps as those aspirations fail to be achieved) to the trough during the 'Winter of Discontent'. While we warn against super long-run comparisons (aspirations have changed so much over the last few centuries) we can see much in our index that makes sense.

Why is a quantitative history of wellbeing important?

The fledgling state of wellbeing data has limited our collective ability to understand how wellbeing responds to different historic events. This has in turn limited the use of wellbeing in public policy, health initiatives, and financial decision making. In practice, if subjective wellbeing is to become a key factor in guiding our collective behaviour, then we need accounts of wellbeing on par with those of GDP.

Using wellbeing as a measure to guide behaviour, however, takes more than the desire to simply improve wellbeing. As noted by Daniel Gilbert in *Stumbling on Happiness*, people have problems understanding what is called *affective forecasting* – the ability to understand how one will feel in the future – and with this also comes a limited capacity to understand how prior events and decisions influenced our past happiness. To overcome this, especially at the level of government, we must develop our capacity to predict how wellbeing responds to both deliberate and unexpected events. Better predicting economic fortunes was the motivation of the national income accounting following the depression in the 1930s, which later became GDP. Of course, now numerous decisions are based on GDP, despite a near global acceptance that, in the words of John F. Kennedy, "it measures everything in short, except that which makes life worthwhile" (Presidential Library and Museum, n.d.). Thus, like GDP, governments and other agencies recognise the importance of this additional 'emotional accounting' and, by all accounts, they want to understand how better to use it to improve future wellbeing. But to do that, we need historical data. ◀

We can also see the boost to happiness after World War II (a period of high aspirations) and the fall back (perhaps as those aspirations fail to be achieved) to the trough during the 'Winter of Discontent'.

The Author

Daniel Sgroi is an associate professor at the University of Warwick and a research theme leader at the Centre for Competitive Advantage in the Global Economy

Publication Details

This article is based on research presented in CAGE Working Paper 236/2015, "Historical Analysis of National Subjective Wellbeing using Millions of Digitised Books", by Thomas Hills, Eugenio Proto and Daniel Sgroi.