

Is there madness in our method?

History of Economic Thought Society: 33rd Conference, 22 September 2022

Michael Brennan, Chair, Productivity Commission

The emotion of art is impersonal. And the poet cannot reach this impersonality without surrendering himself wholly to the work to be done. And he is not likely to know what is to be done unless he lives in what is not merely the present, but the present moment of the past; unless he is conscious, not of what is dead, but of what is already living. (TS Eliot)

Why should we bother with the history of economic thought?

It's not an idle question. If we regard economics as a pure 'objective' science, with consistent progress and each discovery building on the last, then we shouldn't really care *how* we came to know what we know – it would be enough just to know it.

Do we read dead economists because all their theories are equally valid for their time and place? Do we think of economics as like an evolving literary canon, with each new work interpreted in light of the old and the old re-interpreted in light of the new as TS Eliot suggested?

Maybe, but this can't be the whole story. After all, it's possible to be wrong in economics. A purely relativist account doesn't cut it.

Economics isn't physics, but it does progress over time. There are genuine discoveries and breakthroughs, but they tend to come in an uneven, non-linear and context-specific way. Popper contended that all scientific 'truths' are conditional; they seem particularly so in economics. The profession is more loosely bound around its shared corpus of knowledge and its understanding of technique than the 'hard' sciences of physics or chemistry.

But would any of us deny that economics has made advances?

Think of Smith's discussion of the division of labour; Ricardo's theory of land rent, later generalised by the marginalists; the marginal revolution itself, a remarkable case of simultaneous 'discovery' by economists as different as Carl Menger, Leon Walras and WS Jevons; the theory of public goods; game theory more generally. Don't all of these examples represent in some way a discovery or advance, an improvement in our understanding of economic phenomena in the real world?

But this progress is jagged. There can be two steps forward and one step back; we embrace the new but forget something of the old. We get carried along by a fresh intellectual current and we can fall for the fallacy that there is but one true way to 'do economics', only later realising what we have inadvertently left behind.

So economics, perhaps more than other sciences, has to be methodologically self-aware, and knowing your history helps.

In the face of any new wave, it pays to ask not just 'where will this get us?' but also 'where could this take us off course?'. What is it that we risk forgetting in all the excitement?

To do that you need some vague sense of true north. In what follows, I merely offer my own perspective. I am not a scholar, but a policy practitioner. Moreover, these views are just mine: I cannot talk on behalf of the Productivity Commission, which is full of smart people who (one would hope) would disagree with at least part of my account.

I would emphasise three things about economics, traits that, put together, are central to the discipline and which we would not want to lose.

First, it is about **people**. Economics (as per one treatise) is a study of *human* action. Not microbes or molecules; not billiard balls or falling objects; and not heroes or saints. Just people – who act purposively, with the usual mix of strengths and frailties, cognitive and moral.

Second, it is a **social science**. People don't just act – they *interact*. They tend to truck, barter and exchange, as Smith put it. One person's actions shape another's incentives and vice versa. Economics highlights the way individual actions aggregate into patterns, which in turn shape and re-shape the very context for individual action in the manner of a feedback loop. This is a trait shared by traditions as diverse as Walrasian general equilibrium, Hayek's *Use of Knowledge in Society* and Keynesian insights into macro adjustments in the face of uncoordinated expectations. Economics has an account of equilibrium and, importantly, of disequilibrium.

Third, it aspires to **generality**. Economics aims to find regularities; to better understand the world, not just interpret an individual event.

This is my own amateur typology – by no means perfect. My contention is that economics risks going off course when it loses sight of any or all of these things.

I will illustrate this by talking about two intellectual currents: the rise and dominance of neoclassical economics, with its tendency to mathematical formalism, and the ongoing empirical revolution, with its focus on experiments and quasi-experimental econometric techniques.

Both have been positive developments. I think the benefits of neoclassicism outweighed the costs, and that will probably be true of the empirical turn. But there *are* costs, and there are risks if we don't remain humble about the inevitable limitations of each intellectual wave.

The Neoclassical method

Neoclassical economics really grew out of the marginal revolution of the 1870s but arguably achieved its full dominance, including the mathematical formalism, after the Second World War. It had rigour and epistemic ambition.

In terms of my typology, it nailed the **generality** component: aiming for widely applicable insights (mainly theoretical) to the point of economics imperialism – uniting law, politics and the family under the economist's analytic gaze.

