

Happiness and Productivity: Understanding the Happy- Productive Worker

SMF-CAGE Global Perspectives Series
October 2015

Daniel Sgroi

ABOUT THE AUTHOR

Daniel Sgroi is an Associate Professor in the Department of Economics at Warwick University. He is a theme leader of Warwick's ESRC Centre on Competitive Advantage in the Global Economy (CAGE). He obtained his DPhil at Nuffield College, University of Oxford, and worked at the University of Cambridge prior to joining Warwick. Most recently, he has published on a variety of applied economics topics including the economics of wellbeing, the optimal use of information and herding in financial markets in numerous journals including the Economic Journal, the Journal of Labor Economics, the European Economic Review, Games and Economic Behavior, Experimental Economics, the Journal of Economic Behavior and Organization, and the Journal of Economic Theory.

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 ten-year programme of innovative research. The Centre's research 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 in both 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 seeks to draw lessons for policy-makers from economic history as well as the contemporary world.

www.warwick.ac.uk/cage

ABOUT THE SOCIAL MARKET FOUNDATION

The Social Market Foundation is a leading cross-party think tank, developing innovative ideas across a range of economic and social policy. The SMF's current research themes are: Productivity and Growth; Cost of Living; and, Public Service Reform. We champion policy ideas which marry markets with social justice and take a pro-market rather than free-market approach. Our work is characterised by the belief that governments have an important role to play in correcting market failures and that a sustainable market economy rests on social and political foundations that are widely regarded as fair. www.smf.co.uk

SUMMARY

- “Happiness” is now an accepted and important policy objective for governments alongside big aggregate targets such as economic growth or unemployment. However, there is surprisingly little work on the importance of happiness as an *input* and not just an output.
- This briefing reports on experimental work linking happiness to productivity, by Andrew Oswald, Eugenio Proto and Daniel Sgroi undertaken at the University of Warwick involving over 700 participants and recently published in the Journal of Labor Economics (October 2015).
- In three different styles of experiment, randomly selected individuals are made happier either through the use of a short (10 minute) comedy clip or through the provision of drinks and snacks. We check that these methods make the subjects happier (they do) and then go on to show that these individuals have approximately 12% greater productivity than a control group.
- A fourth experiment studies major real-world shocks (bereavement and family illness) and the impact this has on current productivity. Lower happiness is systematically associated with lower productivity. This effect lasts for approximately 2 years. These different forms of evidence, with complementary strengths and weaknesses, are consistent with the existence of a causal link between human well-being and human performance.
- The productivity measure we employ allows us to differentiate between productivity that comes through effort and ability. We find that the main route from happiness to productivity is through increased effort by workers.
- Private sector firms cannot ignore potential productivity gains in the current economic climate, and this might help public sector departments to offset the negative impact of spending cuts. Our findings make it simpler to see what sort of policies might work and which sort of workers will benefit most. Having scientific support for generating happiness-productivity cycles within the workforce should also help managers to justify work-practices aimed at boosting happiness on productivity grounds.

INTRODUCTION

There is now no doubt that “happiness” as a policy objective is here to stay. This should not be surprising. Happiness is after all something that everyone wants – almost by definition. Consider how central the concept of “happiness” is to the United States Declaration of Independence and countless tracts by moral philosophers and political scientists from Plato, Aristotle, and Confucius onwards.

To date the focus has been largely at the aggregate level. Objectives such as building subjective wellbeing measures to run in parallel with GDP, or adding happiness or life satisfaction measures to national surveys have all been about positioning happiness alongside the likes of GDP, unemployment, inflation or economic growth as key government objectives (most notably in 2011 when the UN released a World Happiness Report and the OECD launched a “Better Life Index”). However, just as we have worried in the past about the links between national income or economic growth and national level happiness measures, so too should we be interested in the links between individual-level income or productivity and happiness. The reason why is simple: at the individual level it may be possible to establish the direction of causation: while much of the debate so far has focussed on whether income or economic growth at the aggregate level generate happiness, we turn instead to look at the converse question: can happiness at the individual-level generate productivity improvements? This briefing paper reports on an attempt to address this question (Oswald, Proto and Sgroi, 2015) by running a series of laboratory experiments using subjects in controlled environments.

