

Experimentation, Behaviour Change and Public Policy

**Inaugural lecture to the Department of
Political Science, University College London**

Peter John

Professor of Political Science and Public Policy

November 20, 2012

Abstract

Many of the greatest challenges in public policy can only be fully addressed by changing the behaviour of citizens. Climate change, poor economic growth, obesity and crime are just some of the problems that need action on the part of the public as well as from governments. But governments do not know enough about how to influence behaviour; in particular, the need to find out the most appropriate ways to engage citizens so they consent to and participate in decisions that are made about them. Greater use of experiments randomized controlled trials (RCTs) is one way to increase policy relevant knowledge on behaviour change. Peter John will give examples of experiments he has undertaken, which illustrate how academics can lead the way in the creation of relevant knowledge. He will acknowledge both the strengths and weaknesses of the experimental method, and set out how governments and academics can design and use RCTs in better ways.

Thank you David, thank you Steven. And thank you – colleagues and friends – so many of you - for coming today, and I know many of you have travelled a long way to get here. I am very grateful.

Today we live in a world of large problems of public policy, many of which can only be solved by the willing cooperation of citizens who need to share a common purpose and who are prepared to change their personal behaviour over the long term. Consider the challenge of climate change where a set of human processes operating for over a century have created – and will create – problems that affect us all, and which can only be fully addressed by changes in individual behaviour albeit helped and in partnership with strong actions carried out by public authorities. Or take health, where in spite of rising levels of health world wide and greater technological advances, most societies

are facing more heart conditions, strokes and cancer episodes, which affect both the individuals concerned and also the wider communities that have to bear the cost. Or in the political sphere, where in spite of rising levels of education and ease of transfer of information, there is declining trust in politics and a reduction in traditional forms of political participation, such as membership and engagement with the political parties. The missing ingredient in these three examples of public problems is changed behaviours on the part of the citizens.

My claim is that relatively small adjustments in personal lifestyles and routines, which are not very costly for the individuals concerned, can add up to a massive impact in collective social welfare. This is because small individual changes can add up to a large combined effect. So with climate change small adjustments to consumer behaviour can lead to a large contribu-

tion to carbon emission reductions; relatively small changes in diet and exercise can have health-enhancing effects; and modest rises in participation and voice, such as in complaining, can keep bureaucrats and political leaders on their toes.

There is an added advantage that these small actions might have a chain reaction, for as one behaviour change takes place, people take notice, and a bandwagon effect can take hold, whereby these small actions feed on each other, and the result is greater than the sum of the original acts. So the argument is that many of the negative features of the world today could be addressed by changes in individual behaviour that are not costly for the citizens, do not fundamentally affect the pursuit of personal goals and represent broadly where most people want society to go.

So the question arises – if it is that simple – why do so many

of these challenges remain in public policy? One take on this is that citizens tend to follow their private preferences in spite of saying they believe in publicly-regarding actions, as for example as expressed in opinion polls. So people act to benefit themselves, such as maximising income and increasing leisure activities. It reminds for governments to follow the public interest by providing services run by experts and professionals who can use their knowledge to devise policies that determine what the role of citizens should be. In this conception, while the citizen will be allowed to follow their private interests where possible, when it comes to the role of the state, public officials and politicians should be allowed to limit individual freedom if it serves the public interest. This can involve using the power of the state to regulate and to control vigorously the actions of the individuals. An example would be to be very tough on en-

forcing rules on recycling household waste by fining people who do not sort their refuse into different sacks. Such an approach to behaviour change has a long pedigree, and comes from the idea that the guardians of the state have access to knowledge and a set of public values — what we now call a public service ethos — to implement policies on the citizens' behalf, using their reason to find out what the citizens would consider to be right if they were given a chance to deliberate with full knowledge. On the left, in Britain we have been influenced by the Fabians who believed the state could be harnessed in this way, staffed by the intelligentsia educated with the knowledge from the sciences and social sciences. On the right too, before the Hayekian turn in Conservative thinking that emerged in the early 1970s, there was a strong belief in intervention for the common good. Both these views draw on to the utilitarian tra-

dition of thinking in public policy, which was influential in the founding of UCL — and after this lecture the Dean will invite you to a reception where you will walk past the body of Jeremy Bentham. In fact the popularity of behavioural sciences might be seen to fit with such a technocratic vision whereby the tools of behavioural economics and psychology become harnessed to refashion the state to deal with these problems. One can find echoes of this approach with the recent House of Lords inquiry on behaviour change, for example.