It was also conscious of the **social** component. The focus on equilibrium formalised the process I mentioned earlier: how price taking agents, through their combined and unintended actions, generate market outcomes which in turn set the prices shaping individual choice.

But it arguably lost sight of the **human**. I don't mean that neoclassical economics was uncaring. Quite the contrary: it brought welfare (i.e. well-being) considerations into the foreground. Once Menger and co. had 'discovered' marginal utility, it allowed the Paretos, Pigous and Samuelsons to explore the welfare implications of markets, market failures and corrective policies.

Where neoclassicism lost the human was in the specifics of the model itself. The people (or representative agents) in the model are generally assumed to solve a constrained optimisation – maximising utility subject to a budget constraint, or profits according to a production function and factor prices. They generally only change behaviour in the face of an externally generated 'shock'.

This has been central to the neoclassical approach. The models were neat, and powerful in generating predictions based on their assumptions. But one cost was that the agents within them ceased to resemble people, so much as automatons.

They responded to stimuli rather than actively making choices. Their 'actions' were not really actions at all, dictated as they were by the first order conditions – effectively pre-determined by the parameters and functional form of the model.

Agents (usually) had perfect information, so there was no discovery, no novelty, no innovation. In fact, in the competitive equilibrium there was no competition as such, not in the active, bustling sense we would recognise from the real economy. Attempts to add stochastic elements – say error terms with a specified distribution,

or a stock of knowledge based on new ideas that arrive according to a random process (of given parameters) – don't fundamentally change this reality.

Does this really matter? After all, the theory is just allegory, a stylised ideal against which to better understand the world as it is, including to isolate where the real economy departs from said ideal.

The answer is that it matters when the allegory starts to obscure as much as it reveals. The theory largely left out the frictions, but the frictions (and human response to them) turn out to be among the most interesting and important things going on in a modern economy. The theory also struggled with the brute economic fact – plain for all to see – that economic change, growth and progress are primarily driven from *within* the system rather than from exogenous shocks.

Pyka and Saviotti described this phenomenon in connection with neoclassical growth models:

There is no explicit recognition of changes in the ways economic activities are organised, for example the rise of the factory system, and later the modern corporation and mass marketing or the stock market and other modern financial institutions, the rise and decline of labour unions, and the continuing advance of science, or the changing roles of government as factors influencing the growth process ...

I reflect on Tim Harford's fascinating book, *The 50 Inventions that Shaped the Modern Economy*, in which he tells the story of innovations ranging from barbed wire to double entry book-keeping, and the people and struggles that brought them about.

I wonder what I would have made of it, as a neo-classically trained economics graduate in the mid-1990s. I would have found it fascinating, no doubt. I presume I would have found it economically relevant, but perhaps not the main game – not the central subject matter of economics.

If pressed, I might have noted that all these innovations were effectively exogenous forces: facts about the external world that an economist would take as given. A series of *external* forces that *impacted* on the economy, with implications for key economic variables like relative prices, factor returns and resource allocation.

Perhaps all these innovations could all be put into the bucket labelled 'A' – the exogenous technology parameter in a neoclassical growth model. Perhaps they could be endogenized via an equation linking knowledge to research effort.

But that doesn't really solve the problem. It might improve my intuitive understanding of the process somewhat, but the logic is still essentially circular. If I convinced myself that I had said something important about the drivers and process of technological advance, I would be very much mistaken.

The English author GK Chesterton (certainly no economist), drew an apt connection between the determinist and the madman:

As an explanation of the world, materialism has a sort of insane simplicity. It has just the quality of the madman's argument; we have at once the sense of it covering everything and the sense of it leaving everything out. Contemplate some able and sincere materialist, as, for instance, Mr. McCabe, and you will have exactly this unique sensation. He understands everything, and everything does not seem worth understanding. His cosmos may be complete in every rivet and cog-wheel, but still his cosmos is smaller than our world.

Arguably the weakness in the neoclassical frame stems in part from its strict determinism – its tendency to explain economic phenomena as the pre-determined and inevitable result of the model's core assumptions, occasionally shocked from the outside. It explains everything – or *something* very completely – and yet somehow leaves so much out.

A symptom of the mathematical formalism of the neoclassical approach is that it creates a strict duality between factors that are exogenous (unexplained and determined outside the model) and those which are endogenous (internal but pre-determined, at least within given parameters).

To draw this out, note the stark contrast with economic frameworks that sit *away from* the neoclassical tradition, such as Joseph Schumpeter's treatment of the role of the entrepreneur in economic progress. In part two of the *Theory of Economic Development*, Schumpeter starts to move the discussion on from the circular flow – his steady state straw man, which is ripe for disruption.

