

"Happiness economics" in reverse: Does happiness affect productivity?

Daniel Sgroi

Happiness economics typically looks at how macro-level variables such as economic growth affect happiness. This column turns such thinking on its head and asks whether a rise in happiness might change behaviour at the micro-level, looking specifically at productivity. Experiments suggest that happiness raises productivity by increasing workers' effort. Economists may need to take the emotional state of economic agents seriously.

One of the biggest growth areas in economics over the last few years has been "happiness economics". A plethora of intriguing results, starting with Easterlin's famous paradox breaking the link between well-being and income (Easterlin 1974) through to recent counterintuitive findings showing that "happier countries" produce more suicide cases (Oswald 1997), all show that we still do not have a clear mental grip on how mood links to economic variables such as income or economic growth.

What we do know is that these links are often far from what we expect. Some economists have taken this to suggest that we need to broaden our key macro-economic objectives and include some measure of subjective well-being (Stiglitz et al. 2009).

Working with my colleagues Andrew Oswald and Eugenio Proto, I am seeking to better understand the micro-level links between happiness and economic behaviour. Crucially we turn the traditional causal relationship between production and happiness on its head—rather than investigate how economic variables, like income or growth affect happiness, we have instead studied how happiness affects economic variables. Thinking about which economic variable to examine first, we had only to consider Paul Krugman's words, "Productivity isn't everything, but in the long run it is almost everything."

With this in mind, my colleagues and I address the question: does happiness make people more productive in a paid task? (Oswald et al. 2009). Using an experimental methodology, we find that it does. We demonstrate this in a piece-rate "white-collar" setting with otherwise well-understood properties. The effect operates through a rise in sheer output rather than in the per-item quality of the laboratory subjects' work. Effort increases. Precision remains unaltered. In our first experiment, we induce short-run shocks to happiness and find a pronounced positive effect on productivity. In our second, we turn to the longer run when we ask whether subjects who have experienced recent "life shocks" perform significantly differently, and again, we find that they do.

Experiment 1: A short-run boost to happiness

Let's start by focusing on Experiment 1, i.e. a short-run mood-induction experiment. We used a video clip, composed of a ten-minute comedy routine, to induce a short-run happiness shock in some of our subjects. Our subject-pool was composed of a group of 276 subjects drawn from the University of Warwick. Of these, 182 took part in the main experiment, while the others participated in further sessions designed to test robustness to different payment regimes and alternative movie clips. The subject pool in the main version of experiment 1 was made up of 110 males and 72 females. Our measure of productivity was to ask subjects to correctly add 5 2-digit numbers as many times as possible within 10 minutes. The timeline for Experiment 1 is summarised in Table 1.

Table 1. Timeline for Experiment 1

Stage	Activity
1: Entry	Subjects are welcomed and given basic instructions on etiquette.
2: Happiness Shock	Subjects in treatment 1 view a comedy clip for 10 minutes, otherwise not.
3: Instructions	Subjects are given additional instructions, including a statement that their final payment relates to the number of correct answers.
4: Productivity Task	Subjects undertake numerical additions for 10 minutes.
5: Ability Control	Subjects undertake a GMAT MATH-style test for 5 minutes.
6: Questionnaire	Subjects complete the final questionnaire.
7: Experiment Ends	The experiment ends and subjects are paid.

The first question we might ask is, were we successful in boosting happiness using the comedy clip. Glancing at Figure 1 the answer seems an unreserved yes. In fact those subjects who saw the movie reported a happiness level of 0.99 points higher on a 7-point scale.

Next and most importantly, what was the impact of the treatment (viewing the happiness-inducing video) on productivity? Glancing at Figure 2 we see the effect was a significant and sizeable boost to productivity. In fact, 10.49% more additions were successfully completed by our treated group.