To put all of this into context, in the post-2008 climate of spending cuts in the public sector, intense competition in the private sector and concern about future sources of economic growth and profit, an opportunity to boost worker-productivity at relatively low cost would seem hard to dismiss. The hope is that this report will provide an impetus for governments and private sector alike to think about what policies might work and which types of worker to target. We will of course have things to say on both of these issues and do so in the final sections of the report.

WHY THE LAB?

A key advantage of examining behaviour in a laboratory is **control**. In the real world a great deal can happen and does happen which makes it hard to establish causation or even a clear relationship. So many factors may come into play that are hard to perceive or measure. For example is happiness boosting productivity, or is productivity boosting happiness? What if a third factor (say salary or the work environment) is boosting both?

In the laboratory these confounding factors disappear since we can take two randomly selected groups of subjects and make one change to their situation: any resulting change in behaviour is then highly likely to be due to that one change. In our experiments we “treated” one group to a “happiness shock” (commonly called a “mood-induction procedure” in Psychology) but left a second control group “untreated”, so when we look at differences in productivity across the two groups it makes sense that the “happiness shock” was causal. Moreover, we can check that it had the effect we wanted, and also examine other groups with placebos or alternative shocks. All of this control is a unique feature of a laboratory experiment which makes them an ideal way to establish relationships and causation.

With this in mind we ran four different types of experiment:

- Experiment I: the focus was on short-run “happiness shocks”, a happiness question to check the “shock” worked, and a productivity task;
- Experiment II: which was similar to Experiment I but also asked happiness questions throughout the lab experiment to try to understand how happiness evolves through time;
- Experiment III: identical to Experiment I but using a different form of short-happiness shock in the laboratory.
- Experiment IV: the focus was on measuring severe (un)happiness shocks from the real-world and combining this with a productivity task in the lab.

Following convention we randomly assigned our subjects to different groups; no subject took part in more than a single experiment; individuals were told that the tasks would be completed anonymously; they were asked to refrain from communication with each other; and they were told not to use electronic devices for assistance. Subjects (Warwick University students) were told in advance that there would be a show-up fee (of £5) and the likely range of bonus (performance-related) payments (typically up to a further £20 for the hour’s work). Following

the economist's tradition, a reason to pay subjects more for correct answers was to emphasize they would benefit from higher performance. We wished to avoid the idea that they might be paying back effort - as in a kind of 'reciprocity' effect - to investigators.

One possible criticism of our approach is a lack of realism and this is why it often makes sense to follow a laboratory experiment with a similar experiment "in the field" – in a real-world company. While this will result in a loss of the sort of control we have in the lab, combining the results from the lab and the real world will help to give a complete picture. We discuss this more in the section entitled "Next Steps".

The second issue we need to address is what sort of shocks can we use and how we can measure productivity. We will address these issues next.

MAKING PEOPLE HAPPY

It is easy enough to ask how happy people may be and this is now common across a number of national-level surveys. A standard question (which we use in our experiments) is as follows:

How would you rate your happiness at the moment? Please use a 7-point scale where 1 is completely sad, 2 is very sad, 3 is sad, 4 is neither happy nor sad, 5 is fairly happy, 6 is very happy and 7 is completely happy.


However, what if everyone answers with a "6"? That gives us no variation and so no ability to explain anything using reported happiness. A second issue may be that we might worry that subjects are not giving us a true answer. We can address both of these questions by using a "mood induction procedure" to generate the variation we need in happiness. By making some people happier and leaving some "untreated" by our procedure we can create a happy group and a neutral group. We can check this has worked (using the reported happiness question) and see whether the happier group are more productive.

A mood-induction procedure is a "treatment" or undertaking specifically designed to change someone's mood or temperament. A typical approach would be to subject one group to some text, a movie, a verbal speech, music or some such stimulus designed to elicit a particular mood, and subject a control group to an alternative (which could be nothing or a placebo). Research within Psychology (such as the meta-study by Westermann et al, 1996) indicates

that of the many ways that have been used to induce different moods in experimental subjects the most effective is the use of a movie clip particularly for the induction of a positive mood state (such as making someone “happier”).