Here are two quotes, the first describing his process of economic development:

It is spontaneous and discontinuous change in the channels of the flow, disturbance of the equilibrium, which forever alters and displaces the equilibrium state previously existing.

Which sounds a lot like an exogenous shock. Until you read on to find:

By development therefore, we shall understand only such changes in economic life as are *not forced upon it from without but arise by its own initiative, from within*. [emphasis added].

So, Schumpeter's entrepreneur, the prime mover of economic development, is a disrupter who brings about discontinuity, who blows up the pre-existing order. And yet Schumpeter is adamant that this happens from *within* the system, within what we can loosely call his 'model'. Endogenous but undetermined.

I confess when I first read this, I thought it a contradiction. But that really just reflects a narrow view of what constitutes an economic 'model'.

Geoffrey Hodgson has pointed out that innovation relies on acts of individual creativity and choice, but as he said: "Genuine creativity, real choices and willed

changes of purpose mean that human action must contain an element of indeterminacy in the sense of an uncaused cause”.

This is not to say that modelling based on general equilibrium (and optimising agents) has no place. Used well, it can support intuition in ways that are not obvious to the naked eye. We have used it at the PC to better understand the potential impact of external shocks or large-scale policy change. We are using it as part of our current productivity review.

But it has limits, most notably in describing the spontaneous change that comes from within the economic system itself, which as it turns out, is where much of the action is.

Arguably economics has made its greatest strides when preserving a space between the exogenous and (my term) the *strictly endogenous* or model-determined: an in-between category of free action within the ‘model’ (broadly conceived).

That is, an analysis of genuinely ‘choosing’ agents, whose actions are not predetermined by parameter values and functional forms but who retain the basic tendency to purposive (or weakly rational) behaviour in their interactions with others (both competitive and cooperative).

This ‘in between’ is a region where agents can roam, discover, make and change plans, trade, make promises, renege, show faith, agree on rules or gradually shape norms. One where economists can still seek findings that are generalisable in a useful way, even with some quantification, if not mathematical precision. And ideally where economists retain the ability to say *something* about welfare implications.

I think of economists as diverse as Hayek, Coase and James Buchanan on the one hand, and Joan Robinson on the other; or Richard Nelson and Sidney Winter on growth; or (noting my earlier Schumpeter reference) Israel Kirzner’s description of entrepreneurship. I think also of Armen Alchian’s famous piece on uncertainty, evolution and economic theory.

All of these contributions (and they are a pretty random selection) share something of that vision of purposive, rather than parameter-driven, behaviour and change from within the system or model, neither exogenous nor pre-determined.

The point is that formalism eliminates that in-between. It shrinks the space. And we lose some important understanding as a result, whatever benefits we might also gain.

The Empirical Revolution

So what should we make of the next great intellectual current, namely the empirical turn in modern economics? It is 40 years since Ed Leamer talked about ‘taking the con out of econometrics’.

In 2021, Joshua Angrist, David Card and Guido Imbens won the Nobel Prize for their work on experimental and quasi experimental approaches in applied microeconomics. So a lot has changed.

Angrist describes it as the Credibility Revolution, driven by new data, better econometric technique, and a renewed focus on thoughtful research design.

Any observer of the profession can see the influence and impact of this movement, particularly on young graduates. The focus is squarely on causal inference, that is, identifying causation – usually after the event. Having observed effect B, can we say with confidence that it was caused by action/policy/event A? And how big was the effect *on average*, relative to a (real or synthetic) control group? That is, what is the average treatment effect (ATE)?

Where the stylised neoclassicist might assume a parameter which affects behaviour in the model, the modern empiricist starts with observing human behaviour to try and work out the sign and size of the parameter.

I note in passing, that both approaches imply a slightly mechanistic world view; a presumption that the central task of economics is to generate knowledge about underlying causal forces. A sort of billiard ball view of economics.

It’s a well-known trope that correlation is not causation. Finding causal effects in data – to discern the actual impact of a policy or event – is hard because the counterfactual is usually unobservable. Once a ‘treatment’ is administered, we cannot directly observe what would have happened without it.

But randomised trials and quasi-experimental methods (like instrumental variables, regression discontinuity and difference in difference approaches) can give us some insight. When we randomise (or find instances of naturally occurring randomisation), we can have more confidence that the counterfactual is proxied by the experience of the ‘control’ group.

