Summary Paper for CEPR, IFS, UCL, BREAD, TCD Workshop in Development Economics 2022

Ruimin Ao, Marco Castelluccio, Sebastian Espinoza Rojas, Anusha Guha, Hao Hu, Ignasi Merediz Sola, Emma Scandolo, Rachel Tan, Yasuka Tateishi and Davide Zufacchi

The conference was held in London on 14-15 October 2022. Cutting-edge research in low- and middle-income countries was presented, and the participants enjoyed discussing a wide range of topics, including agriculture, climate change, credit market, education, family, firms, health, labour market, and political economy and public procurement. In this piece, PhD students at University College London (UCL) recapture the papers presented and provide a brief overview of what they have learned from the conference.

Agriculture

The Long-Run Development Impacts of Agricultural Productivity Gains: Evidence from Irrigation Canals in India

Sam Asher, Alison Campion, Douglas Gollin and Paul Novosad

For a long time, agriculture has been the most important (if not the only) sector in human society. Structural transformation, specifically the transition from a farming-based society to one with well-developed industrial and service sectors, is heavily dependent on productivity growth in the agriculture sector. Increased agricultural productivity can free up labour and boost wealth, allowing people to pursue careers in other sectors or invest in machinery that will ultimately facilitate industrialisation.

In this paper, the authors focus on the extensive canal construction in India to unravel the long-run effects of agricultural productivity growth. They find that access to canals substantially improves irrigation intensity and land productivity in the long run. While population density became higher in the treated regions, such increases did not translate into structural transformation.

The agricultural productivity shocks come from the canal constructions in India in the 1970s. The canals span over 300,000 kilometres and benefit over one in four Indians or 130,000 villages. The authors combine different datasets as sources for demographic characteristics, settlements, agricultural productivity, and water resources from canals. The outcomes are measured between 2011 and 2013.

The analyses employ regression discontinuity design (RDD), in which only locations with elevations lower than those of nearby canals are defined as treated, since these are the areas where water from the canals can naturally flow to. The results are threefold. First, looking at the agricultural outcome, canal construction increases the irrigation areas and enhances the likelihood of growing water-intensive crops. It also leads to more land being cultivated and to higher yields. Specifically, such increases in the yield come from winter crops, which grow during dry winter months.

Second, the time horizon of the paper allows for the analysis of the long-run effects, especially on non-agricultural outcomes. In line with the increased farming activities, treated areas also experience higher population density due to increased labour demand. However, the population increase does not translate into the development of other sectors in treated areas. When comparing towns (urban) and villages (rural), the population gain is concentrated in towns, although there are few observed shifts in activity between sectors. Despite some positive urban growth, the back-of-the-envelope calculations show that agricultural intensities in rural areas increased, equilibrating the gains from agricultural productivity growth.

Finally, the impacts of canal construction on consumption vary substantially based on land ownership. While there is no detectable impact at an aggregate level, detailed analyses show that landowners (around 40% of the population) enjoy positive consumption gains. Such effects tend to increase with the size of the land owned. Similar effects are found for educational attainment, where landowners tend to have higher education. This has important distributional implications as wealthier households seem to benefit more from the canal.

The authors perform additional analyses using different specifications and definitions of treatment areas. Examples include excluding the areas within 2.5m of the RDD cut-off, because those areas, even located above the canals, can easily pump water from canals for farming. They also try limiting the sample distances to nearby canals, subsampling by the ruggedness of land and redefining the treated as those inside the canal command area. The results are robust to these different adjustments.

— Summary by Ruimin Ao

Impact of Market Information on Cashew Producers in Guinea-Bissau

Brais Álvarez Pereira, Adewusi Mendonça, Dayvikson Raiss Laval Tavares, Sebastian Schäber and Giulio Schinaia

When agents make decisions based on imperfect information, market outcomes may be inefficient. Can providing agents with more information improve market outcomes? This paper explores this question in the context of the cashew industry in Guinea-Bissau, where producers hold imperfect information. In this industry, price fluctuations are poorly communicated, and prices vary both within and across seasons. These information frictions may limit the ability of producers to engage in temporal arbitrage. Specifically, this paper examines whether and how providing market information to producers affects their ability to determine the right quantity and the right timing to sell.

To investigate this, the authors use a two-level clustered randomised control trial among cashew producers in 290 villages in Guinea-Bissau. Among the sampled villages, two-thirds of the producers were treated, and the remaining served as the control group. The treatment group received three pieces of information via mobile phone every week: (1) simplified news on the national cashew market (e.g. delayed deliveries); (2) a range of farm-gate prices (the prices at which producers sell raw cashew nuts); and (3) news and advice on whether it is a good time to sell (e.g. we advise to sell if price X is offered).

Compared with the control group, treated producers sold more frequently and opted for multiple sales rather than concentrated transactions to reduce risks. In addition, the treated sold cashews earlier in the trading season than the control group. They did this to smooth their incomes and to face less uncertainty about current and future prices. They also attained higher sales revenues than the control group on average, even though they did not receive higher prices for the cashews sold.