Movie clips were used in Experiments I and II. The treatment was a 10-minute clip of comedy sketches by a well-known comedian. The questionnaire results indicate that the clip was generally found to be entertaining and had a direct impact on reported happiness levels. We also have direct evidence that the clip raised happiness through a comparison of questionnaire happiness reports directly before and after the clip. Experiment II in particular asks subjects to report their happiness level throughout the session and provides even stronger evidence that the movie clip is having the effect we want. As a control, we used either a calm “placebo” clip or no clip at all. In fact those subjects who saw the movie reported a happiness level of about 1 point higher on a 7-point scale as compared to those who did not see the clip as shown in Figure 1. Finally, it is useful to notice that the reported level of happiness before the clip for the treated group is not statistically significantly different from the happiness of the untreated group.

Figure 1: Those exposed to the randomized happiness treatment in the laboratory have higher happiness in Experiment II



Notes: here the happiness treatment is a comedy movie clip in the laboratory (the lines indicate 95% confidence intervals).

MEASURING PRODUCTIVITY

The flip-side to being able to induce happiness is being able to design a task that allows productivity to be measured. The task we chose is based on a similar task in Niederle and Vesterlund (2007). Subjects were asked to answer correctly as many different additions of five 2-digit numbers as possible in 10 minutes, with a typical example as follows:

31	56	14	44	87	Total =
----	----	----	----	----	---------

Figure 2: Adding Five 2-digit Numbers under Timed Pressure

This task is simple but is taxing under pressure. We think of it as representing, admittedly in a stylized way, a white-collar job: both intellectual ability and effort are rewarded. One nice feature of our chosen task is that it allows us to measure productivity in three ways:

- The number of correct additions
- The number of addition attempts
- The percentage of correct additions

The first of these can be considered the overall measure of productivity and will be determined by the other two. The second (attempts) is a simple proxy for **effort**: if this increases holding the percentage that are correct constant, the overall number of correct additions will still rise, but we can argue that it is effort that is driving any productivity increase. The third measure (percentage of correct additions) is a simple proxy for **ability** (or precision). If the percentage correct rises, but the number of attempts stays constant, the overall number of correct additions will rise, but we can argue that it comes through a rise in the ability to get the additions correct and not a rise in effort.

To supplement our core measure of productivity, we also used a short mathematical reasoning test as a supplementary control on ability. This involved asking subjects 5 questions along similar lines to that of Gneezy and Rustichini (2000) where a mathematical GMAT-style reasoning test was administered. Subjects had 5 minutes to complete this and were paid at a rate of £0.50 per correct answer. We also asked our subjects to tell us how well they performed in school-level mathematical ability and for their overall exam performance. Finally, the fact that subjects were randomly allocated to groups, and we used several hundred subjects, means that ability-types should be evenly distributed between the group that was exposed to

the mood-induction procedure and the group that was not.

SHORT-RUN SHOCKS (EXPERIMENTS I-III)

We can group the first three sets of experiments together since they all deal with short-run shocks induced within the lab and share similar features.

In **Experiment I**, in which 276 subjects participated in total, we used a comedy clip as our “mood induction procedure”. The control-group individuals were not present in the same room with the treated subjects; they never overheard laughter or had any other interaction. The experiment was carried out with deliberate alternation of the early and late afternoon slots to avoid time-of-day effects. Table 1 provides a timeline of events during the experiment.

Table 1: Timeline for Experiment I

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 mathematical reasoning test for 5 minutes.
6: Questionnaire	Subjects complete the final questionnaire and are paid.
7: Experiment Ends	The experiment ends and subjects are paid.

The final questionnaire inquired into both the happiness level of subjects (before and after the clip for treatment 1), and their level of mathematical expertise. The core sessions took place over 4 days. We then added 4 more sessions in two additional days designed to check for the robustness of the central result to the introduction of an explicit payment and a placebo film (shown to the otherwise untreated control group). Subjects received £0.25 per correct answer on the arithmetic task and £0.50 for each correct mathematical reasoning answer.

We used two different forms of wording:

- For days 1-4 we did not specify exact details of payments, although we communicated clearly to the subjects that the payment did depend heavily on performance.

- For days 5-6 the subjects were told the explicit rate of pay both for the numerical additions (£0.25 per correct answer) and the mathematical reasoning questions (£0.50 per correct answer).

