

Experimental innovation and growth policy: Why do we need it?

Albert Bravo-Biosca, May 2016 (Version 1.0)¹

Despite the importance of innovation and high-growth entrepreneurship for economic growth, there are still many open questions on the best approaches to support them. There is a need to experiment with new approaches and, crucially, learn what works.

European governments alone spend around €150 billion every year trying to make their economies more innovative and entrepreneurial, as do many other governments around the world.² Yet the programmes that governments put in place to increase innovation, support high-growth entrepreneurship and accelerate business growth suffer from two weaknesses:

1. There is insufficient innovation in innovation and growth policy:

Despite a fast changing world in which new technologies and business models continuously emerge, the tools that governments use to support them have evolved little. New programmes are regularly launched, but changes are often purely cosmetic (e.g., a new name, or a few tweaks at the margin). Also, supposedly new policy approaches are sometimes the reinstatement of old ideas that had fallen out of fashion (the 300-year old idea of challenge prizes being the latest example). And even when there is genuine innovation, be it incremental or radical, it is not of much use if we cannot tell whether it's better or worse than the programmes that preceded it (since in contrast to the private sector, there are no markets to guide us).

2. There is limited evidence on the effectiveness of innovation and growth policy:

The uncomfortable truth is that surprisingly little policy in this domain is backed by hard evidence. A few years ago Nesta supported the Compendium of Evidence on the Effectiveness of Innovation Policy, led by researchers at Manchester University.³ With some exceptions, it showed that the evidence was scarce, often of poor quality, and typically inconclusive. More recently, a systematic review by the What Works Centre for Local Economic Growth examined almost 15,000 evaluations in this policy domain, and found that only 2.5 per cent of them provided a credible answer, and of those, only one in four found a positive effect on productivity or employment (or 0.6 per cent of the total).⁴ Without good evidence, it is impossible to allocate our limited resources to the programmes that have the greatest impact.

1. In the spirit of continuous iteration and improvement, we see this brief as a living document that will evolve over time as new examples, arguments and debates emerge in the field.

2. Firpo and Beevers (2016) 'As much as €152 billion is spent across Europe supporting businesses: but does it work?' Available at www.innovationgrowthlab.org/blog/much-%E2%82%AC152-billion-spent-across-europe-supporting-businesses-does-it-work.

3. See www.innovation-policy.net.

4. See www.whatworksgrowth.org. Credible refers to evaluations that satisfy the level 3 of the Scientific Maryland Scale, which requires that the evaluation method used has a credible counterfactual (note that random allocation is not a requirement for level 3, it is sufficient to have a clear justification on why the control group would have performed in a similar way as those benefiting from the intervention if the intervention had not happened).

There are surely more effective ways to use the large budgets that governments around the world devote every year to supporting innovation, entrepreneurship and growth. We just need to be willing to try new programmes and really be open to finding out whether they are working. Otherwise large amounts of money may continue being wasted on ineffective schemes. Even more worryingly, ideas that could lead to the big innovations of tomorrow may never be developed if we lack the right instruments to nurture them.

The case for an experimental innovation and growth policy

Governments typically introduce large new programmes without prior small-scale testing, and very limited knowledge of whether they will be effective. Instead of recognising that they don't know what impact the programme will have, they typically assume that it will work. In addition, there are always many possible design choices for a programme, yet despite not knowing what the best design is in order to achieve the greatest impact, a ministerial announcement follows shortly and millions are poured into it.

The alternative is to set up small pilots to experiment with new programmes, evaluate them using rigorous methods, and scale up those that work. The learning should not end when the pilot ends, it is important to continue evaluating and experimenting, since the fact that a programme is proven to work well doesn't mean that it couldn't work better, or that it will continue to work when implemented at a larger scale. Ultimately, this experimental approach is a smarter, cheaper and more effective approach to develop better innovation and growth policy instruments.⁵

In order to become experimental it is not sufficient to pilot new programmes. Trying new things is important, but real experimentation requires putting in place the systems to learn whether they are working or not. Not having an evaluation strategy, as it often happens, defeats the purpose of using pilots.

