advantage

The Magazine of CAGE Research Centre / Spring 2021 / Issue 11 / ISSN 2397-8309

In vaccines we trust?
The ties that bind
An early influencer?
Putting people's welfare first
Essay competition winner

In this issue ...

Welcome to the Spring issue of *Advantage*...

The articles in this issue showcase CAGE's policy-relevant economics at its best. While each project featured has taken advantage of data opportunities in very specific geographical or historical contexts, the findings and insights drawn from the data have a much wider relevance.

Our cover story analyses the aftermath of the shocking revelation in 2011 that the CIA undertook a fake vaccination programme to hunt for Osama bin Laden in Pakistan. The revelation and ensuing anti-vaccination propaganda campaign by the Taliban provided a unique opportunity to understand how anti-vaccination sentiment affects the uptake of vaccinations. Monica Martinez-Bravo's and Andreas Stegmann's findings are particularly meaningful in the context of the current COVID-19 vaccination drive.

Next, we move to Spain, where Manuel Bagues and Christopher Roth take advantage of a natural experiment — the Spanish lottery for military conscription — to understand how interregional contact affects sentiments of cultural cohesion and national identity (very pertinent for the UK at the moment). Then on to Germany during the Protestant Reformation: Sascha O. Becker and Jared Rubin consider the effects of Martin Luther's networks on the spread of Protestantism, revealing details about information flows and the spread of ideas that can tell us much about human interaction today. Finally, we move to Jordan, where Stefano Caria has tested out a new algorithm proposed as a modification to the randomised controlled trial — which enables larger numbers of participants to benefit from the policy schemes being tested.

In this issue we're also delighted to include the CAGE Essay Prize Winner 2020, Julia Tattersall. Her summary of Mark Harrison's paper on KGB methods in Lithuania during the Cold War is beautifully crafted and a really engaging read.

At a time when our social spheres have changed considerably, these articles offer intriguing insights into the power of human interaction, social networks and the way in which we assimilate — and act on — information. We hope you enjoy the issue.

Stephanie Seavers, *Editor*

CAGE Director: **Mirko Draca**
CAGE Research Director: **Bishnupriya Gupta**
CAGE Impact Director: **Arun Advani**
Editor: **Stephanie Seavers**

Contributors

Manuel Bagues is Professor of Economics at the University of Warwick and a CAGE Associate

Sascha O. Becker is Xiaokai Yang Chair of Business and Economics at Monash University, Melbourne, part-time Professor at the University of Warwick, and a CAGE Associate.

Stefano Caria is Associate Professor of Economics at the University of Warwick and a CAGE Associate

Mirko Draca is Professor of Economics at the University of Warwick and CAGE Director

Monica Martinez-Bravo is Associate Professor at CEMFI

Christopher Roth is Assistant Professor of Economics at the University of Warwick and a CAGE Deputy Theme Leader

Jared Rubin is a Professor of Economics at Chapman University

Andreas Stegmann is Assistant Professor at the University of Warwick and a CAGE Associate

Julia Tattersall is a fourth year undergraduate student in French and Economics at the University of Warwick.

Advantage is published bi-annually and is freely available in electronic format at **warwick.ac.uk/cage**

CAGE Research Centre
Department of Economics, University of Warwick, Coventry CV4 7AL, UK

In vaccines we trust?

By Monica Martinez-Bravo and Andreas Stegmann

In July 2011, the Pakistani public learnt that the CIA had used a fake vaccination campaign as cover to hunt for Osama Bin Laden. The disclosure of the ruse led to an increase in anti-vaccination propaganda by the Taliban. Empirical evidence suggests these events damaged the reputation of vaccines in Pakistan, eroding parental confidence in vaccines and health workers, and ultimately causing a decline in vaccination rates.

Vaccines are responsible for some of the largest improvements in human welfare in history.

The ongoing rollout of COVID-19 vaccination programmes around the world is a case in point. These programmes are regarded as a key step in ending one of the deadliest pandemics in recent history. Despite the benefits that society can derive from immunisation, vaccine scepticism is common. This holds true even in the case of well-established vaccines, for instance those that prevent measles or polio, despite the fact that these vaccines have proved safe and effective for decades.

Vaccine scepticism has frequently been fuelled by groups that oppose the use of vaccines for ideological reasons. These groups have often spread misinformation about vaccine safety and effectiveness. This is the case both in developed countries in Europe and North America as well as in developing countries such as Nigeria, Afghanistan or Pakistan, where religious extremist groups have engaged in anti-vaccine propaganda.

Despite the potentially large negative consequences of vaccine scepticism, we still have a limited understanding of how the spread of information that discredits vaccines affects immunisation rates and the demand for formal medicine more generally. To investigate this research question, we analyse data from a sequence of recent events in Pakistan.

As part of the operations to capture Osama Bin Laden in 2011, the CIA organised an immunisation campaign as cover for their espionage activities. The objective was to obtain DNA samples of children living in a compound in Abbottabad where Bin Laden was suspected to be hiding. This would have allowed the CIA to obtain definitive proof that Bin Laden was hiding there. In July 2011, two months after the actual capture of Bin Laden, *The Guardian* published an article reporting on the vaccine ruse and describing the collaboration of a Pakistani doctor with the CIA.

> **"Vaccine scepticism has frequently been fuelled by groups that oppose the use of vaccines for ideological reasons."**

News of the CIA operation caused uproar in Pakistan. Taking advantage of this, the Pakistani Taliban launched an anti-vaccination propaganda campaign to discredit medical workers and to cast doubt on vaccines. They accused health workers of being CIA spy agents and claimed that Pakistan's polio vaccination campaigns were a conspiracy to sterilise the Muslim population.

We evaluate how disclosure of the CIA operation and the subsequent Taliban anti-vaccination propaganda campaign affected immunisation rates and demand for other types of formal healthcare. Our main results on vaccination rates take advantage of a rich household survey that contains detailed immunisation records for a large sample of children, as well as precise information on their date of birth and district of residence. Our research design compares the evolution of immunisation rates from children who were born in the months before the disclosure to children who were born after the disclosure. We also look across regions that differ in the level of their initial ideological affinity to the Taliban. Parents in districts that exhibit higher support for Islamist groups are likely to have been more exposed to the anti-vaccine propaganda campaign. It is also likely that parents with an initial ideological affinity to the Taliban granted greater credibility to their anti-vaccine messages.

Our findings indicate that the disclosure of the vaccine ruse and the subsequent anti-vaccination campaign had substantial negative effects on vaccination rates: districts in the 90th percentile of the distribution of support for Islamist groups experienced a decline in vaccination rates between 23% and 39% relative to districts in the 10th percentile of support for Islamist groups.

Figure 1 provides a visual representation of these effects: there is a decline in vaccination rates in areas with high support for Islamist groups after the disclosure, while vaccination rates remained stable in areas with low support for Islamist groups.

We also present evidence of effects on disease prevalence. Exploiting variation at the district and year level, we find that a one standard deviation increase in support for Islamist groups is associated with an additional 0.8 cases of polio per district. This is equivalent to doubling the average number of cases per district. ▶

We are able to show that our estimates on vaccination rates are mainly explained by changes in the demand for vaccines rather than changes in supply. Using detailed administrative data on the timing and scope of vaccination drives, we show that the intensity of vaccination activities did not systematically differ across districts with different levels of support for Islamist groups after the disclosure of the vaccine ruse. We also find evidence that households' health-seeking behaviour changed in a way that is consistent with lower demand for formal medicine: households became less likely to consult formal doctors when their children got sick.

These results are consistent with the hypothesis that the disclosure of the vaccine ruse damaged the reputation of vaccines and formal medicine. There is substantial anecdotal evidence describing the vaccine ruse as the triggering factor and the Taliban propaganda as fueling vaccine scepticism.

Indeed, we provide a number of pieces of evidence that are consistent with the Taliban's propaganda playing an important role. We show that our effects are driven by districts where a large fraction of people do not trust or consume mainstream media. We also show that the negative effects on immunisation rates are larger for girls than for boys. This result is consistent with the notion that some parents granted greater credibility to a rumour spread by the Taliban that the polio vaccine was intended to sterilise Muslim girls.

Our findings highlight the importance of safeguarding trust in health systems. Events that cast doubt on the integrity of health workers or vaccines can have severe consequences for the acceptance of immunisation, despite its widespread positive effects. Moreover, our results suggest that events that cast doubts against vaccines can be magnified in the presence of motivated groups with the objective of seeding mistrust. These lessons are particularly relevant at a time when public acceptance of the new COVID-19 vaccines is crucial in the fight against the virus. ◀

Figure 1: Share of vaccinated children, before and after the disclosure of the CIA vaccine ruse, in areas with high Islamist support compared to areas with low Islamist support.