However there are a number of problems with this approach.

The first is that it has an idealistic account of those who are in power. If the task of the state is to control the private actions of the citizens, then whose task is it to control the private acts of politicians and bureaucrats? The solution is not to give more power to those in public life, but to think of ways that those in

power can be held to account and to be made responsive to the citizens, so both are engaged in a joint endeavour to discover the public purpose and to find ways of achieving it.

The second problem is that there may not necessarily be a right answer about what behaviours are best. There is a problem of those in power deciding in advance what is desirable behaviour and then implementing it on an unsuspecting population, because behaviour changes often imply a moral as well as an instrumental purpose. So we find that those who live in deprived estates are subject to the moral decisions about their lifestyles dictated by politicians in power (who may not even privately uphold those values they inculcate to others). This point illustrates that we do not now have a society based on consensus in values: these need to be re-discovered both by citizens and elites in the process of grappling with the large

public problems I started with.

The third reason is that such a top-down approach will not work in today's society is because people believe in primary of their own choices and simply will not be told how to live their lives. The more the state issues commands and controls the more likely it is that people will resist and evade them. There has to be an alternative that does not give up on enforcement and tough action, but brings people with them.

The way forward is to not to see behaviour change as trapped between private vice and public virtue. In practice, citizens **do** want to mix public and private actions together, but often do not carry out as many public acts as they would ideally like to do. There are a variety of reasons for this. One is to see it as a problem of cooperation. Citizens want to do things that are more pro-social but they need to know they are not acting

alone and that others round the world are also going to act at the same time. But they do not know this and they could easily be in the minority that provides the public good while others free-ride. This is what is called the collective action problem, which has dominated much of social science in the twentieth century. When faced with a structure of incentives, individuals will chose the dominant strategy and will not cooperate: the public good is not provided. Common areas get degraded, fish stocks go down, the planet gets hotter, anti-social behaviour get worse, and volunteering and charitable giving do not happen because individuals cannot realise joint benefits.

There is much that can be done to overcome the collective action problem. And a key task of politics is to provide leadership so as to confront it. In the view of the late Elinor Ostrom, the way forward is to redesign institutions so they give the

right signals to help cooperation, and this may need decentralisation rather than the centralisation of power. So rather than the state dictating what the citizens should or should not do, the role of those in power is to provide the institutional venues whereby cooperation can be brought about, so that the citizens decide their own political objectives and discover the way forward. The role of the state is as a facilitator and designer – to encourage the realisation of the common interest rather than to impose it top down.

One interesting example of this approach that occurred in recent years has been the ban on smoking in work places and in many public locations, which came into force in England in 2007. On the face of it, it looks like the state telling people what to do with a law, but what really happened was the mobilisation of consent helped by government and official reports,

and a change in public opinion, so that the ground was prepared for Parliament to act. The role of the law was to signal decisively the change to all concerned, rather than to impose it, dealing with the very small minority who permitted or carried on smoking. The result is clean air in public places, the outcome that has the willing cooperation of most smokers.