Obviously, experimentation is only one of the ingredients for delivering good innovation and growth policy. Better use of data would also help to develop more effective policies,⁶ and good judgement will always be required, since in an uncertain world where information is incomplete, the evidence base can only take us so far. But it could take us much further if governments made it a priority.

It is surprising how little R&D is done by governments themselves to make sure that their innovation support programmes are having the greatest effect. In other words, not only R&D that advances our scientific knowledge, but also R&D aimed at developing better tools to transform this scientific knowledge into innovation that increases economic growth and helps to address our societal challenges.

Governments routinely ask businesses to increase their R&D investment in order to improve their products and boost their competitiveness, yet they fail to do so with their own activities. If European governments applied to their own programmes the 3 per cent national R&D target that they have all committed to, they should be investing €4.5 billion every year to develop more effective tools to support innovation and growth. While we don't know how much governments spend on R&D focused on making their programmes more impactful, it is certainly only a very tiny fraction of that.⁷ The result is that innovation policy is not very innovative, and we don't know what works and what doesn't. But it doesn't need to be this way.

5. See for instance Breckon, J. (2015) 'Better public services through experimental government.' London: The Alliance for Useful Evidence.

6. See for instance Nesta (2016) 'Innovation Analytics: A guide to new data and measurement in innovation policy.' London: Nesta.

7. In fact, knowing how much governments spend overall on innovation and support programmes is already a very difficult question to answer, given the limitations of the data available. See for instance this blog by Teo Firpo on the challenges of getting data in this space: <http://www.innovationgrowthlab.org/blog/business-support-research-methodology>.

Learning from policy experiments

There are several methods to generate good evidence, and the choice of the method depends on the characteristics of the programme and the circumstances under which it's implemented. But among the different methods available in the evaluation toolkit, there is one that has been particularly underutilised in the domain of innovation and growth policy, namely randomised controlled trials (RCTs), and the quality of the evidence has suffered as a result.

One distinction between high- and low-quality evaluations is who they manage to convince. Low-quality evaluations are only convincing to those that are already convinced about the virtues of the programme, but fail to convince those who think otherwise. In contrast, a high-quality evaluation is robust enough to change people's views on a particular programme, and it is therefore more likely to impact the choices that are made (and ultimately lead to better decisions).

Unfortunately, typical evaluations of innovation, entrepreneurship and business growth programmes only give a good answer to the question 'how well did the programme participants perform before and after the intervention?' They commonly fail to provide a compelling answer to the more important question: 'what additional value did the programme generate?' Or in other words, is the improved performance of firms receiving the intervention the result of the programme itself, or does it reflect some unobserved characteristics of the firms that choose (or were selected) to participate in the programme? Answering this question requires good knowledge of how participants would have performed in absence of the programme, which is difficult to find out unless you have a credible control group.

The idea behind randomised trials is a simple one. Participants are randomly placed in a 'treatment' group and a 'control' group, and the impact of the programme is estimated comparing the behaviour and outcomes of the two. The lottery used to assign participants to each group addresses potential selection biases, so that both groups are comparable and any differences between them are the result of the intervention. Therefore, they provide an accurate estimate of the impact of the programme, and this is the reason why they are often referred to as the 'gold standard' for evaluation.

Randomised trials have been used extensively in health to test the effectiveness of new pharmaceutical drugs as well as medical procedures. But they have also been widely adopted in several other policy areas, such as development, education or social policy. For instance, the MIT Poverty Action Lab (J-PAL) has run over 700 randomised trials of poverty-reduction interventions, radically transforming the development field in the process. The UK-based Education Endowment Foundation is conducting over 100 randomised trials involving more than a 1,000 schools to test different ways to improve educational outcomes.⁸ And the French government runs an experimentation fund for young people,⁹ a bottom-up approach to identify innovative interventions to improve youth outcomes (crowdsourced from organisations across the country), implement them at a small scale, and rigorously evaluate them to find out whether they work, before deciding whether they should be scaled up.