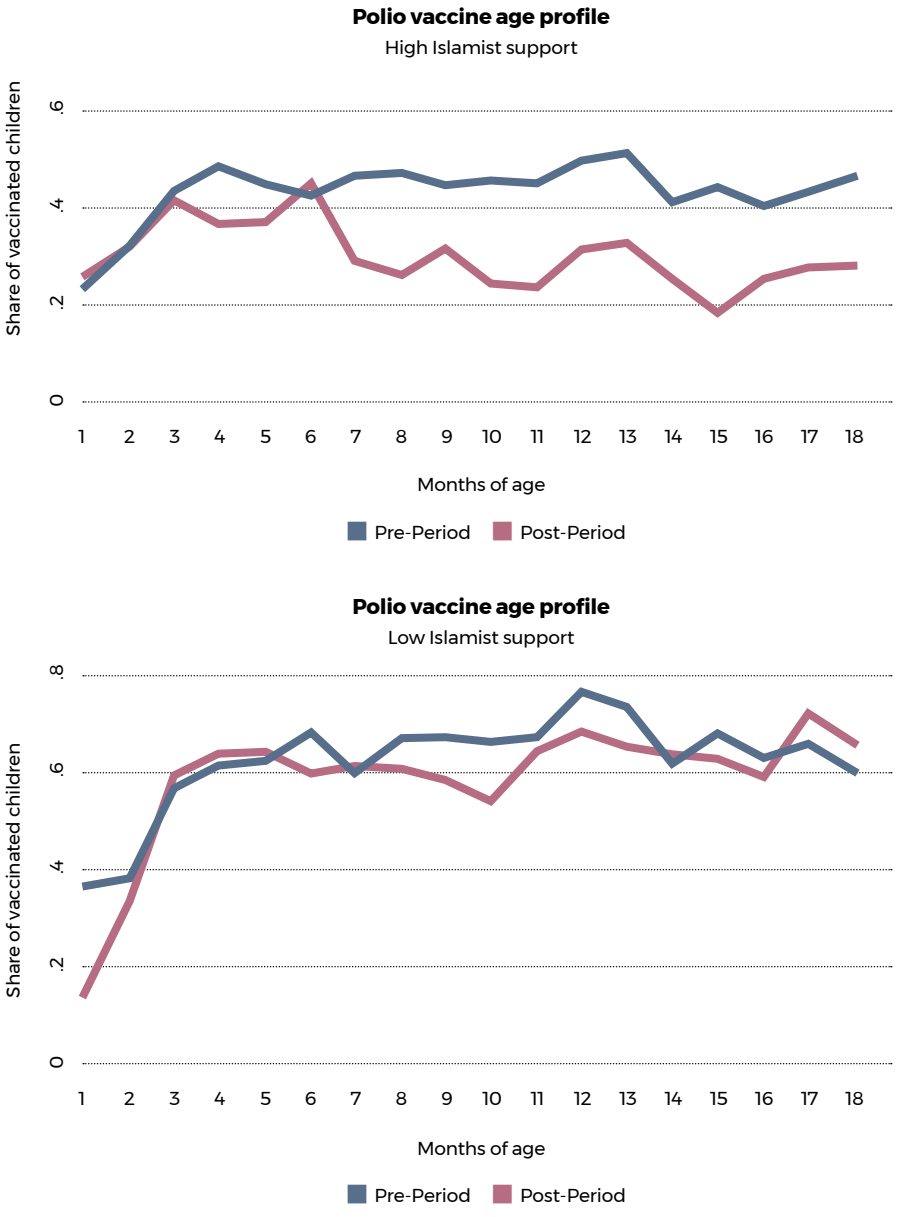

Note: High (low) support for Islamist groups is defined as the top (bottom) quartile of the distribution.

About the authors

Monica Martinez-Bravo is Associate Professor at CEMFI. **Andreas Stegmann** is Assistant Professor at the University of Warwick and a CAGE Associate.

Publication details

This article is based on the paper Martinez-Bravo, M. and Stegmann, A. (2021). In vaccines we trust? The effects of the CIA's vaccine ruse on immunisation in Pakistan. Accepted for publication in *Journal of The European Economic Association.*

The ties that bind: how contact with other regions can strengthen feelings of national identity

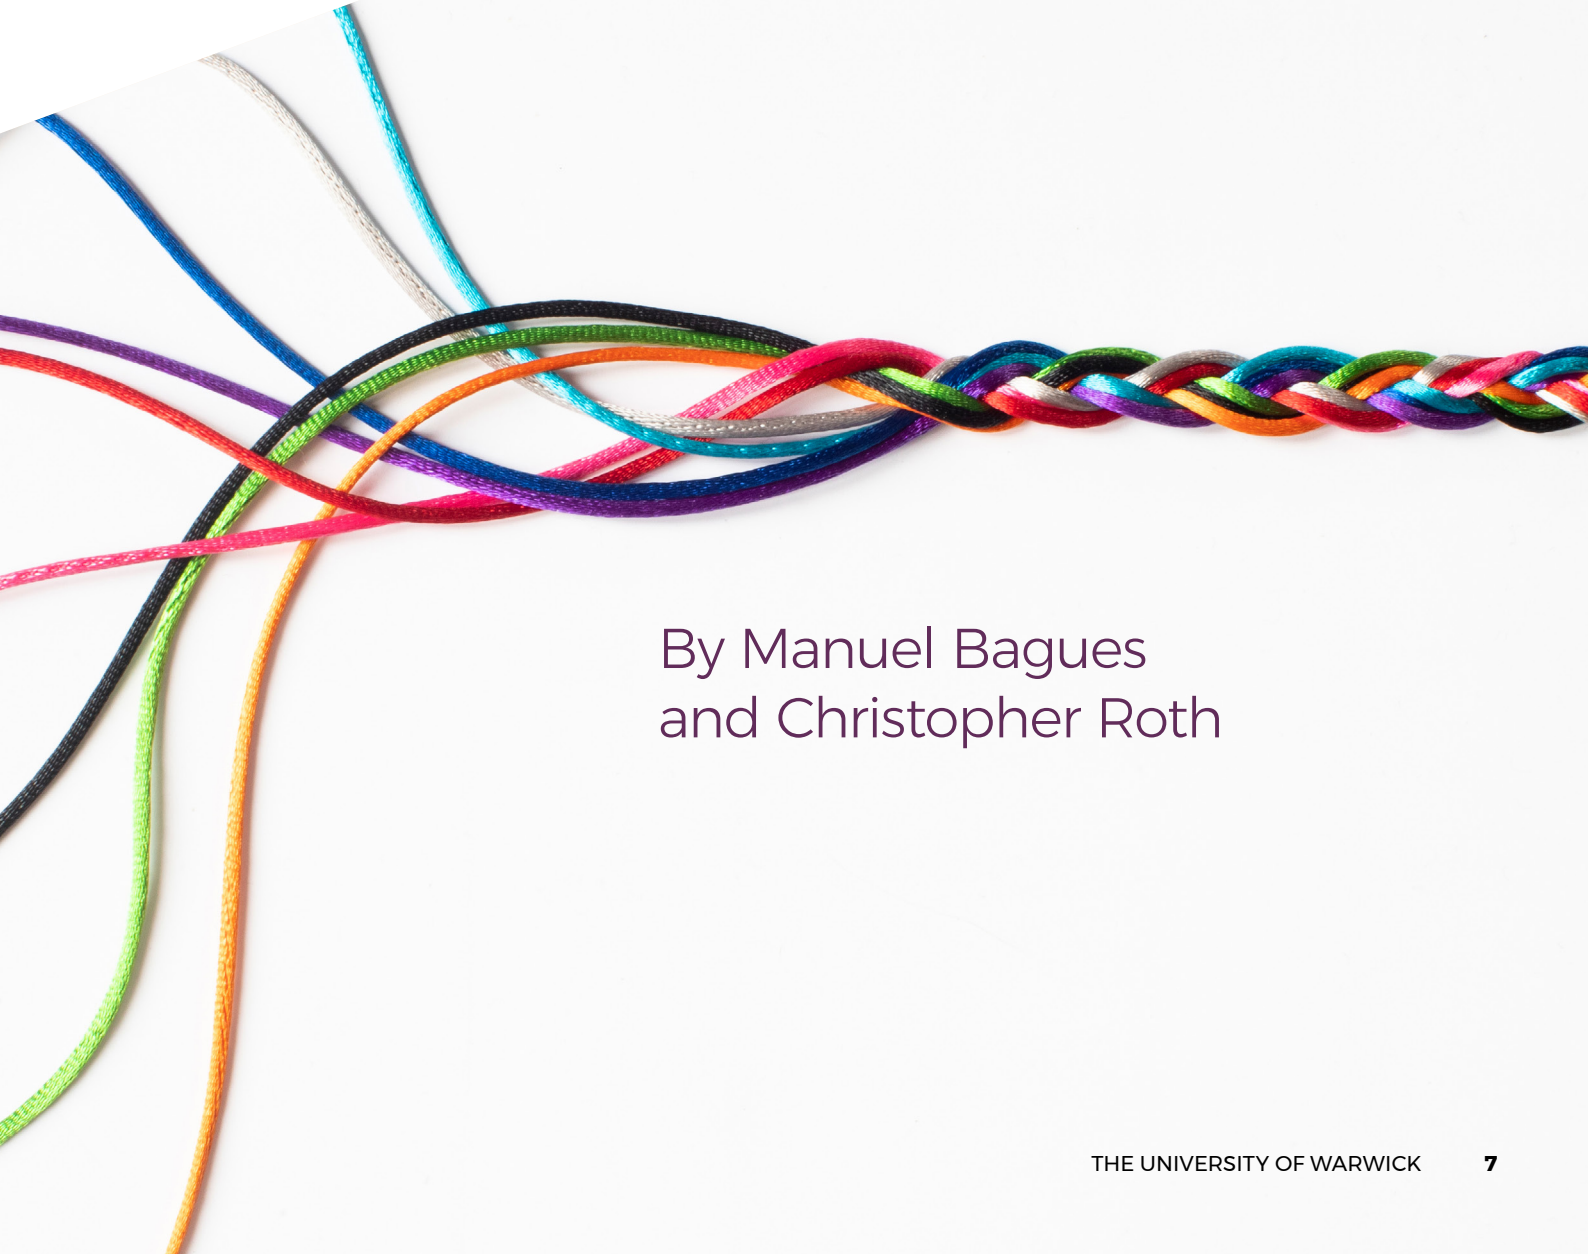

By Manuel Bagues and Christopher Roth

In January, *polls revealed* that 49% of people in Scotland and 47% of people in Northern Ireland supported independence from the UK. A majority of voters from both countries would welcome a referendum on independence. *The Times* called it a *constitutional crisis*, while former Prime Minister Gordon Brown warned that *'dissatisfaction is so deep it threatens the end of the United Kingdom'*. How can the government strengthen social cohesion and encourage a sense of shared national identity in the UK? New research shows that increased interregional contact could make a difference.