Figure 1. Reported happiness in Experiment 1





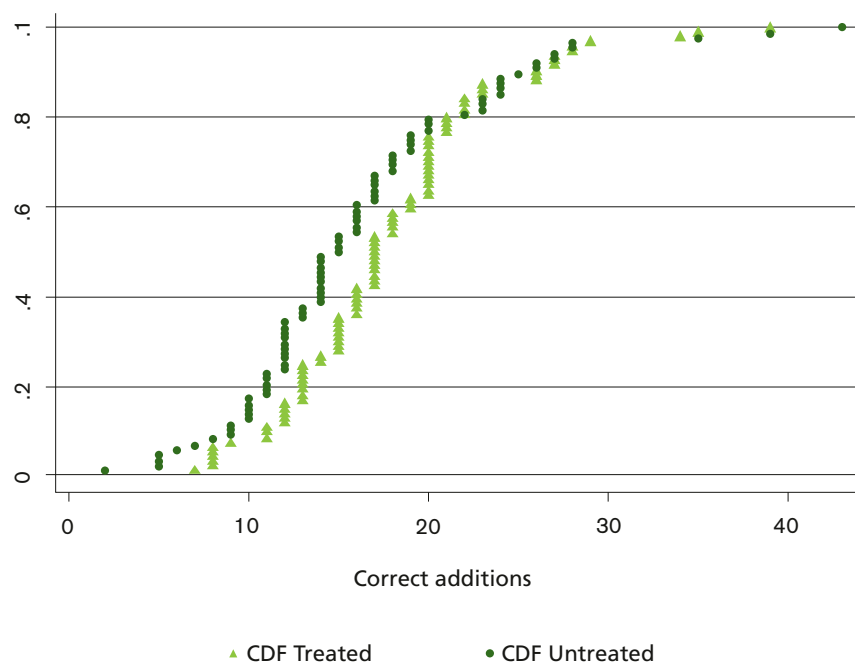
Figure 2. Number of correct additions in Experiment 1



To better understand the underlying causes we ran probit regressions to get a feel for what was going on. This confirmed that witnessing the clip was indeed highly significant even controlling for underlying ability (for which we used school-level results, prior maths ability, and performance in the GMAT MATH-style test), gender, and session. The impact remained at around a 10% boost to productivity. The effect was equally strong across male and female subjects, and it was strong across the full-range of ability. Looking at the distribution of correct answers in the untreated and treated group we have something that looks like first order stochastic dominance, i.e. a rise in performance from the lowest to the highest ability and all the way through the distribution as shown in Figure 3.

Wanting to delve deeper we split productivity into two parts: the number of attempts, which we call "effort", and the probability of each attempt being correct, which we call "precision". What we find is that there was no significant increase in precision, but a highly significant increase in effort.

Figure 3. Cumulative density function of subjects' performances (Experiment 1)



Finally, we tested for the robustness of these results. First, we compared explicit payment (revealing the per addition fee) to masked payment (revealing it would be performance-related but nothing else) and then we compared using a 10 minute "placebo" clip composed of dull visual images to no clip for the control. These variations made no difference to the main result.

Experiment 2: Real life-shocks

The shock we exploited for Experiment 1 was short-run and artificially induced. In Experiment 2 we instead exploited natural life shocks such as recent family bereavement, parental divorce, and health problems. A group of 179 subjects drawn from the University of Warwick participated in Experiment 2, none of whom participated in Experiment 1. The timeline for Experiment 1 is summarised in Table 2.

The subjects for Experiment 2 reported similar initial happiness to the untreated subjects from Experiment 1, a useful check on the consistency of our subject pool. We found that happier subjects performed better, and as we might expect those who experienced significant negative life shocks were less happy and those who experienced significant positive life shocks were happier than those who experienced neither. The pattern was consistent with Experiment 1, and also with the notion of hedonic adaptation, i.e. an effect that is positive, significant, and declining through time. The effect was strong if it happened less than a year ago, and becomes insignificantly different from zero after approximately 2 years.

Table 2. Timeline for Experiment 2

Stage	Activity
1: Entry & Instructions	Subjects are welcomed, given basic instructions on etiquette and a statement that their final payment relates to the number of correct answers.
2: Well-being Report 1	Subjects are asked to report their happiness on a 7-point scale.
3: Productivity Task	Subjects undertake numerical additions for 10 minutes.
4: Ability Control	Subjects undertake a GMAT MATH-style test for 5 minutes.
5: Questionnaire (including life shock reports)	Subjects complete the final questionnaire including questions on recent positive and negative life shocks and are paid.
6: Well-being Report 2	In some treatments we ask subjects to report their well-being for a second time before payment.
7: Experiment Ends	The experiment ends and subjects are paid.