Exploring several potential mechanisms, the authors determine that the findings are not driven by treated producers changing their sale locations, selling to different buyers from before, gaining the upper hand in price negotiations, keeping better records or adopting different attitudes towards risk from the control group. Moreover, information spillovers could explain why treated producers failed to receive higher prices for their cashews than the control group.

These findings suggest that providing reliable market information can affect the timing at which producers choose to sell their goods and improve market outcomes for producers.

— Summary by Anusha Guha & Rachel Tan

Climate Change

The Effects of Climate Change in the Poorest Countries: Evidence from the Permanent Shrinking of Lake Chad

Roman Zarate, Remi Jedwab, Federico Haslop, Takaaki Masaki and Carlos Rodríguez-Castelán

The permanent shrinking of Lake Chad, a large lake in sub-Saharan Africa, has caused slower economic growth in areas closer to the lake than in other areas. This study shows that, as real income and welfare decreased in the area after the shrinking, households migrated away from the lake.

The permanent shrinking of Lake Chad is a unique setting to study the long-lasting effects of climate change on the economy. According to the Intergovernmental Panel on Climate Change (IPCC), 'climate change refers to a change in the state of the climate that can be identified … by changes in the mean and/or the variability of its properties and that persists for an extended period, typically decades or longer'. The novelty of the paper lies in the fact that it studies a phenomenon that is (1) almost three decades long – in contrast to the vast literature studying sudden climate shocks, (2) plausibly exogenous because of the natural reasons causing the shrinkage and (3) considerable in magnitude – Lake Chad lost 90% of its area between the mid 1960s and the late 1980s.

The authors construct a novel dataset that tracks population counts over 70 years in Cameroon, Chad, Niger and Nigeria at the granular (subdistrict) level. They first exploit the panel data to conduct a difference-in-differences regression and second develop a quantitative spatial model to reconcile the empirical findings with multiple locations and four economic sectors (fishing, farming, herding and urban). The model also allows for the analysis of the aggregate and spatial effects of policies related to migration, trade, land use, roads and cities.

Although a priori the effects of the shrinking on population and local economy are ambiguous, the authors find that the decrease in water supplies and the rise in temperatures outweigh the benefits induced by increased land supply. The effects were particularly pronounced in Chad, where locations close to the lake grew 46% more slowly than other areas. Furthermore, the effects were persistent – they are negative and statistically significant even in 2012, when the lake recovered a part of its surface – and decreased with distance from the lake, being null in locations 300km or farther away. The population drop was driven by rural areas and the model helps explain the reason why it was the case.

Because the model is static, the authors also examine the counterfactuals of what would have happened if the lake had not shrunk by comparing the 1965 market allocation steady state. They also divide the main shock into separate sub-shocks and assess the relative importance of each, with the fishing sub-shock being the most consequential. Finally, they examine how various interventions may mitigate the adverse effects of the shrinkage and show that migration can alleviate the negative effects of climate change, calling for an internationally coordinated migration policy.

— *Summary by Marco Castelluccio*

Credit Market

How Do Borrowers Respond to a Debt Moratorium? Experimental Evidence from Consumer Loans in India

Stefano Fiorin, Joseph Hall and Martin Kanz

According to a recent report from the World Bank, more than 80 countries have made debt moratoria available for individuals and firms in response to the COVID-19 pandemic and economic downturns. These debt forbearance policies, which allow borrowers to postpone loan payments, help borrowers during a liquidity crisis and strengthen banks' relationship with their borrowers. However, this may increase moral hazard by generating expectations of lenient credit enforcement and future forbearance, which undermine the effect of goodwill.

In this paper, the authors employ a nationwide field experiment in India to study borrowers' responses to debt moratorium offers. They find that such offers do not necessarily cause moral hazard: credit discipline improves and none of the moratorium offers reduces loan repayment. Their findings support the positive side of the assumptions on customer relationships and borrowers' beliefs.

The authors designed two experiments on a large lender in India and its late-paying customers with loans on mobile phones. In the first experiment, they randomly offered a three-month payment deferral to treated borrowers in two different ways: a 'relationship offer', which told customers that they were being valued by the lender, and a 'regulator offer', which told them that the offer was made due to guidance issued by the regulator. The authors also designed a placebo that sent the same texts as the relationship offer, except that no deferral was offered. Their findings suggest that the relationship offer reduced default rates by 6.9%, while there was no effect of the regulator offer or the placebo offer. In terms of the long-term effect, the relationship offer improved timely repayment, while no effect was found in groups with the regulator offer and the placebo offer.

A consistent repayment record on a single loan is not sufficient to demonstrate that forbearance does not generate moral hazard. To rule out evidence of expectation of lenient enforcement from borrowers, the authors used credit bureau data to check whether borrowers' overall credit discipline affected other loans. The evidence suggests that improved repayment did not come at the cost of reduced repayment of other loans.

What are the reasons for these positive repayment results? The authors propose four explanations: enhanced ability to repay, short-term reciprocity that drives repayment prioritisation without future benefits, long-term reciprocity that strengthens the relationship between lenders and borrowers, and a strategic behaviour of 'pay now to borrow more'. The authors conducted a second experiment on product adoption in addition to the attributes on borrowers' updated beliefs. They show that the third reason dominates in this case: borrowers who were offered a moratorium by the lender have a higher demand for non-credit products marketed by the lender, and such effect lasts for several months.