Countries across Europe face pressure from citizens who place more trust in local political movements than they do in national government. Spain, France, Belgium and Italy, in addition to the UK, have all faced secessionist movements in the last few decades. The challenge for governments — and particularly for the UK at the moment — is how to encourage citizens to look beyond regional matters and support measures for the benefit of the nation. A strong sense of national identity can be critical to encourage social cohesion; our research shows that this can be achieved through increased exposure to people from other regions. We find that interregional contact can increase a sense of national identity and grow feelings of sympathy and trust for people in other regions.

The case of the Spanish military service lottery

Historically, governments have tried to strengthen national identity through education, propaganda and resettlement schemes. Another method is conscription, in which conscripts are purposefully mixed with individuals from other regions. The hypothesis is that this mixing can strengthen interregional relationships and identification with the nation state. But contact between individuals of very different backgrounds could also exaggerate feelings of difference. To understand the effects of interregional mixing, we exploit a natural experiment: the random assignment of conscripts to national military service in Spain.

Between 1940 and 2001 young men in Spain were required to take part in military service in the year they turned 20. Until 1992 the location of service was randomly assigned through a lottery. Around one third of conscripts were assigned to serve in their own region, while the rest were allocated to other regions around the country. 97% of young men complied with the lottery outcome. Those serving outside

Figure 1: Fraction of fellow conscripts from the home region

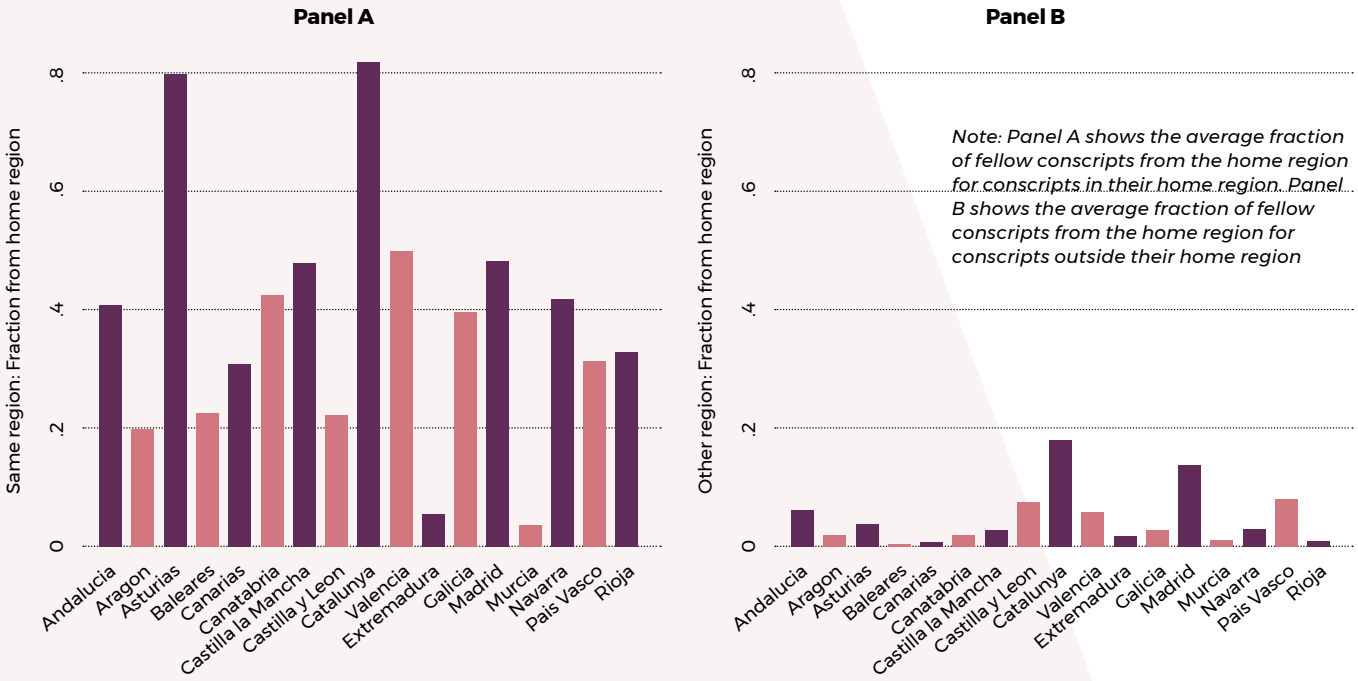

their own region were more exposed to individuals from other regions through interactions with other conscripts, as Figure 1 shows.

Examining the effects of interregional contact

We wanted to find out the extent to which conscripts were affected by the location of their military service. In particular, we wanted to know if being stationed outside their home region during early adulthood increased a sense of national identity and sympathy for other regions in later life. Spain has a number of regions with their own strong regional identities, many of which have expressed desire for independence (Catalonia, Basque Country, Galicia, Navarra and the Balearic Islands). What would be the effect of serving in another region for conscripts originating from these areas?

To find out how conscripts were affected by the location of their military service, we conducted a survey of 3231 former conscripts. We collected rich background data on each conscript including details of their birth, education, income and province of residence. We then asked them qualitative questions to determine their sympathies for Spain and its different regions. First, we asked our respondents hypothetical questions to determine respondents' sympathies towards, and trust for, people from other regions. Next, we asked i) whether they identify with Spain or their local region, ii) whether they are proud to be Spanish and iii) how they feel when they see the Spanish Flag.

Interregional contact can build trust

For conscripts who served outside their home region, we found increases in perceived honesty and sympathy towards people from the region of service, several decades after completion of service. This effect was stronger in conscripts from regions with strong regional identities.

Figure 2: Effect of serving outside the home region on a conscript's identification with Spain

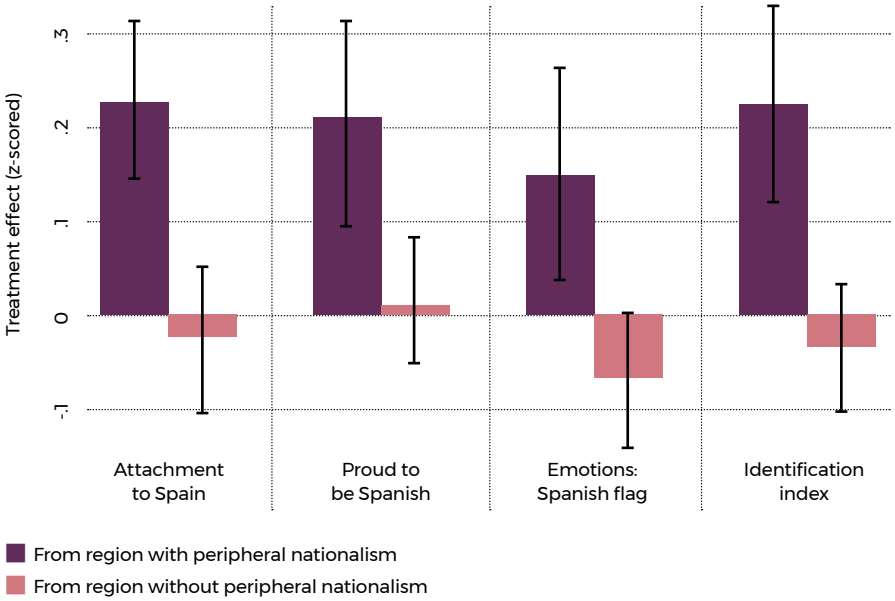

Note: This figure displays average treatment effects separately for respondents originating from a region with peripheral nationalism and for respondents from other regions. The figure also includes one standard error bars.

Interregional contact can increase feelings of national identity

We found that conscripts who served outside their home region had a stronger sense of national identity and love for Spain. However, this was only statistically significant for those conscripts originating from areas with a strong sense of regional identity (fig.2). For these conscripts, the effects of being stationed outside their home region was three times larger than for other groups. The specific region they served in, and the cultural difference between the region served and the home region did not alter the effect.

Implications

Our results suggest that experiences during early adulthood can have long-lasting and persistent effects on people's formation of a shared national identity. Governments aiming to foster interregional cohesion and policies should consider facilitating interactions between individuals from different regions. This could be by encouraging young adults to gain exposure to different areas of the country through education and apprenticeship schemes. With young voters in Scotland expected to be *more likely to vote for independence*, these findings could be particularly pertinent for the UK. ◀

About the authors

Manuel Bagues is Professor of Economics at the University of Warwick and a CAGE Associate. **Christopher Roth** is Assistant Professor of Economics at the University of Warwick and a CAGE Deputy Theme Leader.

Publication details

This article is based on the paper Bagues, M. and Roth, C. (2020). *Interregional Contact and National Identity*. *CAGE working papers* (no. 526).