In contrast, the use of randomised trials to test innovation, entrepreneurship and business growth programmes has been very limited, particularly in advanced economies. There are several reasons for this, ranging from inertia to some misconceptions about trials (see box on page 6). While running experimental pilots in this policy area may sometimes involve additional challenges, there are already several examples that demonstrate both the feasibility and benefits of using them.¹⁰

8. See more at www.povertyactionlab.org and www.educationendowmentfoundation.org.uk.

9. Fonds d'expérimentation pour la jeunesse (www.experimentation.jeunes.gouv.fr).

10. See the Trials Database maintained by IGL (www.innovationgrowthlab.org/igl-database-map).

Becoming more experimental: what does it mean in practice?

A thought often pops up the first time that someone is presented with the idea of using randomised trials in innovation and growth policy. Namely that this involves randomising ten million euro R&D grants, including having to give some of them to unworthy applicants.

While some programmes can easily be randomised, for others it may not be feasible due to political or practical reasons. But even when it is not possible to do a randomised trial to evaluate the full impact of the whole programme, it is often possible (and useful) to use randomised trials to improve the effectiveness of the programme. For instance, you cannot typically randomise R&D tax credits, but you may still be able to use a randomised trial to test the impact that providing personalised advice on how to fill the tax credit form has on firms' likelihood to apply for an R&D tax credit, the amount they claim for, and the validity of their claims. Similarly, it is not possible to do a randomised trial on a system-wide intervention that seeks to create or support a new industry, but the actual instruments that the intervention uses may well be amenable to randomised trials.

Randomised trials can be used at different stages of the policy development cycle. Therefore, it is helpful to consider three broad ways of using them that policymakers can adopt in order to improve the effectiveness of innovation, entrepreneurship and business growth support programmes:

1. Full-fledged randomised trial to evaluate the overall impact of a new programme:

A standard experimental evaluation in which a new programme is piloted at a small scale, with a treatment and a control group, and the impact of the programme on the outcomes of interest is measured in order to find out whether the programme works and should be scaled-up. For instance, at IGL we have been working with the Danish Government to evaluate their new pilot programme to support high-growth businesses, and see the impact that it has on firm performance.

2. Testing tweaks to existing programmes:

Most existing programmes face some challenges that limit their impact, and policy makers and programme managers involved in their delivery typically have ideas on how to make them more impactful. Some of these are easy to test, such as:

- a.** Increasing the take-up of the programme by using A/B trials to test the effectiveness of different marketing materials. This may be as simple as experimenting with different newsletter formats and the specific language used on them, to trying out different marketing channels.
- b.** Trialling changes in the programme design and delivery, instead of rolling them out to all participants simultaneously. This may include tweaks in the user journey, changes in the delivery mode (online vs. face-to-face), or adjustments to the internal processes. For instance, a trial funded by IGL is looking at how changing the intensity of mentoring that entrepreneurs at a tech incubator receive impacts their future performance.
- c.** Incorporating add-on programmes on top of an existing scheme. For instance, several governments are contemplating the option of enhancing their R&D and innovation grants schemes by offering management coaching to recipients, with the aim of increasing successful commercialisation. But to our knowledge the impact of this add-on programme is not being tested first on a subset of recipients, despite being perfectly feasible.

3. Experimental development of new programmes:

Designing a new programme involves making choices with regards to many different design features. While it is not feasible to test each and every option against each other, and therefore good judgement plays a role, it is often possible to test several of them at a time. This rapid experimentation does not necessarily test the effectiveness of each option on the ultimate outcomes of interest, since these may be years away. Instead, it starts from the theory of change of the programme, and identifies early indicators that are expected to correlate with the ultimate outcomes of interest, but which can be observed relatively quickly. This enables rapid learning and continuous iteration at an early stage in the development process, leading to more effective programmes.

This list is only a simplified menu of options for organisations interested in becoming more experimental with their own programmes. There is actually a continuum of options that spans across this range, and that extends even further if organisations want to consider not only their own portfolio of programmes but also the wider mission that they aim to achieve. For example, another way to become more experimental is recognising that there are good ideas elsewhere, and that you need mechanisms to identify, test and support them (such as experimentation funds like the Education Endowment Foundation or the French Experimentation Fund for Youth discussed above).