This research shed light on how people would react to debt forbearance offers. As a policy implication, there is room to explore the optimal moratorium policy design, because (1) moral hazard effects seem to be overstated and (2) an emerging offer during a crisis would make customers prioritise their repayment.

— Summary by Hao Hu

External Validity and Implementation at Scale: Evidence from a Migration Loan Program in Bangladesh

Harrison Mitchell, A. Mushfiq Mobarak, Karim Naguib, Maira Emy Reimao and Ashish Shenoy

Development economists often start a new project from a pilot experiment with an aim of later implementing at a large scale. Funding opportunities for explanatory or small-scale projects are available to support researchers reformulate a project design before extending it. A natural assumption of scaling up would be that results from pilot studies will be replicated, particularly if a target population has a similar profile or a project is carried out in a similar context. However, a considerable number of studies have failed to establish external validity when they scale up.

This paper makes a unique contribution to the literature by identifying the factors that contribute to replication failure. The authors analyse the consequence of scaling up a migration loan project in Bangladesh. The low-interest loan was intended to relax the credit constraint on landless households participating in temporary migration during the agricultural lean season. An earlier intervention in selected districts found that the project facilitated seasonal migration by offering up-front costs, which translated into an increase in earnings and consumption. In 2017-18, the authors scaled up the randomised control trial to cover all the divisions in Bangladesh. Surprisingly, the large-scale intervention induced less migration than the pilot, although the migration rates of the control group were comparable. While the pilot experiment found an increase in seasonal migration by 25-40 percentage points, the large-scale project only increased migration by 5 percentage points, and the effect was concentrated in previously targeted regions; newly treated regions had only a marginal impact on migration probability.

The authors further explore the mechanisms behind the replication failure and identify that systematic mistargeting is the leading problem. Beneficiaries can be classified as always-takers (regular migrants), compliers (those who migrate only with a loan) and never-takers (non-migrants regardless of a loan), with the project having the largest benefit on compliers. Nevertheless, with a capacity constraint due to the scale-up, principal-agent issues arise as the compliance status cannot be easily verified. As a result, the project directed its effort to always-takers, whose incentive for seasonal migration was more salient to project implementers. The authors find no evidence of cross-village spillover effects or changes in population characteristics, and the fact that the migration pattern before treatment predicts the likelihood of recalling the loan provides evidence of this mechanism.

This research has several significant implications. First, when a project is scaled up, the implementers' capacity should not be limited by an increase in administrative burdens. A problem of mistargeting, which is otherwise not present, can be a confounder of outcomes. Second, it is crucial to identify an incentive structure of beneficiaries to avoid the delegation risk in an experiment. The study finds that mistargeting has a substantial impact on take-up and should be considered when designing policies.

— *Summary by Yasuka Tateishi*

Education

Educational Investment in Spatial Equilibrium: Evidence from Indonesia

Allan Hsiao

Investments in education not only target students locally, but also provide them with portable human capital which allows them to obtain high- paid jobs in urban areas. Differences in mobility across locations can thus induce substantial heterogeneity in returns to education. When designing policies aimed at increasing aggregate output through educational investments in low-schooling areas, governments often face an equity-efficiency trade-off due to mobility. Greater mobility reduces the urban-rural inequality in individuals, who can migrate to earn higher urban wages; however, the urban-rural inequality in regional output increases due to out-migration of educated workers, so that gains are not evenly distributed across regions.

In this paper, the author studies the aggregate and distributional effects of education investments through Indonesia's INPRES programme, a primary school construction programme targeting low-enrolment regions. The findings suggest that individuals who were unable to migrate from rural to urban areas will lose their gains from education, so that accounting for the complementarity between education and migration is essential when designing such programmes.

The paper uses a difference-in-differences method to compare, first, educational returns of young students who benefited from primary school constructions with those of older students who did not and, second, the effects between high- and low-impacted districts. This leads to two main findings: returns to education vary greatly across districts and much of this variation is explained by variations in mobility. To explain these facts, the author builds a spatial equilibrium model where governments choose school construction, which builds human capital; then individuals choose education investments, which are affected by governmental choices as districts with more schools have lower education costs; finally, individuals decide whether to migrate for work.

Applying the variation induced by the programme, key model parameters are estimated. First, the programme is used as an instrument to estimate the effect of education on wages; as wages are proportional to human capital, a log-log model is employed to retrieve the parameter reflecting the elasticity of human capital with respect to education. The average elasticity is 0.7, yet there is great heterogeneity across locations. Second, exploiting the difference-in-differences variation, education and migration costs are estimated. Keeping origin and cohort-specific factors fixed, the programme reduces education costs for treated cohorts in school construction districts, and the estimates of this parameter suggest the programme gains from lower education costs. While the programme does not directly impact migration costs, they are key to the estimation as they proxy gains from market access, which significantly amplifies each effect of the programme.