An early influencer? How Martin Luther's network helped him make the Protestant Reformation a success

By Sascha O. Becker and Jared Rubin

Networks are everywhere. Of course, today we are used to social networks where information and ideas can spread at an extremely rapid rate. If someone sends a meme to two friends, and they send that meme to two friends, and so on, the idea gets to thousands of people (or more) with little effort. Social scientists call these 'network effects'. But how did networks operate before the rise of social media? We recreate the networks of Martin Luther, a key influencer of the 16th century, to find out the role they played in the spread of the Protestant Reformation.

Figure 1: Luther's spatial network

Cities to which Luther sent a letter 1501–1522.

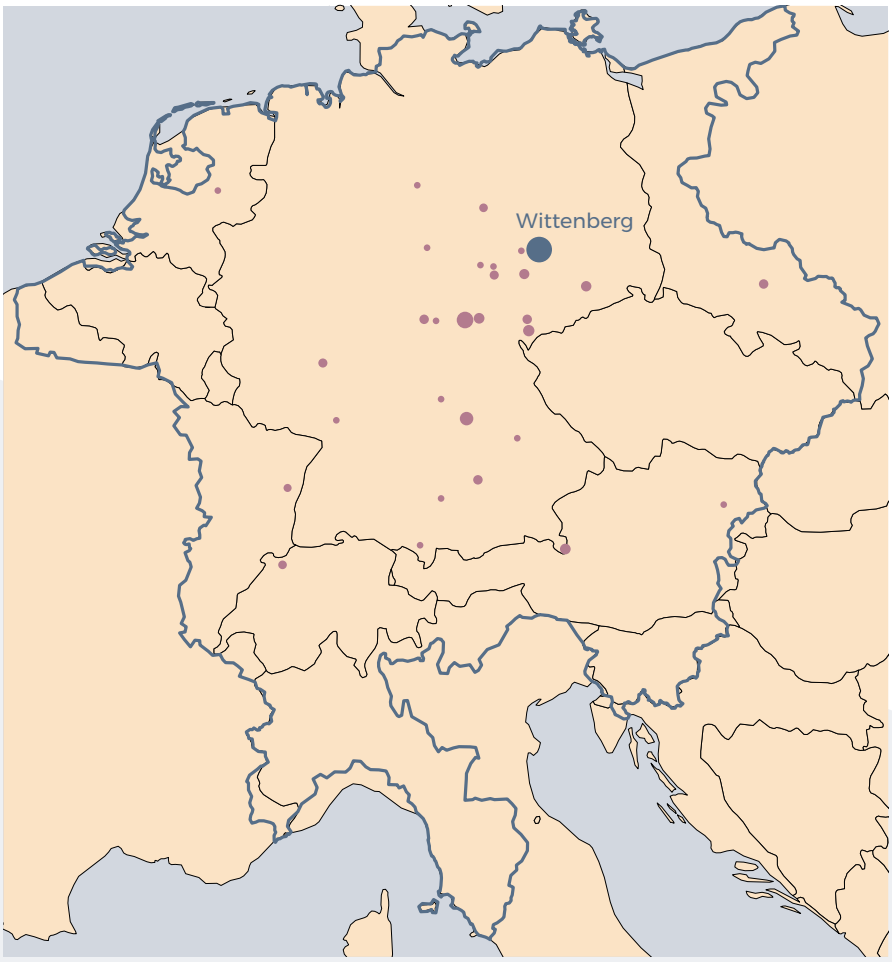

The Protestant Reformation in the 16th century was a classic situation in which networks might have really mattered. Initially, adopting the Reformation was *costly*: the Catholic Church denounced Protestant ideology as heresy, Protestants were burned alive and religious warfare was endemic. When adoption is *costly*, knowing that your neighbours adopted is useful — at least you will have allies. Networks also help facilitate the spread of information, which is incredibly important for a movement steeped in ideology.

Recreating Martin Luther's network

It is hard to recreate personal networks as we need a *lot* of information about the person in question. Who did they know? How well did they know them? Where were their contacts located? This information is hard enough to get for people in the contemporary world (unless you work for Facebook): it is next to impossible for most historical figures. ▶

Fortunately, Martin Luther (1483–1546), the leader of the Protestant Reformation, left enough traces where it is possible to recreate his network. There has been much work done by a panoply of historians documenting numerous aspects of Luther's life. Many of the *letters* he wrote survived and have recently been digitised. For the sake of reconstructing Luther's spatial network, we coded every city in which Luther sent a letter (fig.1).

Due to the *meticulous work of Georg Buchwald*, we know each city that Luther visited and when he visited it. From the *work of Hyojoung Kim and Steven Pfaff,* we know the location of each of Luther's students at the University of Wittenberg. While this data of course does not tell us every person Luther ever met, it does give us a pretty good indication of whom he had some influence with.

The role of Luther's networks in the spread of the Reformation

Luther was *influential*. He played the role of a global opinion leader based in Wittenberg. He had ties with local elites in towns across Central Europe, who in turn exerted influence in their towns. We can conceive of the process as leader-to-follower, originating with Luther and flowing to local elites through personal ties. Who Luther was connected to may therefore have mattered to the success of the Reformation.

But the Reformation may have also spread independently of Luther. Because it was costly to adopt, towns were only likely to adopt if they were connected to another town, via a network, that also adopted. We could therefore envision the Reformation as a 'virus' that spread out of Wittenberg.

So, which was it? Was Luther completely unimportant to the spread of the early Reformation? Would any 'heretical' movement coming out of Wittenberg have lit the anti-papal fuse, even in the absence of Luther? This is unlikely. As we show in Figure 2, Wittenberg was not very well connected within the Holy Roman Empire (via trade routes). It was, in Luther's words, 'on the edge of civilisation'. So the trade network was probably not enough to spread the Reformation.

Figure 2: The Holy Roman Empire trade network

The Holy Roman Empire trade network, mapped by linkages. Wittenberg, on the relative outskirts of the network, is circled.

