

advantage

The Magazine of CAGE Research Centre / Autumn 2020 / Issue 10 / ISSN 2397-8309



Public Policy & Data

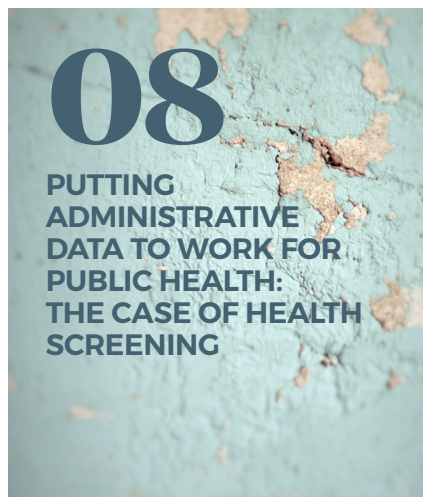
NEW RESEARCH THEME SPECIAL

In this issue ..

04
A TALE OF TWO CRISES:
THE COVID-19
PANDEMIC VS THE
GREAT RECESSION
AND THEIR IMPACTS
ON LABOUR MARKETS



08
PUTTING
ADMINISTRATIVE
DATA TO WORK FOR
PUBLIC HEALTH:
THE CASE OF HEALTH
SCREENING



11
SHOULD THE UK
LOWER ITS VAT
THRESHOLD?



14
EXPOSING THE GLASS
CEILING: HIRING
PRACTICES AND
GENDER EQUALITY
IN THE UK LABOUR
MARKET



18
WHO POLICES THE
POLICE? BODY-
WORN CAMERAS
DO WORK
AFTER ALL



21
HOW EARLY
ADOPTERS
SPREAD TWITTER
ACROSS THE
UNITED STATES



24
PARTING SHOT



26
ABOUT CAGE



27
CONTACTS



Welcome to this special issue of *Advantage Magazine*...

This autumn we are excited to introduce CAGE's new Research Theme, entitled 'Public Policy and Data'. The theme investigates how best to design policies to tackle the most pressing challenges developed economies face. It features researchers across multiple disciplines who go out of their way to collect and leverage novel data sources and employ a wide range of empirical methods, structural modelling, policy evaluation and machine learning. The outcome, we think, yields exciting insights at the frontier of policy research.

As the world continues to grapple with COVID-19 and its associated economic crisis, our researchers are applying cutting-edge analytical methods to shed light on the road ahead:

Mirko Draca, Emma Duchini, Roland Rathelot and Giulia Vattuone investigate possible policy responses to the labour market shock of the pandemic by comparing it to that of the Great Recession. They stress that while retention schemes are needed to smooth the shock of COVID-19, there is a thin line between smoothing and preventing workers from reallocating to jobs that hire.

The pandemic has emphasised the importance of designing effective health screening policies. Ludovica Gazze presents learnings from a similar context and shows that travel costs reduce childhood screening for lead poisoning in the United States and are potentially most harmful for those children who have the highest risk of lead exposure.

Changes to tax policy may be key in supporting recovery in the next few years. Ben Lockwood looks at VAT and the effects its design has on growth incentives for small firms.

Some public policy concerns — like gender inequality, ineffective governance and enforcement, and barriers to innovation — are sure to outlast the pandemic. Our researchers offer new insights:

Emma Duchini and Stefania Simion leverage text analysis and insightful data on job postings to explore how hiring practices correlate with observed measures of gender inequality. Pedro Souza employs a state-of-the-art field experiment bringing together researchers and government agencies to shed light on the effectiveness of body worn cameras as an intervention to reduce police brutality. Karsten Müller and Carlo Schwarz explore the early take-up patterns of Twitter, with findings that can inform the design and roll-out of new technologies.

As the world around us evolves, our researchers are ready to inform high-stake policy decisions by doing what they do best: turning to new sources of data and rigorous methods.

Roland Rathelot, Lucie Gadenne and Ludovica Gazze



CAGE Director: **Mirko Draca**

CAGE Research Director: **Bishnupriya Gupta**

CAGE Impact Director: **Arun Advani**

Editor: **Stephanie Seavers**

Contributors

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

Emma Duchini is Research Fellow at the University of Warwick and a CAGE Associate.

Ludovica Gazze is Assistant Professor of Economics at the University of Warwick and a CAGE Deputy Theme Leader.

Ben Lockwood is Professor of Economics at the University of Warwick and a CAGE Associate.

Karsten Müller is a Postdoctoral Research Associate at the Juli-Rabinowitz Center for Public Policy and Finance at Princeton University and a CAGE Associate.

Roland Rathelot is Associate Professor of Economics at the University of Warwick and a CAGE Theme Leader.

Carlo Schwarz is Assistant Professor of Economics at Bocconi University and a CAGE Associate.