Using the model and its parameters to simulate a scenario of zero school construction reveals that the programme increased aggregate output by 8%, decreased inequality across individuals by 5% and increased regional inequality by 12%. Importantly, the latter effect is driven by urban gains being larger than rural gains, so the programme remains a Pareto improvement. A scenario with no mobility reveals that the programme only increases output by 2%, implying that without migration, the programme's positive effects almost vanish. When migration costs are lower, output increases and inequality between individuals decreases, but rural regions are worse off, so that the programme is no longer a Pareto improvement. These findings highlight the equity-efficiency trade-off that governments face: complementing such programmes with policies that facilitate mobility increases output but leaves rural regions behind; yet, without mobility, the programme would have little impact on individuals and output.

— Summary by Emma Scandolo

Managing to Learn

Sabrin Beg, Anne Fitzpatrick and Adrienne M. Lucas

Like other sub-Saharan countries, Ghana is experiencing a learning crisis, where few pupils meet academic standards for their grades despite greatly improved access to primary education in the last decades. Whilst this may seem at odds with growing literature that pedagogical approaches tailoring education at the individual student's level are effective in improving learning outcomes in developing countries, this crisis may be explained by the scarcity of such approaches in existing schooling systems.

In this paper, the authors show that tackling management engagement when implementing targeted instruction can lead to improved productivity in the provision of schooling. The authors partnered with the Ghana Ministry of Education to conduct a three-armed randomised controlled trial to test how to successfully improve management and student learning within the existing schooling system.

The designed programme was implemented in 210 government schools across 20 districts, with 140 schools being randomly assigned to the first or second treatment arm and the remaining schools operating as usual. The first treatment arm focused on differentiated instruction training and instructional management, while the second expanded the first by adding training on people management. The intervention took place during the 2018-19 academic year, and data was collected through a first baseline survey, a follow-up survey in July 2019 and a second follow-up survey in September 2021, as well as two unannounced spot checks throughout.

Their findings suggest that while the two treatment arms yield symmetric improvements in instructional management, the second arm significantly improves people management scores more than the first arm. The two arms have symmetric positive effects on classroom implementations of targeted instruction techniques and on teacher attendance, with teachers becoming less tolerant of lax behaviour. Whilst attendance rates did not differ significantly between treatment and control groups, students in both treatment arms experienced an increase in their English and math test scores by 0.11 standard deviations, learning 30% more than students in the control arm. Importantly, the effects last over time: a year after implementation, teachers are implementing targeted instruction techniques, and the 2021 follow-up survey reveals that treated students continue to outperform counterfactual pupils.

The authors conclude that management is important when implementing targeted instruction, and it can be successfully improved in existing school systems. However, not all management is equally important: targeting people management on top of a programme that already targets instructional management significantly increases costs, while leaving the positive classroom and learning effects from instructional management unchanged.

— Summary by Emma Scandolo

Family

Sex-Selective Abortions and Instrumental Births as the two faces of the Stopping Rule: New Measures and World Evidence

Jean-Marie Baland, Guilhem Cassan and Francois Woitrin

Since Amartya Sen (1990) referred to the male-biased population imbalance in South Asia as ‘missing women’, economists have argued the causes and consequences of such a phenomenon. Sex-selective abortion, female infanticide, and neglect are the direct causes of the imbalance, while it is challenging to identify indirect discrimination behind gender imbalance. With son preference, fertility can be higher as parents continue to have children until they reach the desired number of boys. Resource allocation to girls will be smaller with a large household size when boys are more prioritised.

The authors focus on the stopping rule as a mechanism to explain higher mortality rates for girls in several developing countries. Their argument relies on the claim that sex-selective abortion does not substitute the stopping rule: they are fundamentally similar when focusing on pregnancy rather than births. In other words, sex-selective abortion will affect the relative probability of birth by gender, while it does not affect the decision of continuing childbearing. The authors use the terminology of ‘instrumental birth’ to denote the behaviour by which parents have children until they reach the desired gender composition and they argue that sex-selective abortion and instrumental birth can coexist as the two mechanisms of the stopping rule.

Their theoretical model predicts that girls will have more younger siblings than boys with the stopping rule. However, disentangling sex-selective abortion from instrumental birth is challenging in empirical analyses, as data on pregnancies is unavailable or unreliable in many cases. The authors then establish a simple method to identify the prevalence of stopping rules: sex-selective abortion can be detected by focusing on the share of elder girls and instrumental birth is detected by focusing on the number of younger siblings. With this method, the authors identify countries in which stopping rules are prevalent and show that stopping rules exist not only in South Asia but also in other parts of the world.

Their finding is particularly interesting as it suggests a possibility of hidden discrimination against girls in countries in which the ‘missing women’ phenomenon is not as evident as in South Asia. While further research will help uncover the economic consequences of instrumental birth, this research provides a new perspective on the demographics of developing countries.