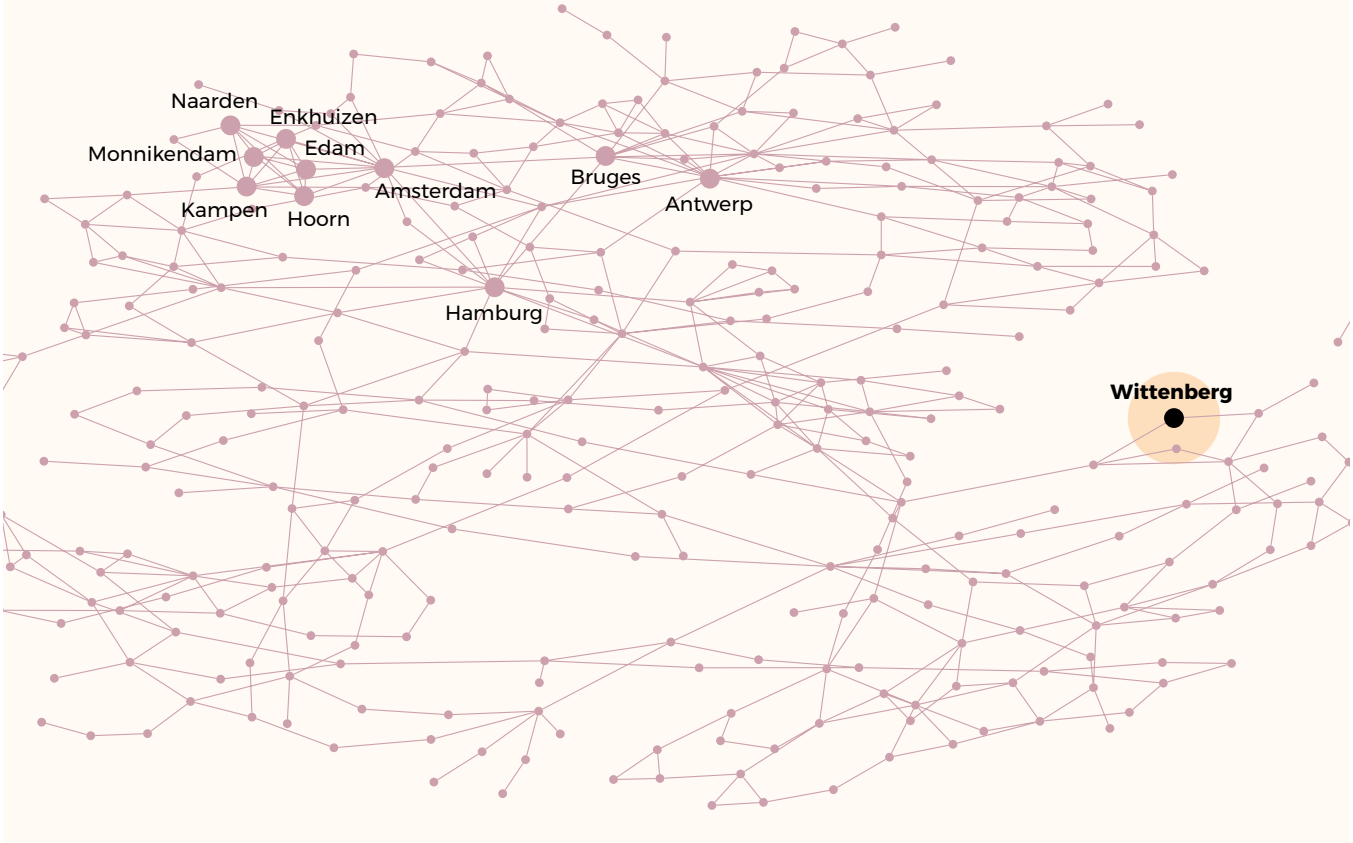

Figure 3: The role of networks in the spread of the Protestant Reformation

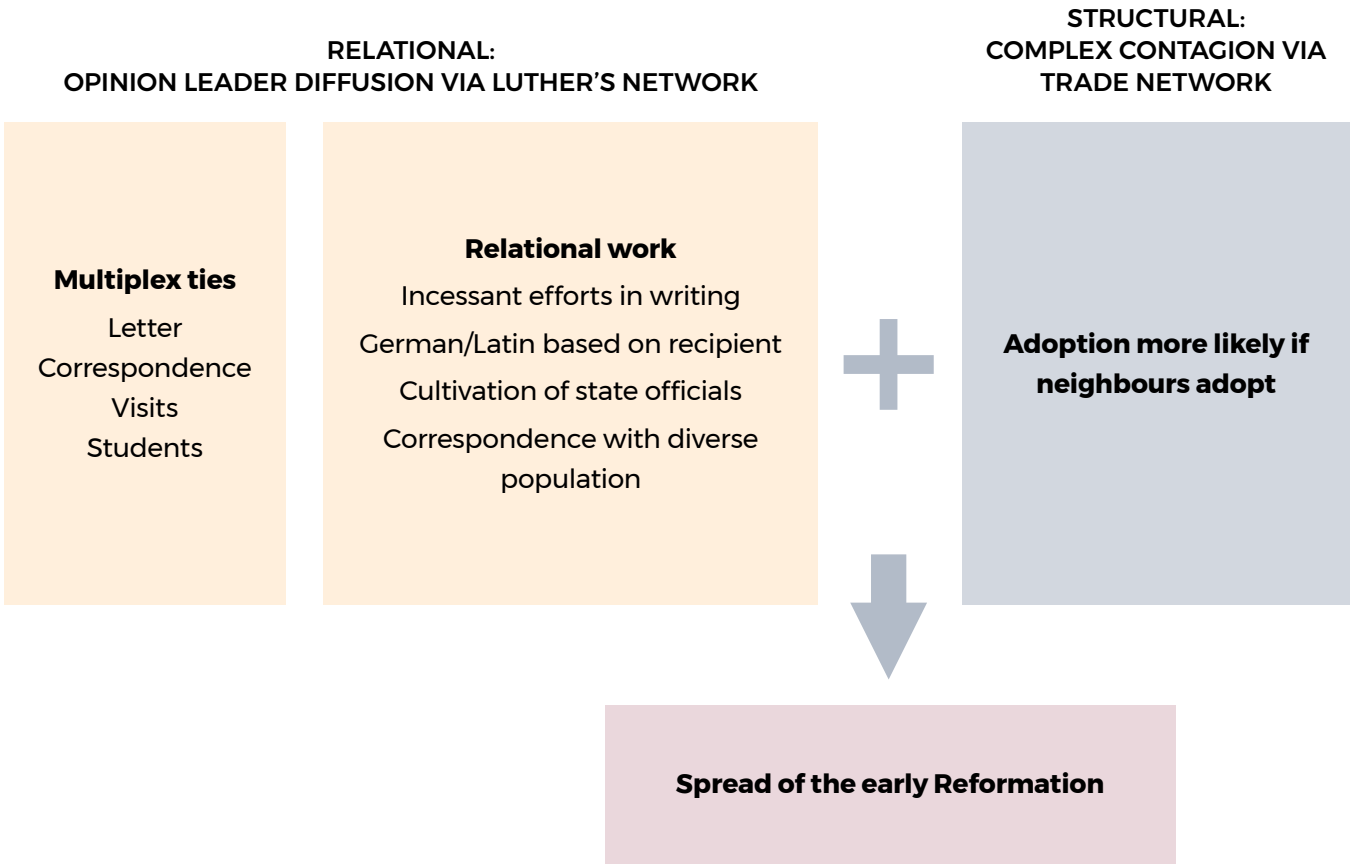

But was Luther's network enough, on its own, to account for the spread of the Reformation? Here, our analysis also finds little evidence in favour. We conduct a simulation analysis and show that Luther would have had to have been improbably 'infectious' for his network to fully explain the spread of the Reformation. In other words, way too many nodes in Luther's network would have had to adopt the Reformation *because* of Luther (and not a myriad of other causes contributing to the success of the Reformation) for this to be the dominant explanation.

So, did Luther's network matter at all? In short, yes. We argue — and our simulations and regression analyses support — that personal/relational diffusion via Luther's network *combined with* spatial/structural diffusion via trade routes helped Protestantism's early breakthrough from a regional reform movement to a widespread rebellion against the Roman Catholic Church.

Figure 3 illustrates this idea. Multiplex ties point to how Luther as an opinion leader mobilised his personal network through an ensemble of letters, visits, and student relationships. It also shows how Luther's network blends with the spatial (trade) network to create complex 'contagion' processes operating at the intersection of information flow and social influence.

Our findings show that *multiple diffusion mechanisms* can reinforce each other to facilitate the spread of ideas within a network, even if each mechanism by itself is not sufficient. This not only provides new insight for the spread of the Protestant Reformation, but also has interesting implications for the understanding of how ideas spread and gain influence today. ◀

About the authors

Sascha O. Becker is Xiaokai Yang Chair of Business and Economics at Monash University, Melbourne, part-time Professor at the University of Warwick, and a CAGE Associate. **Jared Rubin** is a Professor of Economics at Chapman University.

Publication details

The article is based on the paper Becker, S O., Hsiao, Y., Pfaff, S. and Rubin, J. (2020). *Multiplex Network Ties and the Spatial Diffusion of Radical Innovations: Martin Luther's Leadership in the Early Reformation. American Sociological Review*, 85(5): 857–894 (Open Access).