Pedro Souza is Assistant Professor of Economics at the University of Warwick and a CAGE Associate.

Giulia Vattuone is a PhD student at the University of Warwick.

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

Articles in this issue of *Advantage* reflect the opinions of the authors, not of the Centre.

The content of the articles (illustrations excepted) is not copyrighted and may be reproduced freely with an appropriate attribution of source. Please provide us with a copy of any material that you reproduce.


CAGE Research Centre


Department of Economics, University of Warwick, Coventry CV4 7AL, UK

A tale of two crises: The COVID-19 pandemic vs the Great Recession and their impacts on labour markets

By Mirko Draca, Emma Duchini, Roland Rathelot
and Giulia Vattuone

The COVID-19 pandemic has upended the path of the business cycle creating a major unexpected negative economic shock. The new COVID-19 recession is likely to be very different to those that have gone before. As we attempt to devise the best policy response to the economic consequences of the pandemic, it's useful to compare current dynamics to the early phases of the most recent severe downturn – the Great Recession.






After the financial crisis that started in 2007 in the US, the UK experienced a decrease in its Gross Domestic Product (GDP) throughout 2008. Following a decline in labour demand and mass layoffs, the employment rate decreased sharply by 3 percentage points (from 73% to 70%) and remained at that level until 2012, long after GDP started to grow again. The unemployment rate jumped from 5% to 8% and only started to fall in 2013. The impact of the Great Recession was uneven across occupations and industries, affecting workers in manufacturing and 'routine' manual jobs more than others. This continued the trends of the previous major recessions in the early 1980s and 1990s.

The COVID-related economic crisis has a very different source. Starting as a response to a health pandemic, its impact on the labour market has unfolded both as a supply- and a demand-side shock. Workers are hesitant to go to work (or go out at all), affecting production in workplaces that rely on workers being present. At the same time, household consumption has gone down, and some sectors (live entertainment, hospitality, schools) have been mandated to stop their activities completely at various points.

The sources of the two crises are different, and by their nature economic downturns have uneven impacts on activity, affecting some parts of the economy disproportionately more than others. This pattern of impacts is likely to have major implications for the long-term structure of the economy.

**Figure 1: Changes in the labour market by occupation
Great Recession vs COVID Pandemic**



Source: BGT 2019-2020, ONS Vacancy Survey and LFS 2008-2009. Note: Each light (dark) red bar represents the percentage change of the stock of vacancies in the second quarter of 2009 (2020) relative to the second quarter of 2008 (2019) in a given Standard Occupation Classification (SOC) 2010 occupation. The occupational breakdown of vacancies for the period 2008-9 is computed using the LFS shares of newly employed workers in each occupation. A newly employed worker is defined as having less than three months job tenure

Comparing impacts on sectors and occupations.

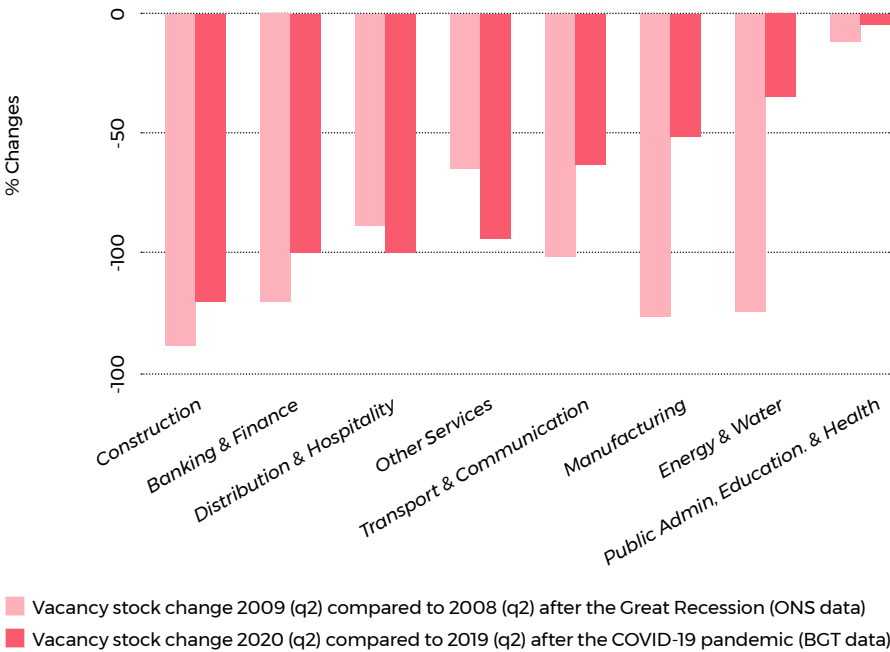
To investigate how much employment in different sectors and occupations has been impacted by the two crises, we must use different data sources. For the Great Recession, we combine the ONS Vacancy Survey with the Labour Force Survey (LFS) to track the quarterly evolution of labour demand in each sector and occupation.¹