— Summary by Yasuka Tateishi

Firms

Misallocation in Firm Production: A Nonparametric Analysis Using Procurement Lotteries

Paul Carrillo, Dave Donaldson, Dina Pomeranz and Monica Singhal

Misallocation is an important topic in the research of development economics; it explains a large part of cross-country productivity differences, which are potentially a source of underdevelopment of low-income countries. Resource misallocation refers to a situation in which inputs of a given good are not efficiently allocated across producers with different levels of productivity. Economically, it is characterised as the marginal products of inputs not being equalised across firms. Misallocation results in aggregate productivity losses, and an economy fails to maximise aggregate output. While the individual values of marginal products of inputs are required to estimate the firm-specific wedges (or distortions), they are not non-parametrically identifiable. For this reason, the existing literature has imposed assumptions on production functions or firms' optimisation behaviour.

In this paper, the authors develop two non-parametric tools for the analysis of misallocation: one to test the existence of misallocation and another to measure the welfare loss from misallocation. In addition, they use the policy shock of procurement lotteries in Ecuador's construction sector to empirically analyse misallocation in the construction services industry.

The conceptual framework of the new methods is to leverage exogenous variation in firm input utilisation to identify the moments of the distribution of individual marginal products of input. By loosening restrictions on optimisation behaviour within firms and allowing heterogeneity in production technologies across firms, the authors derive a wedge condition under unconstrained allocative efficiency of production, which is different from the condition obtained under the aggregate input constraints in existing works.

However, the test for no-misallocation is not feasible since production technologies of each firm are not observable. To solve this problem, the authors propose a solution of exploiting an exogenous change in input over time to perform Taylor expansion on the market clear condition, thereby constructing an instrumental test with standard tests for statistical independence. To measure the cost of misallocation, they apply the same methodology of exploiting input variation under the constraints on elasticity structures of supply and demand. By implementing a non-parametric IV correlated random coefficient model, they identify weighted moments of wedges that are the only components affecting the costs of misallocation.

How can we apply these tools to an empirical analysis? The authors present an example of the procurement lotteries in Ecuador as random shocks to the industry. They combine administrative data from 2008 to 2015, including transaction, sales, wages and employee's data, and corporate income tax filings. In this situation, a canonical assumption on aggregate input constraint would lead to not rejecting the null hypothesis of no-misallocation, while the new tool here that relaxed this constraint would reject the null, suggesting a misallocated production within the construction industry. The cost analysis also indicates that the deviation from efficiency results in a smaller cost than the cost estimated with the existing tool.

The new tools proposed in this paper could be a significant improvement to the existing toolbox of misallocation analysis, since it allows for heterogeneity in the production function and relaxes the constraint on aggregate input, both of which appear to be vital elements in misallocation analysis.

— *Summary by Hao Hu*

Racial Inequality, Minimum Wage Spillovers, and the Informal Sector

Ellora Derenoncourt, François Gérard, Lorenzo Lagos and Claire Montialoux

The effects of minimum wage policy are widely discussed in economics literature from various perspectives. How will such effects on labour markets be derived from the standpoint of economic disparities by ethnicity? Is minimum wage policy effective enough to reduce such gaps across formal and informal sectors? This paper addresses these questions by providing novel evidence from Brazil.

The persistence of large racial disparities is a key characteristic of Brazilian labour markets, as it is in many post-slavery and post-colonial societies. Non-white workers are over-represented in a considerably large informal sector, consisting of 40% of private sector employees in 2009. Informal workers in this case are defined as workers who have no labour contract registered in a booklet issued by the government and accordingly who are not entitled to rights and benefits. By examining a large nationwide minimum wage increase in the 2000s, the authors show evidence of decreasing racial income gaps among formal workers and improving conditions in the informal sector.

The authors analyse the effect of the minimum wage increase on the racial gap in the formal sector by wage percentiles, using a matched employer-employee dataset and a household survey. They also employ quantile regression to show how the cumulative distribution functions of monthly earnings by racial group change over time, and the result showed evidence of a reduction in racial wage inequality among formal workers. The impact of the minimum wage is particularly concentrated among low-earning workers up to the 10th percentile. In the northeast, the poorest region in Brazil, the effect is concentrated even up to the 30th percentile.

The spillover effects of the minimum wage policy appear to be strong in the informal sector. The authors use difference-in-differences to quantify such effects by assigning the strongly treated states as a treatment group and the weakly treated states as controls. There is a large increase in monthly earnings among the bottom 80% of workers, and after comparing such changes between formal and informal workers, there is neither evidence to rule out the possibility of one-to-one spillover effects, nor significant differences in wage increases across racial groups.

To understand these spillover effects, the authors also sketch evidence of different mechanisms. First, the norms on firm wage setting act as a constraint in the formal sector, and the minimum wage strongly binds in the informal sector. Second, similarity in occupation- and industry-level profile between informal and formal workers is consistent with the theoretical prediction of strategic competition among informal and formal employers. The strong spillover effects indicate that existing research may have underestimated the effect of minimum wage policies on racial gap.

A minimum wage policy could have an impact on employment decisions as well. In this paper, however, the authors find no displacement effects of employees. In the strongly treated states, there was no relative increase in informal sector employment either across or within ethnic groups; cross-sector elasticity calculations support this conclusion. In terms of the classic theoretical prediction on unemployment, a small effect is observed across and within racial groups.