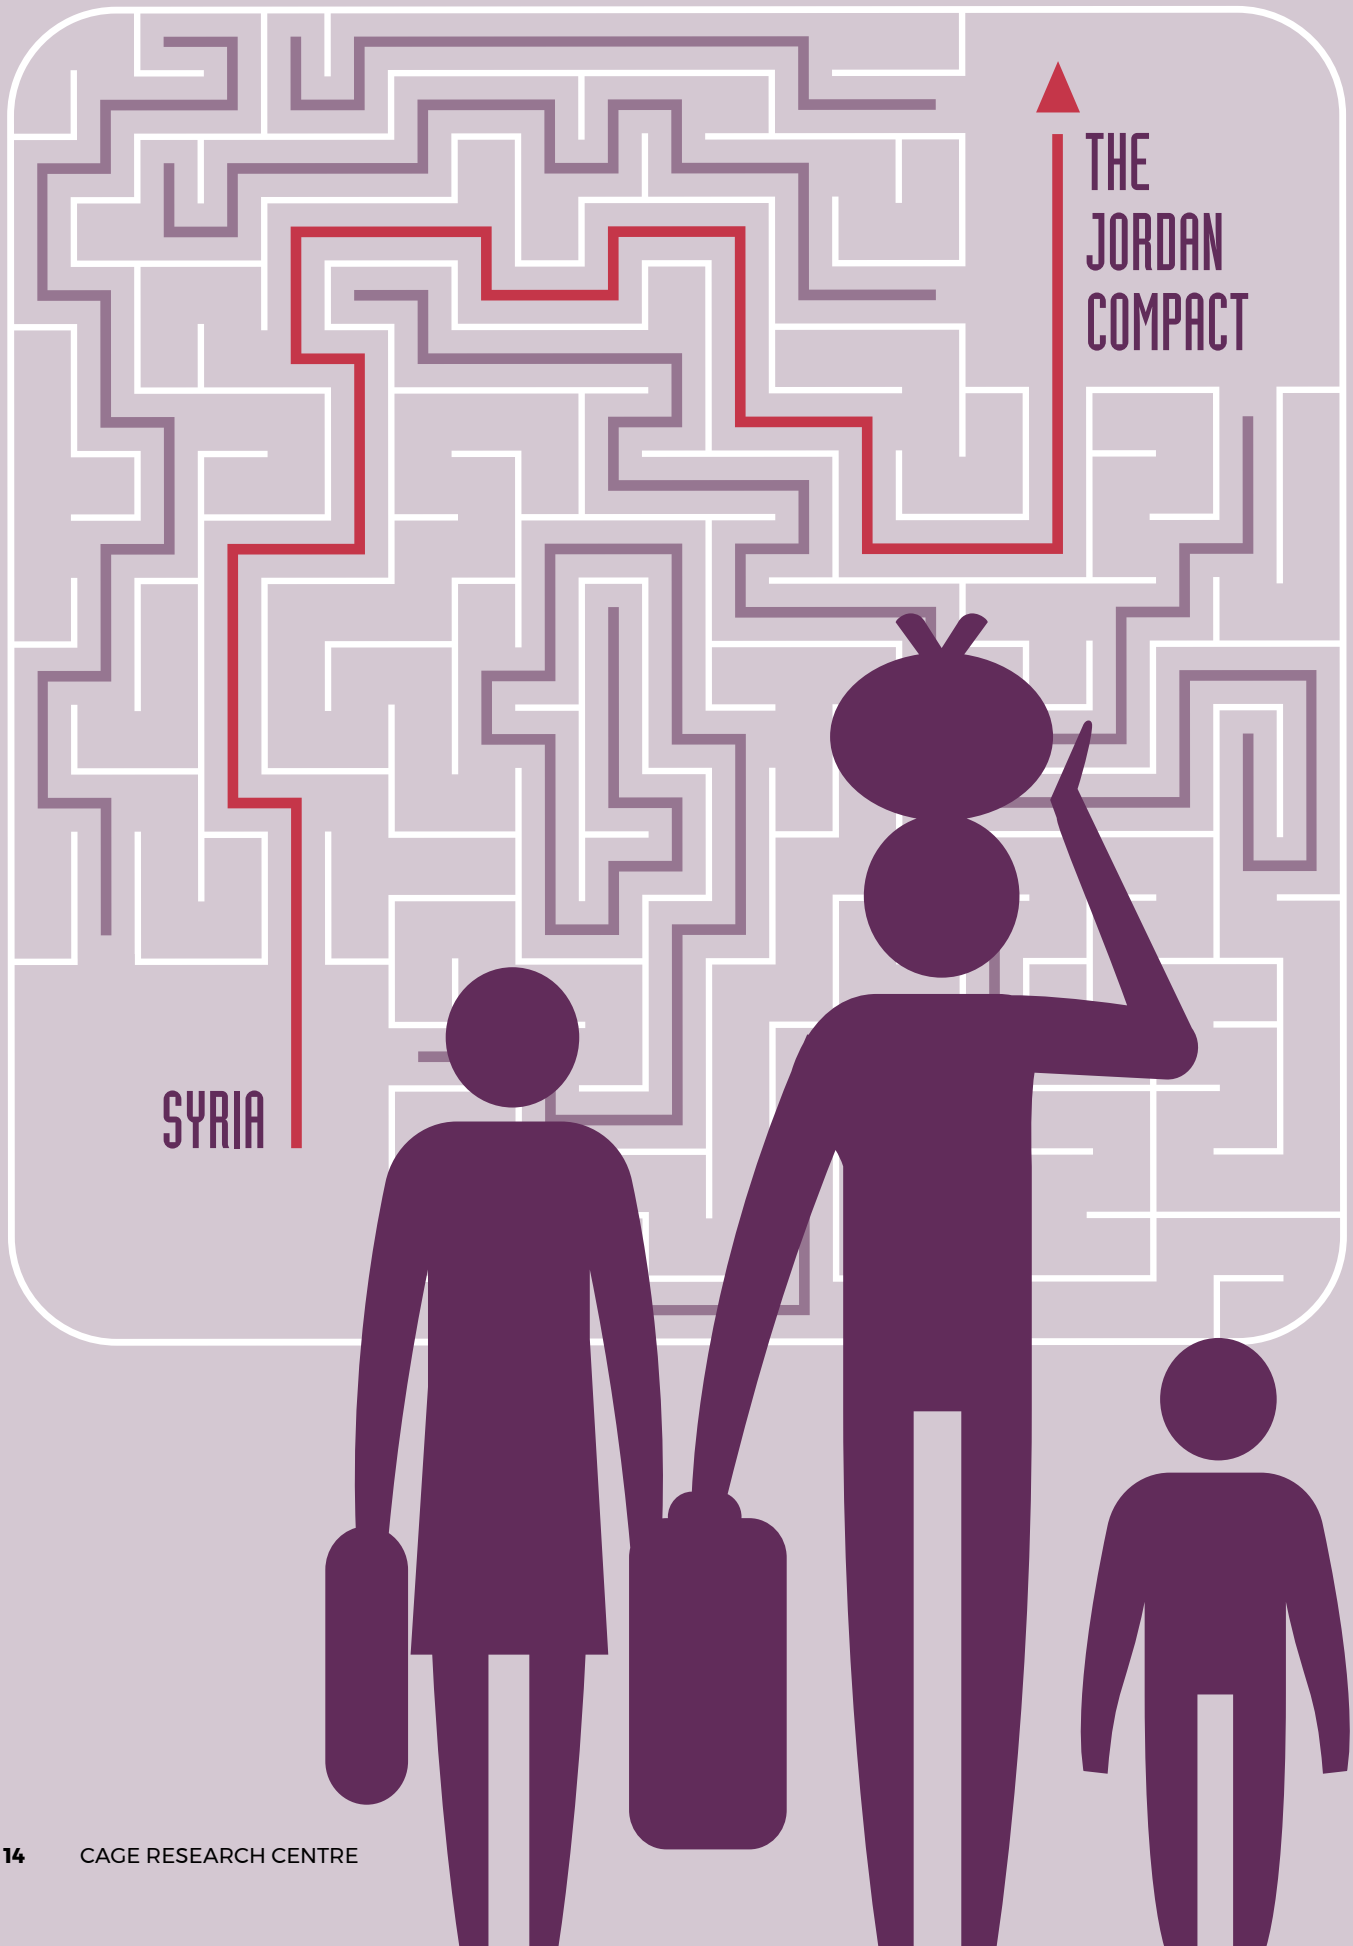
THE
JORDAN
COMPACT
SYRIA

Putting people's welfare first: A new experimental approach to help Syrian refugees find jobs in Jordan

By Stefano Caria

Randomised control trials (RCTs) are one of the most widely used methods for policy evaluation. However, RCTs have recently been *criticised* for sacrificing the welfare of study participants in order to produce scientific knowledge. We propose a new treatment-assignment algorithm that balances these different objectives and show how it can be implemented in the context of the Syrian refugee crisis.

Conventional RCTs are designed to learn about the impact of one or more policies in a transparent and precise way. In the standard protocol, the designer randomly assigns fixed shares of study participants to a treatment and a control group. Individuals in the treatment group are offered the policy, while individuals in the control group are not (or are only offered a weaker version of the policy). If the designer is interested in more than one policy, they will offer each policy to a fixed proportion of people in the study.

This procedure is ideal for learning about policy effectiveness, since it enables researchers to precisely compare the outcomes of a group of people that receives the policy and an equivalent group of people that does not. However, it is *not* designed to improve as much as possible the welfare of the people that take part in the experiment. For example, it fails to use the information on the effectiveness of different treatments that may become available over the course of the study.

Participant welfare is at the heart of current debate on the ethics of RCTs. As the scale, scope, and stakes of experiments grow — with RCTs being used to tackle some of the thorniest social issues, from COVID-19 containment to climate change mitigation (*Muralidharan and Niehaus, 2017* ; *Haushofer and Metcalf, 2020*) — and as *more and more people worldwide take part in experiments*, neglecting participant welfare has become very hard to justify.

A new approach: The Tempered Thompson Algorithm

In our paper (Caria et al., 2021), we propose a treatment-assignment algorithm that allows researchers to balance the objectives of learning about treatment effectiveness and maximising participant welfare. The algorithm generalises two well-known treatment assignment protocols. First, the *Thompson (1933)* protocol, whereby one assigns increasing shares of individuals to treatments that prove to be effective. Second, the conventional RCT protocol, where assignment shares are fixed. This guarantees that, over time, more people are assigned to promising treatments, while no treatment is ever dropped out from the experiment.

Our algorithm has two key features. First, it is adaptive. This means that, over time, it changes treatment assignment probabilities for new participants by using information on the effectiveness of different treatments on the existing participants. New participants are therefore more likely to be assigned to a treatment that will benefit them. Second, it is targeted, in the sense that it recognises that different groups of people may benefit differently from the same intervention. It uses group-specific information about treatment effectiveness to help assign participants to treatments most likely to improve their welfare.

Testing the algorithm: Employment policies for Syrian refugees in Jordan

We implement this methodology in a field experiment designed to help Syrian refugees in Jordan find employment. Jordan has been at the forefront of the Syrian refugee crisis and, since the start of the conflict, has received close to 700,000 refugees: one tenth of its original population (*UNHCR, 2020*).

Jordan has also recently implemented a bold set of reforms, known as the Jordan Compact, that substantially expanded access to formal employment for refugees. Legal restrictions to refugee employment are common in both developed and developing countries. If successful, the Jordan Compact could set an example for similar policy reforms around the world.

In collaboration with the International Rescue Committee, a non-governmental organisation, we test whether Syrian refugees need additional support in order to make the most of the new employment opportunities that have opened up for them. In particular, we study the impacts of offering (i) a small, unconditional cash transfer; (ii) information provision to increase the ability to signal skills to employers; and (iii) a behavioural nudge to strengthen job-search motivation. These interventions tackle a number of job-search barriers that have been identified in the previous literature: the cost of job search, the inability to show one's skills to employers, and the need to keep a strong motivation. We use the Tempered Thompson Algorithm to simultaneously learn about the impacts of these policies and maximise the employment rates of the refugees who take part in our study, as measured in our first follow-up interview, six weeks after treatment.