This achieved several things. First, for days 5-6 we have a revealed-payment setup, which is a proxy for many real-world piece-rate contracts, and for days 1-4 we mimic those situations in real life where workers do not have a contract where they know the precise return from each productive action they take. Second, this difference provides the opportunity to check that the wording of the payment method does not have a significant effect, thereby making one set of days a robustness check on the other.

In Experiment I it is not possible to observe how subjects' happiness evolves in real time, although they are asked some retrospective questions, the happiness levels of individuals before and after the comedy movie clip. To deal with this, we designed **Experiment II**, in which 104 subjects participated. We ask happiness questions before playing the movie clip, and then longitudinally throughout the experiment, which represents a departure from the Experiment I timeline in Table 1 since we asked subjects about their happiness level on three occasions. The initial measurement was at the very start of the experiment. The second was immediately after the comedy or placebo film. The third time was at the end of the experiment. Experiment II used explicit payment instructions and a placebo clip (which provided a gap between two of the happiness questions for the control subjects).


In **Experiment III**, a different treatment was adopted. In an attempt to mirror somewhat more closely, admittedly still in a stylized way, the sort of policies that might potentially be provided by actual employers, our treatment subjects were provided with chocolate, fruit, and drinks as the mood induction procedure, otherwise everything from Experiment I remained the same. The experiment involved an additional 148. We offered the selection of snacks and drinks to the treatment group (comprising 74 subjects in 4 sessions) who had 10 minutes to consume whatever they chose. We provided nothing for the control group (who were a different set of 74 individuals, also in 4 sessions). The control group were still asked to sit for 10 minutes prior to the experiment beginning; this was to ensure that any effect was not due to the additional minutes of experimental time for the treated group.

RESULTS FROM EXPERIMENTS I-III

From Figure 1 in the last section we know that the mood induction procedure was effective at raising happiness levels by about 1 point on a 7-point scale, but does the increase in happiness carry through to a boost to productivity?

Glancing at Figure 3 we see the answer is yes. In fact, over 10% more additions were successfully completed by our treated group.

Figure 3: Those exposed to the randomized happiness treatment in the laboratory have higher productivity in Experiment I



Notes: the happiness treatment is a comedy movie clip in the laboratory (lines indicate 95% confidence intervals).

We next analysed the data in more detail in an attempt to understand better the underlying causes, confirming 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 mathematical reasoning 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 ability range indicating that the boost from happiness applies equally well to high-productivity and low-productivity workers.

Wanting to delve even deeper we exploited our ability to analyse the two different parts of productivity: effort (attempts) and ability (precision). What we found was that there was no significant increase in precision, but a highly significant increase in effort due to the happiness boost.

We also 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 the 10 minute “placebo” clip or no placebo. These variations made no difference to the main result.

The results from Experiment III mirrored those from Experiment I. As before, a positive productivity effect was produced, and again the size of that effect was substantial: having spent approximately two dollars per person on drinks and snacks, productivity was boosted by almost 20% for a short period of concentrated work. This helped to confirm that the result was not reliant on any particular type of happiness shock.

In Experiment II we tracked happiness throughout the session, making it much easier to look at causation (rather than just correlation). As people are randomly assigned to the treatment group, we know that the baseline levels of productivity of the treatment and control group are identical. It is therefore possible to find out whether there is a link between their measured rise in happiness and the measured implied effect on productivity. Glancing at Figure 4 we can see a positive relationship (an upward slope) between the difference in happiness before and after the “shock” and the resulting boost to productivity. This provides a final observation: the more powerful the effect of the happiness shock, the greater the final boost to productivity.

Figure 4: The greatest rise in happiness during Experiment II is associated with the greatest productivity gain



Note: here those not exposed to the happiness treatment have the same baseline productivity; hence the y axis can be viewed as a change in productivity from the common baseline.

EXPERIMENT IV: FAMILY TRAGEDIES AS REAL-LIFE SHOCKS

The preceding experiments studied small happiness interventions. In **Experiment IV** we study real-life unhappiness events assigned by Nature rather than by us for 179 subjects (8 sessions in 2 days). These shocks, for which we use the generic term “Bad Life Events” (or BLE) are family tragedies such as recent bereavement or serious health issues in the immediate family.