— *Summary by Hao Hu*

Health

The Economic Impact of Depression Treatment in India: Evidence from Community-Based Provision of Pharmacotherapy

Manuela Angelucci and Daniel Bennett

Depression is a serious and widespread illness around the world, with various symptoms and significant impacts on socio-economic outcomes. It is more prevalent among the poor, potentially reinforcing poverty traps. Because the provision of mental health services is limited in some contexts, pharmacotherapy could be a valid substitute or complement; it reduces the need for human capital, which is especially important with the shortage of providers. There is limited evidence of the short- and long-term effects of community-based provision of pharmacotherapy in developing countries.

This paper studies the effects of pharmacotherapy on depression, socio-economic outcomes, and the potential link between mental health and economic behaviour. The authors conduct a community-based cluster cross-randomised trial in a region near Bangalore, India, in which 1000 adults with symptoms of depression, 86% of whom are female, are offered psychiatric care (PC) and livelihoods assistance (LA). They address work-related challenges, in combination with personalised assistance to help participants find employment or other income-generating opportunities.

They measure the impact of the different treatments, both during and 16-26 months after the programme, on time use, earnings of participants, human capital investment in children, and consumption, durable goods ownership and hygiene/sanitation of households, as well as several potential pathways that link depression to these outcomes.

They provide three sets of results. First, they show that pharmacotherapy reduces depression severity, especially when paired with a light-touch livelihood intervention, with larger benefits that persist longer after the treatment ends. There is no significant effect on either work time or earnings in the long term, which is consistent with the literature. Second, and perhaps most importantly, they provide evidence of large causal benefits of depression treatment on child human capital investment after the PC; it increases school enrolment and decreases labour force participation, particularly for older children. Third, they show that the treatments mainly affect behaviour through a preference pathway; they increase risk intolerance and reduce the incidence of negative shocks.

This project derives some policy implications. First, it shows the potential of using pharmacotherapy as a complementary tool to address mental health issues in some contexts. An additional contribution is the finding that pairing the treatment with livelihood assistance sessions increases the persistence of the positive effects in a cost-effective way. However, the mechanisms behind the combination of these two treatments are not yet comprehensively understood, suggesting that further research on the topic is needed.

— Summary by Ignasi Merediz Sola

Labour Market

Meet Your Future: Experimental Evidence on the Labor Market Effects of Mentors

Livia Alfonsi, Mary Namubiru and Sara Spaziani

Youth underemployment has become a global issue with relevant consequences on well-being, economic growth and development. One of the things that could explain this phenomenon is the insufficient information that young jobseekers can access in the search process. For instance, they may not know how to do a proper search or where and how to find the vacancies that need to be filled. They may also have an incorrect idea of how much they can earn in entry-level positions or how the market will reward experience in the medium term. In fact, recent studies have discussed how young people have been developing overoptimistic beliefs regarding their potential outcomes in the labour market.

This research attempts to determine whether personalised mentorship by experienced workers could improve young people's labour market results by providing them with relevant information in the job-search process. To answer the question, the authors carried out a randomised mentorship programme in Uganda for students in their last year of vocational education. They find evidence that mentorship could improve young people's labour market outcomes by increasing the probability of finding a job by 27% three months after graduation and increasing their earnings by 18%. They show that, while students work more and earn the same in the short term, their labour market trajectories become steeper in the medium run.

By analysing the mentorship process, the authors find that the primary mechanism behind the improvement in young people's labour market outcomes is the information they receive about entry-level jobs and labour market dynamics. Surprisingly, they do not find evidence that referrals, specific job-searching tips, or entry-level job recommendations are relevant drivers of the results. The authors present evidence that mentored students (1) have less overoptimistic beliefs about starting wages, (2) are more willing to accept unpaid jobs and (3) raise and adjust their expectations regarding the returns to experience. Specifically, they demonstrate that mentored students have lower minimum wages they are willing to accept by 30%, increase their willingness to accept an unpaid job by 13% and reject fewer job offers. Additionally, the authors show that overoptimistic students are contributing to the results, and that the programme is less effective on students with less optimistic expectations.

— *Summary by Sebastian Espinoza Rojas*

Unpacking Neighborhood Effects: Experimental Evidence from a Large-Scale Housing Program in Brazil

Carlos Alberto Belchior, Gustavo Gonzaga and Gabriel Ulyssea

In both developed and developing countries, public housing is a widespread welfare programme. Giving low-income individuals a significant wealth shock and allowing them to relocate away from home can significantly influence their lives. Theories speak of the mixed effects of changing neighbourhoods on labour market outcomes: some argue that such effect is dictated by the quality of peers that the neighbourhood offers, while others claim labour market access in the neighbourhood to be the key driver.

In this paper, the authors find that neighbourhood effects are heterogeneous and range from no impact at all to a 7.5 percentage point decrease in formal employment, and that distance from job opportunities, rather than network quality, is the main mechanism through which neighbourhoods impact employment.

The authors take advantage of a large housing programme in Rio de Janeiro called *Minha Casa, Minha Vida*, which provided subsidised housing for individuals in the city's outskirts. The programme is attractive for evaluation purposes for two reasons. First, houses were randomly distributed to eligible beneficiaries, which makes potential beneficiaries who did not receive the programme benefit a valid counterfactual for those who did. Second, in a single lottery of the programme held in 2015, beneficiaries were randomly assigned to neighbourhoods with different characteristics, ruling out the possibility of results being driven by the self-selection of individuals in certain neighbourhoods.