We find that the cash intervention has large and

Figure 1: The (simulated) impacts of the Tempered Thompson Algorithm

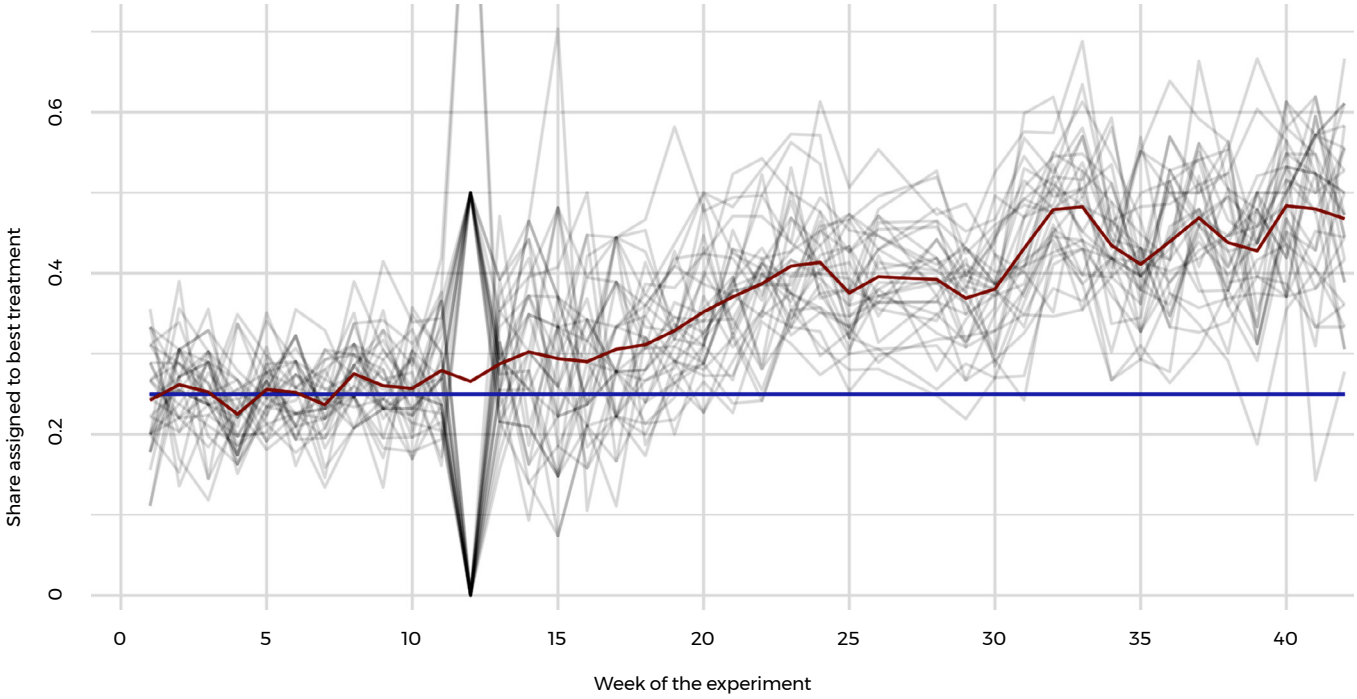

This plot shows simulated shares assigned to the best treatment within each stratum from 32 simulations (red), contrasted with the corresponding shares from a pure RCT (blue). The grey lines show each of the 32 simulation trajectories.

significant impacts on refugee employment and earnings, two and four months after treatment. While employment rates remain stubbornly low in the control group, the cash grant raises employment by 70% and earnings by 65%. These are sizable impacts compared to those documented in recent literature on active labour market policies and are larger than the impacts of the other two interventions, which are not significantly different from zero. However, impacts six weeks after treatment are close to zero for all interventions.

We also quantify the welfare gains from using our algorithm. First, we find that targeting increases employment rates by about 20%. However, since six-week employment effects are close to zero, there are almost no gains from the adaptive part of the algorithm. Second, we show that, if we had set two-month employment as the objective, a measure that responded strongly to the cash intervention, the algorithm would have doubled the employment gains of a standard RCT.

Our results show that the Tempered Thompson Algorithm can successfully measure policy effectiveness whilst also generating substantial welfare gains. However, choosing the right objective is key. The results also show that lifting legal barriers may not be sufficient to fully boost refugee employment in Jordan if this does not come with some additional financial support to help refugees search for employment. ◄

About the author

Stefano Caria is Associate Professor of Economics at the University of Warwick and a CAGE Associate

Publication details

This article is based on Caria, S., Gordon, G., Kasy, M., Quinn, S., Shami, S. and Teytelboym, A. (2021). *An Adaptive Targeted Field Experiment: Job Search Assistance for Refugees in Jordan.* *CAGE Working Papers* (No. 547)

Further reading

Duflo, E. and Banerjee A. (2017). *Handbook of Field Experiments*. Elsevier.

Muralidharan, K. and Niehaus, P. (2017). *Experimentation at Scale*. *Journal of Economic Perspectives* 31(4), 103–24.

Haushofer, J. and Metcalf, C. J. E. (2020). *Which Interventions Work Best in a Pandemic?* *Science* *368*(6495), 1063–1065.

Thompson, W. R. (1933). *On the Likelihood that One Unknown Probability Exceeds Another in View of the Evidence of Two Samples*. *Biometrika* *25*(3/4), 285–294.

UNHCR (2020, March). Jordan: Statistics for Registered Syrians in Jordan *https://data2.unhcr.org/en/documents/details/75098*

Essay competition winner: Julia Tattersall

Summary of Mark Harrison's: *Contracting for counterintelligence: The KGB and Soviet informers of the 1960s and 1970s.*

Cage working paper (no. 408)

Mass surveillance was often employed by communist regimes to empower the state over society. A large aspect of this was human intelligence — a powerful weapon in preventing action against authority. Comprised of a network of informers, each operating undercover, human intelligence was designed to spy on potential threats to society. Any precursors to the disruption of political order would be reported to the KGB. In this way, the state was able to pre-empt adversary action against the regime.

Whilst the collaboration and assignments of individual informers were kept secret, the existence of the network was not unknown. Although not the primary objective, widespread suspicion was of benefit to the KGB. Fear of mistakenly sharing hostile thoughts with an informer led to less hostile discourse in general. However, a low-trust society was not the intended outcome of KGB surveillance. The system relied on trust, which facilitated the process of collecting information. In fact, an informer's principal mission was to befriend members of society with the intent of later breaking their trust in order to supply knowledge to the KGB.

Throughout his essay, Harrison attempts to shed light on the specifics of this underground network, focusing on the recruitment process and the subsequent relationship between informer and handler. Drawing from a dataset of 21 personal narratives found in the archive of the KGB of Soviet Lithuania, Harrison is able to build a picture of the informers' realities.

Within the dataset, Harrison naturally found trends in the agents' characteristics. Besides the inevitable preference for sociability, hostile political records were also deemed desirable, with two-thirds admitting to a corrupt past, which often referred to participation in non-communist activity. There were more extreme examples too, like that of agent 'Neman', who had previously collaborated with Germany during Occupation.

KGB databases from the 1970s termed such propensities as negative, but the narratives show that these were in fact highly sought after. A record of political misconduct favoured both parties. For the state, a history of hostility meant connections to the political underworld, which would undoubtedly prove an advantage. For the agents, 'Neman' included, KGB collaboration was a means of redemption: an opportunity to cancel out one's past.

Whilst some citizens offered their services to the KGB, the majority were not volunteers. Instead the identification of candidates was strategic, with the KGB first determining a need and then searching for a suitable citizen to meet it. This investigation process could take months, with recruitment reserved only for those who succeeded the lengthy character evaluation.

The contract of recruitment featured a binding agreement of cooperation with the security police signed by the newly hired informer. It symbolised the written half of a two-sided commitment, with the handler's half being verbally implied. The agent's codename was incorporated into the contract, but there was neither mention of specific tasks nor duration of employment. The incompleteness of the agreement placed the informer in a position of ignorance, putting even more pressure on the search and due diligence required before recruitment. The KGB needed a subject that was committed enough to carry out tasks that were not specified in the initial agreement.

Consent was another concern, with the KGB itself acknowledging that recruitment was sometimes coerced. Whilst consent allows the candidate to refuse the contract with no adverse consequences, coercion involves an explicit threat to impose costs on those who fail to comply. The line dividing these becomes blurred when retribution is expected but never voiced. Evidence shows that there was no specific penalty for refusal, yet as many candidates had compromised records it is unsurprising that they felt obliged to comply. With this kept in mind, consented recruitment by the KGB's standards was high. However, the prolonged scrutiny of candidates that took place prior to recruitment equally explained the low incidence of coercion. The covert stage of preliminary investigation allowed for the selection of recruits who were already aligned with the KGB's ideologies or who could be aligned with little pressure. Certainly, this was the case with 'Neman'.

Yet Harrison discovered that regardless of the manner of recruitment, shirking was often a problem. When observing the performance of its agents, the KGB would pay close attention to the quality of the informer's cooperation, considering both the willingness of the agent on recruitment and the willingness of his subsequent performance. ▶

The former referred to the motives behind recruitment — whether it was wholehearted or forced. The latter referred to the agent's compliance with instructions of a task and his results thereafter. For unproductive informers, intervention was required. Due to the costs involved with searching for recruits, the KGB was strongly incentivised to persist with inadequate agents and invest in their abilities. The handlers had two instruments at their disposal: trust and re-education.

On one hand, poor performance was perhaps the result of the informer's qualities not being fully matched to the task at hand. In this instance, the handler would employ either *privitie* or *vospitanie*. *Privitie* referred to training the agent in the tradecraft of working undercover, whilst *vospitanie* referred to moral and political re-education, aiming to reform the agent's character and ideological attitudes. Among the 21 narratives, Harrison found that given the necessary *privitie*, many agents still did not perform to a satisfactory level. 'Nevskii' is an example of this, whereby his hard work was offset by his somewhat repellent attitude. Evidently, the only barrier to his success was his persona, which could be overcome through re-education in *vospitanie* in order to design an outward self that would attract the targets of the KGB.

On the other hand, it was perhaps the trust in the agent–handler relationship that was the problem. Non-compliance to tasks was considered a breach of trust that needed to be repaired. Often, the informer was inhibited from full cooperation by some hidden moral or political reservation. This was the case for 'Ruta', an agent who emerged from his depression only after trusting the KGB with his most inner thoughts. Once he had rebuilt the relationship with his handler, he was no longer lonely and was entrusted with more responsibility as an agent. Here, trust was the key that unlocked the informer's ability to perform his role. Yet the cause for non-compliance was not always on the part of the agent. Sometimes it was the KGB's fault. One handler trusted his agent, 'Stanislav', too much, leading him to withhold information about his colleague's criminal activities. The remedy for such was tighter discipline. However, the opposite could have equally been problematic, whereby not enough trust was causing the agent to become demotivated. Hence, although the trust was two-sided, the two sides were not equal. The informer's trust in the handler had to be unconditional, whilst the handler had to strike a balance – neither too much, nor too little – making use of surveillance in order to 'trust, but verify'.