In thinking like this, I have been very influenced by the work of Elinor Ostrom, and it links to my long interest in decentralised structures: how collective action may be achieved through appropriate institutional design, which I believe can be done in a way that is consistent with maintaining forms of equality between citizens. But in recent years I have become convinced that the lack of cooperation is also affected by cognitions of citizens as well as by the structure of interests between them. With behaviour change we know that norms, what oth-

ers do, and the habits of the past are powerful determinants of behaviour. And for this we are in debt to psychologists and economists, who have created the field of behavioural economics, for pointing this out. Given these heuristics, the role of the state changes from seeking to alter the balance of incentives to one that responds to the defaults, priors and norms that people have. So now we are in the realm of nudge – the term that Thaler and Sunstein popularised in their book of that name. The idea is that the state can take advantages of biases in people’s behaviour to help them get to where they want to go without trying too hard, which partly addresses the problem of coercion and the acceptability of behavioural choices I talked about earlier, though it does not quite address the problem of paternalism. As the economist Bob Sudgen has pointed out, nudge does involve the state saying it knows best and designing

policies accordingly. I agree with this argument, and I think that nudge interventions run alongside and may even enhance more traditional authoritative tools of government. But the difference is in the approach of the state, which is more responsive and citizen-focused. I am going to come back to this problem of nudge in a minute, but I now want to talk about experimentation which heads the title of this lecture.

The big problem with behaviour change is that citizens are locked in a variety of different kinds of behaviours that come from their neighbourhood, work place, friends, and family, which are in turn influenced by what government does, such as regulating, providing information, incentivising, and reminding. As social scientists we know a lot about why people behave the way they do. In voter turnout, there is half a library of studies which show the influence of region, social class, family,

gender and so on. We know much less about how to change individual behaviour – the standard observational techniques in social science do not help us understand what interventions work in achieving behaviour change. Most of what government does as a matter of course has been introduced without much knowledge about how citizens react to such rules, so that we get reminders to pay our taxes, notices saying do not litter, but no one knows – at least until recently – what is the impact of these messages. So citizens are constrained by their own environments and in way that might be reinforced by government routines and interventions. So how do we build up the evidence base?

By experiments – as used in the social sciences – I am talking about a research design whereby any influence on the outcome of interest that is not the intervention has been randomised

between an intervention/treatment and control group, where the latter do not get the intervention.

SLIDE

In the world this can happen by chance, such as by a lottery, but more often, we need to create these conditions artificially in what is called the randomised controlled trial, which is where the researcher and policy-maker carry out the randomisation, creating control and treatment groups, where the random allocation of enough subjects – whether they are individuals, households, streets, communities – allows the determination of the effect of an intervention. This is very powerful because other methods find it very hard to separate out the various determinants of an outcome: the worry is that there may be something unmeasured that may be correlated with it. The RCT sets up the counterfactual of what would have happened

without the intervention.

It is possible to find the use of experiments in political science back to 1926 in the experiment by Harold F. Gosnell, ‘An Experiment in the Stimulation of Voting’, published in the *American Political Science Review*. Gosnell did what was what these days we call a Get Out the Vote experiment on 6,000 citizens divided into treatment and control groups, where the treatment group got a mailshot encouraging them to register to vote. Gosnell compared registration after the intervention: he found after several reminders 75 per cent in the treatment group registered compared to 65 per cent who did not which strikes me as an impressive treatment effect. But the recent interest and expansion came with the work of Don Green, Alan Gerber and colleagues at University of Yale from the late 1990s who did a whole series of voter turnout experiments showing

the impact of different kinds of intervention. Now there is a massive interest in experiments to test a variety of hypotheses which cannot be fully tested by other methods, such as whether politicians can be influenced by citizen observation, the impact of lobby groups, whether pay for performance works in delivering social services, the impact of media coverage and so on.

Those who operate in the world of policy have also shown an interest in experiments for much the same reason: an experiment allows policy-makers to ask what would happen if they did X and not Y so they are able to pilot what they do before doing it and evaluate it when they carry out the policy, in a way that generates a valid answer. Most other methods of evaluation suffer from problems of selective interventions on those more willing to implement the policy or on comparing outcomes before and after an intervention when they do not

know whether a comparison area would have fared just as well. In fact, the surprise is that policy-makers have not done more experiments in the past as their benefits have been known for a long time. There have been some great examples of experiments done by government going back to the 1940s and of course the National Health Service commissions and utilises many experiments. But it was only since 1997 that the UK government started to take a strong interest with experimental designs reported as the gold standard, such in the Magenta Book published in 2003. But even here it is only since 2010 that the interest has really taken off – here to address the core thing I am talking about in this lecture: how to introduce more effective behaviour change initiatives. This has been pioneered by the work of the Cabinet Office’s Behavioural Insights Team, many of whose members are here this evening, as they have

sought to show evidence for behavioural interventions. The team has carried out a number of experiments on reminders to pay taxes, court fines, energy savings, health and charitable giving. I think it is fair to say that the team has prompted a wave of experimentation across and beyond government, and to the rest of the world.