In Experiment IV, subjects’ productivities are measured at the very outset with no “mood induction procedure”. At the end of the experiment, these subjects are quizzed, by questionnaire, about recent tragedies in their families’ lives. One other key aspect is that we asked subjects to report *their level of happiness right at the start of the experimental session*. This was to avoid ‘priming’ problems. The underlying logic is that we wanted to see if people’s initial happiness answers could be shown to be correlated with the individuals’ later answers and behavior. A final questionnaire included supplementary questions designed to find out whether they had experienced at least one of the following BLEs: close family bereavement, extended family bereavement, serious life-threatening illness in the close family, and/or parental divorce. The data suggested that it was appropriate to aggregate all of these happiness-shock events by using a single BLE variable. Bennedsen et al. (2010) also considered something similar for CEOs and suggested that company performance may be impeded by traumatic family events. Table 1 presents the timeline for Experiment IV.


Table 2: Timeline for Experiment IV

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 mathematical reasoning 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.

RESULTS FROM EXPERIMENT IV


Those subjects who had been through a bad life event within 3 years were noticeably less happy and less productive as shown in Figures 8 and 9 respectively. The subjects for Experiment IV who did not suffer a BLE reported similar initial happiness to the untreated subjects from Experiment I, a useful check on the consistency of our subject pool. Compared to the control group, those who did suffer a BLE mark themselves nearly half a point lower on the happiness scale, and they achieve on average 2.31 fewer correct additions. They also make fewer attempts which seems to drive the reduction in productivity and the effect seem similar for males and females.

Figure 8: Individuals with a recent Bad Life Event have lower productivity in Experiment IV



Note: a BLE is bereavement or family illness (lines indicate 95% confidence intervals).

Figure 9: Individuals with a recent Bad Life Event report lower happiness in Experiment IV



Note: a BLE is bereavement or family illness (lines indicate 95% confidence intervals).

We also analysed the importance of when the BLE took place. The immediate-run (<1 year) loss of happiness is approximately one full point on our 7-point scale. Therefore, although our subjects may not be aware of it, their happiness answers at the start of Experiment IV are correlated with whether later on they report that a BLE recently occurred in their family. The pattern in the happiness coefficients is itself broadly consistent with hedonic adaptation - the well-being effect declines through time. Overall, the consequence of a bad life event is empirically strong if it happened less than a year ago, and becomes insignificantly different from zero after approximately 3 years. Our results are consistent with a range of adaptation findings in the survey-based research literature on the economics of human well-being (e.g. Clark et al. 2008).

We are especially interested in the effects of a bad life event upon human performance. Having had a bad event in the previous year is associated with particularly low performance on the additions task, our results indicate that the size of the productivity effect is large; it is typically more than two additions and thus greater than 10%. The extent of the deleterious effect of a Bad Life Event upon subjects' productivity is a declining function of the elapsed time since the event.

In the questionnaire, we also asked about parental divorce thinking this would also be a reasonable BLE to consider. However, this turned out to have a tiny (occasionally positive) and statistically insignificant effect on the individual, so the divorce of parents, at least in our data set, does not appear to qualify as a "bad life event". We discuss this finding in greater detail in Proto, SgROI and Oswald (2012) where we also 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. Though we do not know why this was the case, it might have been perceived by our subjects as a release from a more difficult situation and/or may also have been a longer-term issue granting additional time for the subjects to get used to the situation, but these are only conjectures.

POLICY RAMIFICATIONS

Much of what we have done is not just about establishing a link from happiness to productivity but about thinking more deeply about how such a link might work and trying to characterise features of the process. Let's consider the key findings distributed across the four experiments.

- **The effect is big:** In terms of national GDP or economic growth, rises of 3% or so are considered very large. The impact of the shocks we used were of the order 10%-20%. The induced shocks were relatively cheap (e.g. £2 per subject) but the scale of the effect makes it likely that even close to zero-cost “nudges” would have impact. This is an important idea to test in practice.

What does this mean for the private and public sector? The scale of the effect cannot be dismissed. In the post-2008 environment of public sector cuts a key question is the impact spending cuts will have on the long-run sustainability of important welfare and social policies. By boosting worker-productivity it might be possible to partly offset the impact of spending cuts and also to help rebound back as and when more funding is available. In the private sector competitive pressure means that productivity may be decisive for survival.