Handler intervention was often met with a promise to protect the informer through affirmation of status — a powerful incentive in a secretive, low-trust society obsessed with personal record and rank. For exceptional service, the agents were compensated with tangible rewards that signalled gratitude and acknowledgement. In turn agent morale improved, increasing motivation and effort.

> **"The KGB used deception in the initial stages of recruitment, investigating potential candidates from under cover."**

To conclude, having drawn from 21 agent narratives, Harrison is able to piece together a better understanding of the KGB regime of Soviet Lithuania. His essay reveals the rigour involved in the evaluation of candidates that took place prior to recruitment and the post-contract intervention that was used to improve agent performance. The necessity of these lengthy processes could largely be blamed on the incompleteness of the recruitment contract, as it placed the premium on loyalty. In fact, Harrison found that the recruitment process lacked a few things, including the handler's written commitment to the agreement and the freedom of refusal exempt from any threat. Following recruitment, agent analysis was key to ensure performance was satisfactory. In the case of non-compliance or poor service, the handling officer needed to either invest time in re-educating the agent or in strengthening the trust in their relationship. Re-education typically involved the training and character development needed to successfully operate as an undercover agent. Trust, however, was at the crux of the relationship between informer and handling officer, and therefore needed to be monitored closely. Whilst an agent's trust needed to be unconditional, handlers were advised to 'trust, but verify'. No agent was so trusted that his work would not be checked. Rather, a balance had to be found in order to motivate the agents and encourage productivity.

In essence, trust and deception were two sides of the informer's coin. The KGB used deception in the initial stages of recruitment, investigating potential candidates from under cover. Once hired, the agents were then taught the same skills to spy on fellow citizens, whom they would later go on to betray. ◀

About the author

Julia Tattersall is a fourth year undergraduate student in French and Economics at the University of Warwick. She is an executive member of two societies at Warwick: UN Women UK Warwick and Warwick Women's Careers Society.

Parting Shot

Forget 'progress studies', we need a new science of bureaucracy

By Mirko Draca, Director, CAGE

Bureaucracy is something that we seem to accept: it's time to make a conscious effort to understand and control it.

Donald Trump claimed in 2017 that Americans built the Golden Gate Bridge in four years, and the Hoover Dam in five. A completion time of that order for an infrastructure project seems unthinkable today. More recently, the bureaucracy of the EU has come under intense criticism as the time taken to approve COVID-19 vaccinations has left EU countries significantly behind the curve in inoculating their citizens. I regularly hear academic colleagues complain about bureaucratic forms and compulsory training courses. It's been claimed that UK universities have more administrators than faculty staff (Spicer, 2017), and these types of complaints about high administrative overhead seem very common across industries. ▶

In 2019, Stripe CEO Patrick Collison and George Mason University (GMU) Professor Tyler Cowen made a call for a new science of 'progress studies' focused on both analysing the historical rise of living standards and working out ways to actively engineer increases in standards across many domains. By contrast, I think the concept of 'bureaucracy studies' is a better bet for engineering progress, because it targets a pervasive problem in social organisation.

So, what would 'bureaucracy studies' involve? I think the economics of information would need to be at the core of it. The concepts of asymmetric information and moral hazard, especially the 'principal–agent problem', are indispensable tools here. These ideas are mature within the economics discipline — going back at least 50 years — but there is little awareness of them among the general public. 'Asymmetric information' describes a situation where one party to an interaction has more information than the other party. The classic example comes from the used car market and shows how lack of trust in the quality of the cars being offered can lead to a spiral of lower prices and falling product quality (Akerlof, 1970). The resulting 'adverse selection' can be remedied in part by third-party certification schemes, but anyone who has struggled with a cheap, secondhand car knows that this policy has its limits. The 'principal–agent problem' occurs when actions are delegated out to a person who doesn't necessarily have the same interests as the one doing the delegating. A canonical example is politicians: we (the 'principals') delegate social decision-making power to them (the 'agents'), but they maximise according to their own interests, which don't neatly align with that of society as a whole. Theoretical research on the economics of information boomed in economics in the 1970s and 1980s but I would argue that we have not incorporated the insights of this work into the design of modern organisations.

A second element of bureaucracy studies is probably more familiar to the general public: behavioural economics. I think of behavioural economics as giving us a set of tools that allow us to pick apart systematic and frequent departures from conventional 'rational' decision-making. Concepts such as the sunk cost fallacy, framing effects, status quo bias and herd behaviour have been popularised in books by Kahneman (2012), and Thaler and Sunstein (2012). Behavioural economics has established itself pretty well within the public policy industrial complex through groups such as the UK Cabinet Office Behavioural Insights (aka 'nudge') unit, but the complementarities between this and information economics for understanding bureaucracy have not been harnessed.

The final element of bureaucracy studies would be rich empirics. The modern economy provides abundant material for the study of organisations and bureaucracy. A great example of the type of empirical work that is possible is Aral, Brynjolfsson and Van Alstyne's (2012) study of worker productivity in an executive recruiting firm, where they use data on more than 125,000 email messages and 1,300 well defined projects to study multi-tasking. The intra-organisational networks and patterns of communication that are implicitly revealed in email data have the potential to tell us a lot about bureaucracy. The pandemic has also created a new source of data that is sure to be informative, namely information about the frequency and structure of (virtual) meetings.

A big question for bureaucracy studies is whether bureaucracy has gotten out of hand. Trump's mention of the Golden Gate Bridge and Hoover Dam has been scrutinised by the *Washington Post*, who point out the long lead times for approval of these building projects. There's also well-argued pushback against claims of overblown UK university bureaucracy (Greatrix, 2017). Culturally, Franz Kafka was writing about bureaucracy in the early 20th century, well before 'cringe comedy' masterworks like *The Office* and *W1A*. The fundamental problems of bureaucracy have been around for a long time. I suspect that as society *has* become more complex, the situation has gotten worse: but that is a proposition that needs testing. ◀

References

Akerlof, G (1970) 'The Market for "Lemons": Quality Uncertainty and the Market Mechanism'. *Quarterly Journal of Economics*. 84(3): 488–500.

Aral, S., Brynjolfsson, E., and Van Alstyne, M. (2012) 'Information, Technology and Information Worker Productivity'. *Information Systems Research*. 23(3): 849–867.

Collison, P., and Cowen, T. (2019) *'We Need a New Science of Progress: Society Needs to Get Better at Knowing How to Get Better'*. *The Atlantic Monthly*.

Greatrix, P. (2017) 'University administrators — *'"lovely and well-meaning" but still below stairs'*. WonkHE. 23 August 2017.

Kahneman, D. (2012) *Thinking Fast and Slow*. Penguin: London.

Spicer, A. (2017) *'Universities are broke: so let's cut the pointless admin and get back to teaching'*. *The Guardian*, 21 August 2017.

Thaler, R., and Sustein, C. (2012) *Nudge: Improving Decisions About Health, Wealth and Happiness*. Penguin: London.

Ye Hee Lee, M. (2017) *'Fact Checker: Trump's Claim that American's 'built the Golden Gate Bridge in Four Years and the Hoover Dam in Five.'* *The Washington Post*. 14 June 2017.

About CAGE

CAGE is a research centre based in the Department of Economics at the University of Warwick. We conduct independent policy-driven research informed by history, culture and behaviour. Our aim is to move beyond traditional measures of economic success to consider broader influences on global prosperity; from cultural and behavioural attitudes to voter preferences and political institutions. We analyse historical and contemporary data to draw out lessons for modern policy. CAGE is supported by the Economic and Social Research Council (ESRC).

Who we are

We are a small team of experts seeking to apply economic principles to ask new and innovative questions of data. We want to know how and why economies are successful, and the ways in which history, culture and behaviour shape the global economy (and vice versa).

We produce robust evidence to inform policymakers and journalists and influence both policy and debate. Our core team consists of nine Research Theme Leaders and Deputy Leaders who work across four Research Themes. We also have a number of internal and external associates who contribute to our research.

Publications and events

Our academic working paper series showcases the research of our team and our associates. We also publish a bi-annual magazine, Advantage, which highlights the best of our policy-driven research for an informed non-academic audience. Our policy briefings and themed policy reports seek to draw out policy recommendations and findings to inform current debate.

Our event programme focuses on driving impact from our research and we conduct regular briefings in London and across the UK. We also host a policy conference bringing together academics and policy specialists to discuss contemporary economic and political challenges. We support young talent through our annual summer school.

CAGE research uses economic analysis to address real-world policy issues.

advantage

CAGE conducts independent policy-driven research informed by history, culture and behaviour. We produce a wide range of publications which are available to download from the centre's website: *warwick.ac.uk/fac/soc/economics/research/centres/cage/publications*

Join in the discussion by attending our free events: *warwick.ac.uk/fac/soc/economics/research/centres/cage/events*

Stay informed

Join our mailing list to receive our bi-annual *Advantage* magazine and our newsletter which features highlights of our research and impact activities.

Visit us at **warwick.ac.uk/cage**

CAGE Research Centre
Department of Economics
University of Warwick
Coventry CV4 7AL
United Kingdom

warwick.ac.uk/cage

@cage_warwick

cage.centre@warwick.ac.uk

UKRI Economic and Social Research Council

CAGE

WARWICK
THE UNIVERSITY OF WARWICK