The paper studies this 2015 lottery using programme eligibility as an instrument to test the labour market effects of moving. The instrument has a strong impact on take-up, with half of the drafted individuals taking up treatment; moreover, moving changes living conditions and improves housing quality but reduces income and education levels of the residential neighbourhoods. Findings suggest that receiving a house through the programme reduces the probability of formal employment by 1 percentage point and reduces employment hours and months; however, it has no effect on wages conditional on employment. Importantly, the authors find employment responses are dramatically different depending on the neighbourhood to which people move.

Finally, the authors use partial identification tools to unpack why neighbourhoods are important for labour market outcomes. Focusing on market access, network quality, amenities and crime rates as key mechanisms and imposing theoretical sign restrictions on their effects on labour market outcomes, they estimate the bounds on the importance of each mechanism. Labour market access is found to be a necessary mechanism for the observed patterns, explaining 82-93% of neighbourhood employment effects. Thus, the findings provide strong evidence in favour of the spatial mismatch hypothesis.

— Summary by Emma Scandolo

Political Economy & Public Procurement

Bureaucratic Nepotism

Juan Felipe Riaño

In many developing nations, governments often have limited capabilities. This places limitations on what the government can accomplish, the policies it can enact and the services it can provide. The literature has mostly focused on the quality of human capital within bureaucracies, including incentives within the public system (Ashraf and Bandiera 2018, Besley and Ghatak 2018, Bénabou and Tirole 2006, Khan 2021), as well as the selection of bureaucrats (Dal Bó, Finan and Rossi 2013, Deserranno 2019, Ashraf and Bandiera 2018). Bureaucratic nepotism is a phenomenon that can affect both the allocation and compensation of public workers. As a result, it plays a crucial role in determining state capability. Nevertheless, we know surprisingly little about its magnitude, operation and effects. This is because we often lack comprehensive data on family connections, performance and career paths of public sector workers.

The author presents the first comprehensive empirical analysis of bureaucratic nepotism in an entire modern bureaucracy, drawing on extraordinary data on the Colombian public sector for seven years (2011-17). His extensive data collection effort, and a new method of family network reconstruction, allow him to recreate the entire career paths of 1,083,714 civil servants and their extended family networks spanning more than 2,400,000 individuals via consanguinity and affinity ties (i.e. through blood or marriage ties).

The author first presents new empirical data on family ties in public administration, then studies the impact of family connections within bureaucracies using two sources of identifying variation.

Family connections are pervasive. He finds that around 38% of workers have a relative in public administration, 18% have a family connection to a top bureaucrat and around 11% work with a family member in the same agency. Furthermore, when there are family connections, they are between close family members. Finally, more nepotism is associated with lower performance, as expected.

Building on Xu (2018), the author looks at top bureaucrat turnovers to evaluate how changes in family connections to public sector managers and advisors impact the allocation and compensation of public sector workers. He finds that, when family-connected to a top bureaucrat, a public sector worker is 40% more likely to be hierarchically promoted and receives a 2% to 5% higher salary. This is not necessarily evidence of favouritism. It is possible that family connections simply allow top bureaucrats to do a better job at screening employees. Nevertheless, this seems not to be the case. Looking at pre-promotion characteristics, there is at best no difference in education, public sector experience and general conduct.

Finally, the author looks at the effectiveness of anti-nepotism legislation enacted in 2015, which prohibited top bureaucrats from appointing, designating, nominating or contracting (directly or indirectly) any family member up to the fourth degree of consanguinity. The legislation seemed to be effective in reducing the number of (illegal) family connections (15% decrease). However, it did not improve workforce quality (40% of illegally connected bureaucrats kept their jobs, and 30% were simply reshuffled across public sector agencies) or prevent kin favouritism. Indeed, top bureaucrats strategically reacted to this reform by substituting hierarchical promotions for wage increases, which were not explicitly prohibited by anti-nepotism legislation.

— *Summary by Davide Zufacchi*

Political Power-Sharing, Firm Entry, and Economic Growth: Evidence from Multiple Elected Representatives

Harsha Dutta, Pulak Ghosh, Arkodipta Sarkar and Nishant Vats

Does the number of politicians affect the effectiveness or ineffectiveness of political institutions? Having multiple politicians improves checks and balances and therefore minimises corruption, while it can also lead to coordination failure or duplication of effort. This can relate to the argument made by Bandiera, Prat and Valletti (2009), who distinguish between the preference of government officials for corruption (active waste) and the incapacity to reduce costs in public procurement (passive waste). Theories would predict that political power-sharing decreases active waste while increasing passive waste. However, it has not been widely discussed whether one weighs more than the other.

This study contributes to the discussion by analysing how political power-sharing affects private firm activities in India. The authors use the quasi-random variations in the number of politicians that bureaucrats report to. The electoral boundaries in India are not identical to the administrative boundaries, as they are exogenously set to equalise the land sizes under each politician's control. By overlaying the two boundaries, the authors distinguish between split (multiple politicians) and unsplit (single politician) blocks.