- **The long-run:** Happiness shocks (our “Bad Life Events”) from the real world, if they are significant enough, can generate an effect on productivity of up to 3 years in duration.

A more or less permanent “shock” will generate a long-run impact. This would suggest that significant environmental changes might work into the long-run while smaller shocks like the ones we trialled in the lab will produce a strong but possibly temporary effect. Given the scale of the short-run effects, it is possible that this might be driven into the long-run through a positive cycle of happiness and productivity.

We would warn against misrepresenting the “bad life event” study as suggesting that firms should avoid hiring workers with a recent tragic past. It could just as easily be argued that it is just these sorts of workers who would benefit the most from a positive happiness shock if their happiness levels are low to begin with: so we would not advise using any of our findings to help choose *who* to employ (which would not be acceptable morally nor most likely legally in any case).

- **Effort not ability:** Making subjects happy seems to encourage them to put in more effort, it does not make them better at the task. This is an important distinction as some tasks will benefit more from increased effort than others.

Our findings to date focus on the type of work for which effort might be especially important such as simple piece-rate easily divisible tasks which require some mental capacity. We can combine this with existing psychological research, such as Amabile et al (2005), which suggests that creativity is also boosted by happiness. This yields a wide range of possible jobs to consider, but at least from our research we would prioritize tasks where effort is key.

- **The effect works across the entire distribution:** Irrespective of how highly productive workers may be initially, making them happier should help make them even more productive. Moreover, the greater the boost to happiness, the greater the resulting boost to productivity.

The universal impact of happiness-boosting policies suggests that it would make sense to target the entirety of workers in the firm rather than have selective policies or focus on especially high or low productivity workers. To some extent this might complement policies such as “worker of the month” awards or bonus schemes but care needs to be taken to avoid clashes of impact such as the negative happiness implications (“I just missed a bonus this month!”) vs the powerful incentive-effects of performance-related pay. It may well be easier to boost happiness towards the bottom end of the happiness distribution, and so a policy of trying to bring everyone up to a certain happiness level would seem reasonable, and good human resources divisions should be thinking in terms of how to spot and help those who are suffering especially low levels of happiness.

PRACTICAL CONSIDERATIONS

Any laboratory experiment is a simplification of a real firm. Moving to real-world settings introduces complexity, which is one of the reasons why the lab is a good place to start any such research. However, attempting to study these ideas in the workplace means thinking about several important practical considerations.

- ***How can we measure productivity?*** The firm needs to have a simple measure of productivity. This can be formed either directly from output or through a measure of the time taken to achieve a task.
- ***How can we boost happiness in a workplace?*** Many possibilities exist from improvements in the working environment, better interaction between workers, between

management and workers, and positive feedback or encouragement. The issue may be not whether a sensible procedure for boosting happiness exists, but which one to select.

- **The procedure needs to be cheap.** Any costly boost to happiness may crowd out the benefits to productivity. This brings us into the realm of “nudge economics”: in which the aim is to influence behaviour through policies which cost almost nothing. One option might be short verbal or written positive feedback designed to offer encouragement.
- **Time.** Many issues relate to time, such the duration of the mood-induction and how often it needs to be repeated as well as how quickly workers become habituated to any procedure or change. On the output side we would need to consider the duration of any productivity impact and how this relates to the nature of the procedure used.
- **Control vs treatment.** In the real world we need to make sure that the control group are as similar as possible to the “treated” group. Ideally we need two groups of workers who do the same thing in more or less the same way and in similar environments or we might confuse the direct impact of happiness with complex interactions between mood-induction and the environmental differences between groups.

Another important issue for any firm or manager considering applying happiness boosting policies is to understand the pathways from happiness through to productivity. The results of Killingsworth and Gilbert (2010) suggest the possibility that unhappiness may lead to a lack of mental concentration. The appendix to Oswald, Proto and Sgroi (2015) provides a model that explores this idea further. Its main result stems from internal resource-allocation by the worker. In the model, a positive happiness shock allows the employee to devote more attention and effort to solving problems at work (essentially because the worker can switch from worrying) which can explain our key findings.