There is one surprising caveat to this concerning parental divorce. In a second paper (Proto et al. 2010) we report that recent parental divorce does not have an impact on reported happiness or productivity. This is not due to the low power of our method as the other life shocks did have pronounced effects. We checked for consistency with the British Household Panel Survey and found that there is indeed no clear causal link from the incidence of parental divorce through to reported happiness.

In a third paper (Sgroui et al. 2010) we address a related issue, i.e. the danger that certain survey questions might change reported happiness by bringing to mind recent disturbing or enjoyable events. In short, can we trust survey answers to well-being questions? In some treatments in Experiment 2 we specifically asked for a second subjective well-being report at the end and found that asking our subjects to report significant positive or negative life events did not change their reported well-being, indicating that they had likely already factored these

into their initial reports. We conclude that we can have confidence in research that uses reported happiness data from surveys—an important result given the growing prominence of such work.

Conclusions

“Happiness economics” should not be just about whether macro-variables raise or lower well-being, we also need to look to the micro-level impact of happiness on behaviour. So far we know that happier individuals are more productive, the effect coming largely through increased effort, whether we consider a short-run shock induced in a laboratory or longer-run real-life shocks.

There are various implications for economics in general. First, research may need to pay more attention to the influence of emotions. Second, more effective bridges will have to be built between applied psychology and applied economics. Third, if happiness in the workplace raises productivity, the paper's findings have consequences for firms' promotion policies and internal labour markets. Fourth, if happiness boosts human productivity. This raises the possibility of self-reinforcing spirals that might operate even at a macroeconomic level. Finally, there is still much to do! With my colleagues at Warwick, I have already started looking at how mood affects other forms of behaviour, but based on our productivity findings it seems safe to assume that happier people are very different economic agents from their less contented peers.

References

Easterlin, Richard (1974), "Does economic growth improve the human lot?" in Paul A David and Melvin W Reder (eds.), *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*, Academic Press, Inc.

Oswald, Andrew (1997), "Happiness and economic performance", *Economic Journal* 107:1815–1831.

Oswald, Andrew, Eugenio Eugenio and Daniel Sgroi (2009), "Happiness and Productivity", IZA Discussion Paper 4645.

Proto, Eugenio, Daniel Sgroi, and Andrew Oswald (2010), "Are Happiness and Productivity Lower among University Students with Newly-Divorced Parents?", *Warwick Economic Research Papers* No. 937.

Sgroi, Daniel, Eugenio Proto, Andrew Oswald and Alex Dobson (2010), "Priming and the Reliability of Subjective Wellbeing Measures", *Warwick Economic Research Papers* No. 935.

Stiglitz, Joseph, Amartya Sen, and Jean-Paul Fitoussi (2009), "Report of the commission on the measurement of economic performance and social progress", CMEPSP.

About CAGE

Established in January 2010, CAGE is a research centre in the Department of Economics at the University of Warwick. Funded by the Economic and Social Research Council (ESRC), CAGE is carrying out a five-year programme of innovative research.

The Centre's research programme is focused on how countries succeed in achieving key economic objectives, such as improving living standards, raising productivity and maintaining international competitiveness, which are central to the economic well-being of their citizens.

CAGE's research analyses the reasons for economic outcomes both in developed economies such as the UK and emerging economies such as China and India. The Centre aims to develop a better understanding of how to promote institutions and policies that are conducive to successful economic performance and endeavours to draw lessons for policy-makers from economic history as well as the contemporary world.



This piece first appeared on Voxeu on 26 July 2010.
www.voxeu.org/article/does-happiness-affect-productivity

VOX

Research-based policy analysis and commentary from leading economists

© 2013 The University of Warwick.

Published by the Centre for Competitive Advantage in the Global Economy
Department of Economics, University of Warwick, Coventry CV4 7AL
www.warwick.ac.uk/cage

Designed and typeset by Soapbox, www.soapbox.co.uk