Ultimately, the key ingredient to becoming more experimental is to change the mindset within the organisation. Being open to new ideas, regardless of where they may come from. Not being afraid to try them out, experimenting with them in different shapes and forms. Recognising that failure is an option, but being aware that useful lessons will be learnt and money will be saved as a result. Being challenging when necessary, particularly when confronted with attempts to pour millions into untested programmes that could have been tested at small scale. And last but not least, never forgetting that even when programmes work, they can always be made to work better, so embracing continuous experimentation as the norm rather than the exception.

Achieving this culture change is not a simple feat. Sometimes this can be led from the top, with a clear commitment to experimentation and some flagship randomised trials to prove it. But even when this is not the case, it is still possible to advance on this journey, starting small with low-risk trials that help to demonstrate to the wider organisation their usefulness, and hence contribute to make the case to take it to the next level.

Concluding remarks

The idea of experimental innovation and growth policy is a simple one. Look for new ideas but don't assume that they will work. Instead, whenever you can, test them at small scale, put in place the systems to learn whether they actually work, and based on this decide whether to scale them up or not.

Randomised trials are a useful tool to learn what works and what doesn't (and when), and a very much underutilised one in the domain of innovation and growth policy (in contrast to other policy areas). Wider adoption would lead to more effective innovation, entrepreneurship and business growth support programmes, and help to ensure that our limited resources are focused on the programmes with the greatest impact.

This is why we launched the Innovation Growth Lab (IGL), a global collaboration of governments, foundations and researchers interested in making innovation and growth policy more experimental and evidence-based (see page 8 to find out more). Experimentation is a simple idea, but putting it in practice in the messy world of policymaking is not always easy. So if you would like to explore how your organisation can become more experimental, please contact us at innovationgrowthlab@nesta.org.uk We will be happy to help.

Some misconceptions about randomised trials

There are several misconceptions about randomised trials that lead many to conclude that it's not possible to conduct them in the domain of innovation and growth policy. While some of these concerns are valid and require careful consideration, there are sufficient examples that demonstrate that it is both possible and valuable to do them. See for instance the Trials Database maintained by IGL or the projects funded by the IGL Grants programme.

1. They are unethical:

An implicit assumption behind this criticism is that trials involve denying some potential recipients an intervention that would benefit rather than harm them. However, this cannot be taken for granted (for instance, a trial examining an entrepreneurship support programme in the US found out that the quality of the training was so weak that, rather than helping firms, if anything the impact was the opposite).¹¹ Moreover, rolling out programmes without knowing whether they are beneficial or harmful is a risk worth preventing. Finally, using scarce resources on programmes that don't work deprives more effective programmes of funding. Balancing these and other similar considerations is the reason why trials that raise ethical issues require the approval of an ethics committee before they can proceed, to make sure that the benefits outweigh the risks. In some cases an approach to circumvent some of these issues is to use A/B trials, in which instead of having one treatment and one control group, there are two treatment groups allowing you to compare two different versions of the intervention.

2. They are politically impossible:

The use of randomised trials is widely accepted by the general public in much more sensitive policy areas, like health and education. The UK has been an early adopter of trials in innovation and growth policy, and its experience demonstrates that, while they require a carefully planned communication strategy, it's perfectly possible to undertake them without sparking controversy. Moreover, there are often insufficient resources to support all potential recipients in any case, and a lottery can be a fairer (and cheaper) approach to determine recipients than some panel-based scoring approaches. Note that this doesn't mean supporting unworthy recipients, but rather taking advantage of excess demand and therefore the oversupply of good proposals to randomise among them (after having excluded the proposals that would not qualify). For instance, this is the approach that an IGL-funded trial is using to estimate the impact of getting into a tech incubator on startups performance, taking advantage of the fact that the number of entrepreneurs interested in joining is much larger than the available space.

11. Fairlie, R.W., Karlan, D., and Zinman, J. (2015) Behind the GATE Experiment: Evidence on Effects of and Rationales for Subsidized Entrepreneurship Training. 'American Economic Journal: Economic Policy.' Vol. 7(2), pages 125-61.