Of course, in the real world, with its many potentially confounding variables, there is still a great deal of complexity in understanding the conditions under which this really holds true.

One of the reasons economists are well suited to this work is that our theoretical framework helps to locate the potential threats to validity – the instances of reverse causality, or omitted variables that might bias a particular specification.

So this approach seems a good fit for economists. It gives us stronger empirical results and it has arguably re-focused the profession on real world issues and policy relevance – not least in program evaluation. This all seems to the good. What's the downside?

Back to my typology, one issue is that of **generality**. Experimental and quasi-experimental research designs can be quite context specific. But policymakers seek transferable, generalisable results – a program that can be scaled up or an overseas policy that can be replicated.

The average treatment effect from a given sample does not always generalise to another context. Nor does it tell you what will happen in the case of a particular individual.

As Angus Deaton and Nancy Cartwright put it, “The literature discussing RCTs has paid more attention to obtaining results than to considering what can justifiably be done with them”.

The risk of over-selling the applicability of results is real, not least when economists get more focused on econometric technique than policy implications. From time to time in policy discourse, one hears the claim that we ‘know’ some general truth as a result of randomised trials or other empirical work. In fact, what we ‘know’ is probably more specific to time, place and sample.

Advocates will acknowledge this but say that the answer is more and better empirical work – more piecemeal results that will gradually build a knowledge base, brick by brick.

It is an open question what sort of knowledge base we will build up by that method over the long term. Plausibly, we will get better estimates of some structural parameters economists use for *other* modelling purposes, such as a price elasticity for a type of good.

And there are times when an average treatment effect tells us all that we want to know, as in a program evaluation. Did the program produce the desired result in general? If yes, we should continue.

Will we also gain transferable, generalised understanding about complex social policy issues in the real world via individual experiments and quasi experiments? I think the answer is a very tentative yes. The question is how broad and how firm that knowledge base will be. As always, economists will have to show humility about what we truly ‘know’.

But there is a further issue, which goes to the **social/interactive** element of my earlier typology.

Usually in economics, an empirical result about A leading to B is not an end in itself. It is just the start. We really want to know what the result means for subsequent behaviour and interaction, and welfare. The Nash equilibrium, say, and whether it looks like a social optimum. That is the distinctly economic bit.

As James Heckman said of the narrow 'program evaluation' approach:

... the economic questions answered and policy relevance of the treatment effects featured in the program evaluation approach are often very unclear.

He goes on to say of these models that:

... they do not allow for interpersonal interactions inside and outside of markets in determining outcomes that are at the heart of game theory, general equilibrium theory and models of social interaction and contagion.

Deaton and Cartwright give a hypothetical example of an experimental result in agriculture that improves crop yields. As they indicate, even if the experiment has internal and external validity, what if demand in the relevant market is price inelastic? Then the widespread roll-out of the treatment could reduce prices and revenues and possibly leave farmers worse off!

One might add that the interesting *economic* question then becomes: what might farmers do about this? And would consumers be happy about the collective bargain farmers might make to overcome their prisoner's dilemma?

Deaton and Cartwright conclude:

... without the structure that allows us to place RCT results in context, or to understand the mechanisms behind those results, not only can we not transport whether 'it works' elsewhere, but we cannot do one of the standard tasks of economics, which is to say whether the intervention is actually welfare improving. Without knowing why things happen and why people do things, we run the risk of worthless casual (fairy story) causal theorising and have given up on one of the central tasks of economics and other social sciences.

Again, to be clear, this focus on data and empirical rigour is to be welcomed and represents a big opportunity. We just have to keep reminding ourselves that it's not everything. It is a complement to a strong conceptual framework and qualitative understanding of real-world context.

A final aside: I know it is just a matter of vocabulary, but I confess to some squeamishness about the language of 'experiment', 'treatment' and 'control' in social policy settings (that human thing again). The knowledge (supposing it to be knowledge) that cause A led to effect B seems relevant to the omnipotent central planner, with the world spread out before them like a chess board. But that seems more like engineering than economics.

To me, the question for economists is not how the planner can fix the world, but how everyday people go about trying to improve theirs. And what are the policies, settings and institutional rules that help or hinder them in the process?

A policy case study: working from home

As noted above, my perspective on method and history of economic thought is that of a policy practitioner, not a scholar. So I think about it in the context of policy issues. Take one issue of a contemporary interest: the rise of working from home during and after the COVID pandemic.

What should we make of it? What, if anything, is the role for government policy? And how might economists think about it?