First, the authors use geographic regression discontinuity to examine the border between split and unsplit blocks. The characteristics of these blocks are similar, and the villages in unsplit blocks that are close to the border can be thought of as a counterfactual to those in split blocks. They find that closer to the border, the number of firm entries, employment and nightlight intensity are higher in split blocks than in unsplit blocks. Moreover, the demand for unemployment assistance is lower in blocks with multiple politicians, suggesting an increase in private sector activities in split blocks.

Second, they examine the impact of revisions in electoral boundaries in 2008. The electoral borders had remained unaltered since 1977, and the modification was completed through a politically neutral process. Using the differences-in-discontinuity approach, the authors find that blocks that changed from unsplit to split saw a higher level of firm entry than those that remained unsplit. They observe the opposite effect for blocks that changed from split to unsplit.

The authors further conduct a series of analyses to uncover the mechanism of such effect. They find that projects in split blocks take less time for approval and that the entry of smaller firms is more salient. These findings show that lower regulatory costs stimulate private sector activities. Furthermore, the authors show that the effect is stronger if the politicians are from different political parties, implying that checks and balances are the driving mechanism for removing the bottleneck for private firm entry.

To conclude, their findings suggest that having multiple politicians will improve the efficiency of political institutions, and ultimately contribute to economic growth through boosting private sector activities.

— *Summary by Yasuka Tateishi*

Donor Contracting Conditions and Public Procurement: Causal Evidence from Kenyan Electrification

Catherine D. Wolfram, Edward Miguel, Eric Hsu and Susanna B. Berkouwer

Having stringent procurement regulations would improve project quality, while at the same time increasing administrative costs. Although regulatory burden itself is to be borne by government officials or contracting agencies, it can inflict a loss on project beneficiaries when it affects the service delivery.

This study sheds light on the role of contracting conditions by analysing Kenya's nationwide electrification project (Last Mile Connectivity Project). The authors exploit the variations in donor regulations of public procurement to identify the causal link between administrative costs and project quality. The World Bank and the African Development Bank (AfDB) were among the first participants of the project in 2015. The two institutions had nearly identical technical requirements for contractors, with the main distinction being in the contracting process. The World Bank had a separate multi-phase contractual process, whereas the AfDB created a bundled contracting. Furthermore, the World Bank undertook ex-post audits, which were not implemented by the AfDB. As a result, the World Bank required greater contracting procedures than the AfDB. To quantify the effect of different contracting conditions, the authors use two layers of (quasi-)randomisation: an overlap in geographical coverage by two donors and a random selection of monitoring sites for electric connectivity.

They find that the World Bank's electrification was delayed by an average of 10 months due to the required staff time. There was no difference in the quality of outage or voltage in the short run. Thus, the AfDB projects appear to be more efficient since they produce comparable outcomes with lower administrative costs. However, the World Bank outperformed in terms of construction quality (e.g. installation of a cap to prevent rain entering a pole), which could affect the long-term quality of electric connectivity. The authors then conducted a randomised audit, and the quality was modestly improved only at the AfDB sites. They conclude that combining less stringent contracts with rigorous audits may be the most cost-effective way to ensure construction quality.

Their findings suggest that a rigorous contracting regulation can have a negative externality. This has significant policy implications, as the administrative procedures of development projects are not designed to prevent such unintended impact. It is not evident whether such externality emerges due to the presence of other donors with less stringent criteria, or whether the contracting exercise has a spillover impact that accelerates or lowers the externality. Further research on these aspects will help uncover the optimal contracting exercises in a donor-funded project or public procurement in general.

— Summary by Yasuka Tateishi

References

Ashraf, N. and Bandiera, O., 2018. Social Incentives in Organizations. *Annual Review of Economics*, 10 (1), 439-463.

Bandiera, O., Prat, A. and Valletti, T., 2009. Active and Passive Waste in Government Spending: Evidence from a Policy Experiment. *American Economic Review*, 99 (4), 1278-1308.

Besley, T. and Ghatak, M., 2018. Prosocial Motivation and Incentives. *Annual Review of Economics*, 10 (1), 411-438.

Bénabou, R. and Tirole, J., 2006. Incentives and Prosocial Behavior. *American Economic Review*, 96 (5), 1652-1678.

Dal Bó, E., Finan, F. and Rossi, M. A., 2013. Strengthening State Capabilities: The Role of Financial Incentives in the Call to Public Service, *The Quarterly Journal of Economics*, 128 (3), 1169-1218.

Deserranno, E., 2019. Financial Incentives as Signals: Experimental Evidence from the Recruitment of Village Promoters in Uganda. *American Economic Journal: Applied Economics*, 11 (1): 277-317.

Khan, M.Y., 2021. Mission Motivation and Public Sector Performance: Experimental Evidence from Pakistan. *Working Paper*.

Sen, A., 1990. More than 100 Million Women are Missing. *New York Review of Books*, 37 (20), 61-66.

Xu, G., 2018. The Costs of Patronage: Evidence from the British Empire. *American Economic Review*, 108 (11), 3170-3198.