3. They don't tell you why there is an effect:

While this is often the case in classic trials set up to measure the impact of a particular intervention, it doesn't need to be this way. Randomised trials can also be designed to help understand the behaviour of individuals/firms and what drives it (these are often called mechanism experiments, and are generally derived from economic theory). Measuring the impact of an intervention and understanding what causes the underlying behaviour don't need to be mutually exclusive aims. In addition, as a general rule it is considered good practice to use mixed methods when undertaking randomised trials, and therefore complement the quantitative analysis with qualitative approaches such as case-studies, which can provide rich insights about the intervention and the participants.

4. They are not useful:

Several reasons are often given to make the case that the lessons learnt don't justify the effort of undertaking a randomised trial. Two of the most common ones are that results are context-specific and that innovation and growth policy is too complex and has too many dimensions to be able to test it neatly with a randomised trial. On the first, it is fair to say that context matters, as in any other type of evaluation (and everyone agrees that this is not a reason not to evaluate). While some insights may be generalisable, multiple trials and evaluations in different contexts are always desirable. On the issue of complexity, trials can answer some questions but not all questions. Knowing what questions to ask, what design features to test, requires good judgement. But if the right questions are identified, trials can provide a useful answer.

5. They are very expensive:

While randomised trials can cost anywhere from a few euros to tens of millions, the additional cost of using a randomised approach compared to another evaluation approach is actually very little. What is often very expensive is the implementation of the programme that is being evaluated, and the collection of the data needed for the evaluation.¹² But these are costs that need to be borne regardless of the evaluation method used. In fact randomised trials require smaller samples, which implies cheaper data collection, and the analysis work is often more straightforward too (even if the initial design requires more work). The main barrier to the adoption of randomised trials is therefore not a budgetary one, but rather one of political will in the organisation. If anything, limited budgets reinforce the case for experimentation, since it is more likely that there is excess demand for the organisation's programmes, which makes the option of randomising easier, and in addition with scarce resources it is even more important to make sure that they are used in a way that achieves the greatest impact.

12. In addition, the use of administrative data and the increasing availability of new sources of data (e.g., online data) can cut the cost of data collection significantly.

About the Innovation Growth Lab - IGL

IGL is a global laboratory for innovation and growth policy, bringing together governments, foundations and researchers to develop and test new approaches to increase innovation, support high-growth entrepreneurship, and accelerate business growth.

Our aim is to make innovation and growth policy more experimental and evidence-based. To deliver on this mission we undertake a range of activities that tackle the different barriers to the wider adoption of experimental approaches. Some of them are open to the wider community while others are limited to our partners, who ultimately guide our work.

Our work with partners

We work with our partners in different ways to help them become more experimental and contribute to increase the effectiveness of their programmes. We do this through:

- Analysing their portfolio of programmes, both on-going and forthcoming, to understand the challenges that they face in each of them, and proposing testable solutions to address them.
- Identifying opportunities for new randomised trials that they can undertake, and supporting them to scope, design and conduct them.
- Carrying out randomised trials jointly with them, with the option of matching them to researchers in the IGL Research Network.
- Building capacity in their organisations through a variety of activities, ranging from on-site capability building workshops to webinars.
- Developing a global community that facilitates exchanging experiences and peer learning from other partners.

Our wider work

We also try to push forward the experimentation agenda globally through a variety of activities:

- Funding randomised trials through the IGL Grants programme, in order to encourage more organisations and researchers to get into this space.
- Bringing together researchers in this field through the IGL Research Network.
- Connecting researchers and organisations interested in collaborating on trials through our matching and brokerage role.
- Creating useful resources that contribute to capability building, such as our new online guide on how to do randomised trials in innovation and growth policy (in beta).
- Disseminating the lessons emerging from trials worldwide through presentations, publications and our online Trials Database (which includes detailed information on trials in this field).

If you would like to find out more about the work of IGL, or are interested in becoming an IGL Partner, please contact us at innovationgrowthlab@nesta.org.uk

Nesta... IGL is an initiative led by Nesta in partnership with:

The views expressed in this publication are those of the author and do not necessarily represent those of the IGL Partners.