I use this example purely to knock down some straw men – that is, to demonstrate the weakness of an exclusive focus on any single approach. If the issue had arisen 30 years ago and was the subject of my honours thesis, I would have done the neoclassical thing: sketched out a basic constrained maximisation model.

In fact, at the PC we did just that (as part of our overall approach): representative firm, price taking and profit maximising. Utility maximising household with a labour-leisure choice and a budget constraint.

That sort of model does demonstrate some insights. For example, if there are now two types of labour (from home and the workplace) but only one wage, then the employer will tend to set the amount of work from home at the point where the marginal product of labour in each setting is equalised.

That is highly unlikely to be the same point at which the household's marginal disutilities of labour (of the two types) are equalised.

So at the margin, there could be a tension. Probably the boss wants less work from home and the worker wants more. If wages could fully adjust, the story might be different. This is a useful insight, but as Chesterton might say: it explains something, but what a lot it leaves out!

The firm and household are assumed to have perfect knowledge and well understood preferences. But this is hardly realistic in a case like the sudden uptake of remote work, where we are all learning as we go.

Also, the firm and household are representative agents, but this too misses the point – it is the heterogeneity of firms and workers that is of interest. Some are better suited to remote work than others. Some want it more. Some not at all. Some workers will switch jobs, some firms will find workers with a more compatible set of preferences. Many will invest in ways to make remote work 'work' for them.

These effects, the switching and the learning, will alter the static effect of my model, no doubt quite profoundly.

What would the honours student of today do? Perhaps run an RCT, or, to showcase some econometric technique, find a quasi-experimental method like a regression discontinuity or instrumental variable to mimic a degree of randomisation.

The outcome will be an average treatment effect. It might find a causal effect: people who started to work from home were, on average and compared to a counterfactual, more (or less) productive, happy, healthy and so on.

Again, this might tell us something useful but also leaves so much out.

First, for the firm, the randomised trial is probably of limited value. Randomisation is helpful when we cannot observe the counterfactual to the treatment. It gives us a proxy in the form of a control group (or the matching sample).

But here the firm has a pretty good sense of the counterfactual, at least in respect of productivity. They know (to a first approximation) how productive people are in the office. They can make an assessment about productivity at home. Moreover, they can do this at the level of the individual, not just the average.

So the average treatment effect on productivity would surely be an extreme lower bound for the firm: they get to choose who works from home, when and how much.

For the policymaker, there is a risk of being misled by a result that furnishes an average treatment effect from a randomised work-from-home trial. Perhaps the trial showed that productivity fell or mental health worsened. If the policymaker then observes that working from home is on the rise in aggregate, they might extrapolate from the trial results and fear adverse societal effects.

But, of course, the average treatment effect would be unreliable. In the real world, the rise in working from home is by people who *are not randomised*. They are *selected*. They are selected by an economic system: a constellation of employers, managers, team leaders and workers themselves who are choosing who works from home, which days and under what circumstances, and choosing largely on the basis of their fitness for remote work.

The key is to understand that selection system and how it selects – is it efficient? Is it fair? What are the rules of the game?

I emphasise that there is nothing wrong with the RCT, just its naïve interpretation.

To fully think through the implications of remote work, one would want to do so from a variety of angles: the labour-leisure model might be one, the experimental research design is another. Then you might want to go back and read Ronald

Coase or Oliver Williamson on the firm. Or Nelson and Winter on growth and technical change. Or Joel Mokyr on the rise of the factory system.

As it happens, there is a great emerging literature on working from home, such as the work of Bloom, Davis and Barrero.

What characterises that work is its pluralism: gathering empirical evidence from surveys, new patents on remote work technologies, and, yes, random trials, alongside a theoretical framework to support intuition and prediction. Multiple strands of evidence to guide some judgment.

As Ed Leamer said of economists, “We seek patterns and we tell stories”.

That, to me, is a big part of why the history of economic thought matters. It is a protection against fads and monomania, the tendency to fall in love with technique, to think there is only one way to ‘do economics’.

For all the strengths of the erstwhile orthodoxies in economics, they can leave gaps. We often only see the gaps after the event, and only then if we know something about what the orthodoxy replaced.

Returning to Chesterton, he described his religious journey thus:

I did try to found a heresy of my own; and when I had put the last touches to it, I discovered that it was orthodoxy.

Perhaps for many of us it was something of the opposite. We went in search of the economic tradition, and as we made our way through the canon, found it led us towards pluralism.