You no doubt have observed that I have used the term experimentation rather than just experiments. The word experimentation captures the idea that field experiments can become more a way of life for both academic and policy-maker partly, encouraging curiosity and a willingness to take calculated risks to find out new knowledge. This is because both the academic and policy-maker are forced to ask precise questions about the social and political world in ways that reflect an active approach to it. For the academic experiments can test theory (such as

different kinds of mobilisation), but they also help us to engage with the real world by crafting interventions that operate say in a local community. Experiments encourage policy makers to design an instrument that might work rather than try to operate in a blunderbuss way with too many interventions at once. They encourage policymakers to stand back and learn from what they do.

The other interesting thing about experiments is that they need considerable practical skills to carry them out, often in co-operation with many partners. In these partnerships is it possible to observe that rather than experiments being rolled out top-down in fact they involve participation and collaboration, with input from the partners and even the citizens and workers themselves. With policies on behaviour change, my argument is that the practice of experimentation is able to confront the

top down and technocratic approach that is often implied in behavioural interventions — and this can help establish a more collaborative form of governance.

Even with these large advantages, it would be wrong of me not to point out some limitations to the use of experiments.

First, not everything can be randomised and it might seem to privilege more individuated interventions based on large numbers with measurable discrete contacts. But I do think creativity can overcome this and the range of subjects for experiments is much larger than one would think.

The second is the need for a large number of trials to increase external validity, because in the medical world there is a need many hundreds and a much more regimented approach, and we in the social sciences might have to accept that something like a NICE body might be needed – which really would be top-down.

It may be the case that there are phases with experimentation: a creative phase then a more systematic one to tie things down more robustly once the early findings have been rolled out.

The third is that experiments sometimes need complex estimation methods, which bring experiments closer to the regression world, so make them less distinctive than sometimes we think.

The fourth is – once proved by experiments – how long do these effects last afterwards? Are the strong effects derived from their newness, a kind of Hawthorne effect, that falls off in time? How can we ensure the effects last for the long term?

My take on these features is that we are just getting used to experiments in public policy and political science, and in other disciplines, which have experiments at their core, work with these limitations and seek to overcome them. In fact, the

examples I am going to talk about show this learning process.

My interest in experiments started in 2000 partly by accident or least by the intervention of a certain David Halpern. I knew David when I worked at the Policy Studies Institute – we used to jog in Regent’s Park at lunchtimes and I remember how David talked about how he wanted to work for the prime minister – little did I know then that he would be working for two prime ministers, both Tony Blair and David Cameron. The project tested whether the provision of civic education would improve levels of political knowledge, social capital and political attitudes (we found that it did).

SLIDE

By the way anybody who thinks that teenagers are not interested in politics and policy should see the results of this project. The experiment happened in the middle of our two

surveys – we were measuring the attitudes of 15-16 year olds in two waves. While we were in the middle of the project David came up with the idea of having a debate about politics among half the schools using IT software.

SLIDE

Each school elected two representatives who summed up all the issues and then we put them on a coach to London where they met their MPs and were shown around the House of Commons then we went over to 10 Downing Street to have a debate with senior policy makers. We did get the date in Tony Blair's diary and we told the students.

SLIDE

So we may have raised expectations and we hoped they were not too dashed when the group met – in the dining room of No 10 – the prime minister's education advisor at the time,

Andrew Adonis.