CONCLUSIONS

If happiness in the workplace brings increased returns to productivity, then human resource departments, business managers and the architects of promotion policies will want to consider the implications. These are not idle questions. With small increases compounding like interest in a bank account, productivity is a crucial component in the quest to raise standards of living. To borrow a phrase from Nobel laureate Paul Krugman, “Productivity isn’t everything, but in the long run it is almost everything.”

In this briefing paper we described four kinds of trial. The last of these is an attempt to estimate the repercussions of life-events assigned by Nature. The design, in that case, has the disadvantage that we cannot directly control the happiness shock, but it has the advantage that it allows us to study large shocks of a fundamental kind in real human beings' lives. The other three experiments examine the consequences of randomly-assigned "happiness shocks". These experiments have the advantage that we can directly control the happiness shock but the disadvantage that shocks are inevitably small and of a special kind in the laboratory. It is conceivable in the last experiment that there is some unobservable feature of people that makes them both less productive and more likely to report a bad life event. Yet such a mechanism cannot explain the results in the other three experiments. By design, the four experiments have complementary strengths and weaknesses.

Various practical implications emerge from the experimental results. If happiness in a workplace carries with it a return in productivity, managers and human resources specialists who may wish to implement their own "happiness shocks" to help maintain productivity in the context of a highly competitive economic climate in the private sector, and spending cuts in the public sector. We have suggested ways to think about this and practical considerations that need to overcome. If well-being boosts people's performance at work, this also raises the possibility, of self-sustaining spirals between human productivity and human well-being which could help even quite small short-run shocks generate longer-run gains.

Finally, we need to acknowledge the work done by our predecessors. To the best of our knowledge Oswald, Proto and SgROI (2015) is the first examination of the relationship between productivity and happiness to implement a monetary piece-rate setup, examine large real-world shocks to happiness (and not solely small laboratory ones), use a range of different experimental designs, and offer various types of evidence and collect longitudinal data in a way that provides us with an opportunity to scrutinize the changes in happiness within our subjects. However, there exist many important papers spanning several decades that have had an important influence on our work and more generally on our understanding of how happiness affects behaviour. Within Psychology, research by the late Alice Isen of Cornell University has been especially important in this area, for instance Erez and Isen, 2002, and there is also a large literature on productivity at the personal and plant level. For a review of the relevant literature please see Oswald, Proto and SgROI (2015).

REFERENCES

- Amabile, T., Barsade, S., Mueller, J. and Staw, B. (2005), 'Affect and Creativity at Work', *Administrative Science Quarterly*, 50: 367-403
- Bennedsen, M., Perez-Gonzalez, F. and Wolfenzon, D. (2010), 'Do CEOs Matter?', Working paper, Stanford GSB.
- Clark, A., Diener, E., Georgellis, Y. and Lucas, R. (2008), 'Lags and Leads in Life Satisfaction: A Test of the Baseline Hypothesis', *Economic Journal*, 118: 222-243.
- Erez, A. and Isen, A. (2002), 'The Influence of Positive Affect on the Components of Expectancy Motivation', *Journal of Applied Psychology* 89: 1055–1067.
- Gneezy, U. and Rustichini, A. (2000), 'Pay Enough or Don't Pay at All', *Quarterly Journal of Economics* 115: 791-810.
- Killingsworth, M. and Gilbert, D. (2010), 'A Wandering Mind is an Unhappy Mind', *Science* 330: 932-932.
- Niederle, M. and Vesterlund, L. (2007), 'Do Women Shy away from Competition? Do Men Compete too Much?', *Quarterly Journal of Economics* 122: 1067-1101.
- Oswald, A., Proto, E. and SgROI, D. (2015), 'Happiness and Productivity', *Journal of Labor Economics* 33(4).
- Proto, E., SgROI, D. and Oswald, A. (2012), 'Are Happiness and Productivity Lower among University Students with Newly-divorced Parents? An Experimental Approach', *Experimental Economics* 15(1): 1-23.
- Westermann, R., Spies, K., Stahl, G. and Hesse, F. (1996), 'Relative Effectiveness and Validity of Mood Induction Procedures: A Meta-analysis', *European Journal of Social Psychology* 26: 557–580.