SLIDE

When we checked back we found the experiment only partly worked except for civic and political knowledge. SLIDE

This is something one needs to get used to with experiments. But I had already started to notice certain things about experiments: the excitement of the students, the role of partners, and of course the difficulties of implementation.

My second experiment came when I was at the University of Manchester. I had just started, and as part of my chair appointment, The Hallsworth Chair of Governance, the university had given me £10,000 for research expenses. With this money I felt here was enacting the parable of the talents as rather than going on long expensive research trips round the world, I invested it in trying out one of the Gerber and Green experiments for

the 2005 General Election, which had never been done in the UK before. What we did was to test whether voters would be mobilised by a doorstep visit or a telephone. Green and Gerber had hypothesised that the stronger the personal contact the stronger the treatment effect – that is the more the person gets a personal touch and interest, the more responsive they will be. They posit there is a hierarchy of treatment effects.

SLIDE

We selected, 2,300 for each treatment group, and one group got nothing. We showed effects of the interventions by checking with the electoral registers afterwards.

TWO SLIDES

The interesting finding is that the hierarchy did not work in our case: less personal forms of contact were just as important, which shows that we need not rely on face-to-face contacts

for mobilisation and participation — contacts we might make through phones and by implication the internet might be as effective.

As ever with these experiments, it is the other things ones discovers when doing them that is as important as the findings: one is the collaboration involved in getting it off the ground, and the enthusiasm during the project, which I discovered when I took command in my office HQ, the cafe in Wythenshawe Sports centre. The other thing I noticed, a bit like the socialisation experiment, was that the residents of Wythenshawe wanted to talk about politics of all things. Even in this white working class rough estate, the residents treated our students with courtesy and interest — and complained that the political parties did not bother to talk to them any more. So I concluded that there is an appetite for politics and debate, and this links

to behaviour change.

With this Wythenshawe project complete, I got completely hooked on experiments, and I want to talk more about others I have done later. But I would like to pause a bit and tell you why I got interested in them. In this kind of occasion one is encouraged to look back at my career and to examine what kind of research one likes to do. I think that the kind of research I like reflects my pleasure in finding new facts with good methods. I like the science in political science: and this is a bit in contrast with what I like to do for leisure which is rather arty. So I can think of nothing more pleasurable than to go hear Wagner's Ring Cycle, usually with my fellow Wagnerian-obsessive, Chris Brown. But when I get up the next day, when the dragon, giants, dwarfs, gods and heroes all go back in their boxes – I want to do science. And in thinking about my career, what

is striking is the number of empirical firsts I have managed, usually – and this is an important point – in collaboration with other people who have influenced me very much, not just in the projects I have done with them, but more generally - and this is my chance to thank them – for they are my friends too.

So with Keith Dowding, I was the first in the UK to find out using survey evidence that people move to local authority areas on the basis of local taxes and services, a test of the Tiebout hypothesis; with Hugh Ward, I was the first to find out in the UK that the government targets resources to marginal Westminster seats and to the seats of ministers, the pork barrel hypothesis; with Helen Margetts we were the first to show that the UK policy agenda was punctuated; with Oliver James, we were the first to find that citizens vote on the basis of measured performance scores; with Tony Bertelli we were the first

to show that these Audit Commission scores were targeted to marginal seats and where the Audit Commission rewarded the authorities of those councillors who sit on the commission with better performance scores; and with Tony I am now applying an investment model of policy making to show how a national government allocates a portfolio of policies to get re-elected.

With these firsts, it is no surprise that I should be enthusiastic about experiments which allow me to make further advances. To have got to this point, I owe a lot of my collaborators, but I also owe a lot of my early teachers and mentors, who indulged my enthusiasms and helped me at early stages in my career. I thank Elizabeth Meehan, the late Jim Sharpe, the late Peter Willmott, Joni Lovenduski and finally Gerry Stoker, who is here today to do the appreciation. I think all these people share that essential quality of a teacher or colleague:

showing a real interest in the person as well as the intellectual enterprise itself. And ironically that interest in the person is what is influencing my experimental designs.

My third experiment or rather set of experiments are about civic participation and behaviour change, where again I have benefited from talented co-researchers as you can see in this slide.

SLIDE

This is the project that got published as the book, *Nudge, Nudge, Think, Think*, published by Bloomsbury under the guidance of Caroline Wintersgill who has published so many of my books.

SLIDE

Here we expand the outcome of interest to a range of policy outcomes, and we move beyond the nudge perspective by sug-

gesting the citizens need to ‘think’ and address the wider issues, that effective behaviour change should rest on a wider reflection rather than government finding various levels to switch on and off citizens behaviours in an automatic sense. We also tested a range of what I call influencing experiments, which take forward Green and Gerber’s work, and are based on the idea that contacting matters, which can vary by mode but not – they claim – by the type of message.

SLIDE

And then what I call second generation experiments involving feedback, which is where through observation and reference to a wider social group we can achieve higher effect sizes. When the individual receives a message that indicates what the wider group is doing, we get more mobilisation, sometimes doubling the effects.

SLIDE

We did several experiments - which showed these different kinds of interventions.

SLIDE

The one that I want to focus on is the one for charitable donations. Here we – and I am particularly indebted to Sarah Cotterill who made this experiment happen.

SLIDE

To recruit children's books for Africa, we randomised groups into a letter, pledge and what we call pledge plus publicity. Here what is interesting is how light-touch is the intervention using reference to others, which acts as a form or recognition for charitable acts.

SLIDE

We found that letting people know that their names would

be posted up in local places, such as libraries, counts. The pledge did not work. This shows that the intervener needs to work with the grain of the individual, to find out what switches someone on to participate. Again, as with all our experiments, we operated in a decentralised and collaborative way, and I think this was the secret of our success.

Most of the nudge experiments worked, but we found the think experiments harder to operationalise – they were more difficult to use for behavioural policy. My response to this is not to say that government should just stick to carefully controlled interventions, but it may be the case that more think-based interventions could be closer to nudge type in a way that they are dynamic and easy for the citizens to do – and in this way I am interested in the process of feedback, especially online, which is influencing my current work with Helen Margetts, and is be-

ing published our book which is being prepared for Princeton University Press.

I have also carried out more voter turnout experiments with Ed Fieldhouse and Dave Cutts, where we show the importance of impersonal methods of canvassing, such as leafleting and telephone calls.

I have played a small role with the Behavioural Insights Team itself, where I am very proud to be one of the members of their academic advisory panel. I am very grateful to the team members David Halpern, Owain Service and Laura Haynes, Elizabeth McDonald, Michael Sanders and Sam Ngano, for allowing me to help design some of their trials and to analyse the data. What is interesting in terms of this lecture is how the team's work embodies the core theme of this lecture: they have found a way of weaving together experimentation, behaviour change

and public policy in the heart of their work, by embracing RCTs, but in a way I think consistent with what we found in *Nudge, Nudge, Think, Think*. I do not have a great deal of time to talk about the experiments I have been involved with the team, such as showing picture of a car to those who are behind paying for their driving licences, and about using mobile texts to remind people to pay their court fines. Both these interventions show the bottom up nature of experimentation, such as the texting idea which came from the courts service itself and is about using information it already has to good effect, such as the mobile phone numbers of those who are behind with their fines.

What the BIT work shows is the large wave of experiments that are coming about. In this context, both academics and governments are at the start of something that is going to get

very big, and I am excited to be in at the start of it. I am also very happy to be at UCL which I think is one of the best places to be able to do experiments and to take advantage of this interest: and it is a happy coincidence that this lecture happens in – and forms part – of UCL’s Behaviour Change Month. In this context, it is a great opportunity to be able to link to those in the medical faculty, engineering, economics, and psychology as well as colleagues and research students in political science. The great thing about an inaugural lecture is that it is both a celebration of my work, but also a starting point for a new wave of research, such as my new projects on volunteering and on crowdsourcing. So my final point to you this evening is to say that I think I have only just began on this journey. And I am very much looking forward to hatching lots more experiments here at UCL in the years to come. Thank

you.