For the current crisis, we use data on the number of job vacancies posted across many different online job boards gathered by the company Burning Glass Technologies (BGT). The main advantage of using BGT for the current period is that it offers a granular time coverage of labour demand dynamics and provides both an industry and occupational breakdown of vacancies. We look at the evolution of the number of vacancies between the second

quarter of 2008 and the second quarter of 2009, and we compare it to the evolution of the number of job ads between the second quarter of 2019 and the second quarter of 2020.

Figure 1 shows that relative to the Great Recession, the shock due to COVID-19 has had a different pattern of impacts across occupations. Administrative and secretarial jobs, managers and directors, skilled trades, plant and machine operators, and elementary occupations suffered similarly from both shocks. However, sales and customer services, and associate professionals and technicians – in short, jobs requiring mostly face-to-face interactions – were much more affected by the pandemic than by the Great Recession. Conversely, professional and caring and leisure jobs (probably driven by healthcare occupations) were much less affected by the COVID crisis than the 2008 recession.

**Figure 2: Changes in the labour market by industry
Great Recession vs COVID Pandemic**



Source: BGT 2019-2020, ONS Vacancy Survey and LPS 2008-2009. Note: Each light (dark) red bar represents change of the stock of vacancies in the second quarter of 2008 (2019) in a given sector (Standard Industrial Classification 2007).

Figure 2 shows the results of a similar exercise at the sector level. There are stark differences across the two crises. The main one is that while the manufacturing and transport and communication sectors were the main victims of the Great Recession, they have been relatively spared by the pandemic.² The same goes, to some extent, for the public administration, education and health sectors. Conversely, distribution and hospitality as well as other service activities (e.g. hairdressing and physical well-being activities) have been hit harder by the COVID crisis than the Great Recession. Banking and finance and construction have been impacted by both crises.

Towards a recovery?

The labour market took four years to recover to pre-crisis levels after the Great Recession. The external public-health origins of the COVID crisis gave hope to policymakers and pundits

that its impact on economic activity would be short-lived. However, the persistence of the pandemic has led the government to keep in place regulatory restrictions that aim to avoid unnecessary physical interactions. Consequently, some of the most impacted sectors have not been able to operate normally and fully restore their activity. Moreover, the domestic and international demand for services and goods produced in the UK have not much increased, slowing down the recovery. The combination of these two forces now has the potential to turn the temporary halt in activity mandated by lockdown policies into a more durable economic crisis.

Figure 3 shows how the number of vacancies has evolved across occupations since February 2020. Note here that we are comparing year-on-year levels in vacancies, for example the level of vacancies in an occupation during February 2020

“The record of economic policy in dealing with previous mass displacements of manufacturing workers in earlier recessions has been a very mixed one. Hence, a focus on reallocation needs to be a priority for analysis as we enter 2021.”

versus levels in the same occupation in February 2019.

The picture is a gloomy one overall. There has been no spectacular comeback in vacancies for any occupation. Patterns are similar across occupations: vacancy rates started falling in mid-March and reached their nadir in May.

Since then the number of new job ads has been increasing, but at a very slow rate. The evolution has been mostly parallel: the occupations that had been most impacted at the start of the crisis are still the ones that employers are the least eager to hire in late July. Still, some occupations are recovering better than others. Vacancies for plant and machine operative occupations have gone from -60% in late April to -40% in late July (compared to the predicted number of vacancies based on the 2019 data). Associate professionals and technicians have gone from -60% to -50%. Conversely, caring,

leisure and other services, which did not fall as dramatically in April (-20%), do not seem to have recovered by the end of July.

In the long run... we're reallocated.

The Great Recession (and previous crises) has taught that economies recover, in the sense that production, employment and unemployment rates come back to pre-crisis levels. However, this recovery can take years and is not necessarily neutral with respect to sectors and occupations. Overall, routine occupations and some manufacturing subsectors have never caught up after the Great Recession, but these jobs have been replaced by others.

The post-COVID policy dilemma is how to smooth the shock while not slowing down too much the potential reallocations that might need to happen. **The Job Retention Scheme** (JRS – aka ‘furlough’) which began on 20 April and has recently been extended to 2021, is designed to aid

the continuity of economic activity after lockdown in March. However, continuity is not always possible and sometimes reallocation is necessary on the road to economic recovery. Simply put, workers who have lost their old jobs will need to find new ones. The challenge is that structural changes in the economy post-COVID will mean that there will need to be many cross-sectoral movements of workers.

For example, it is likely that there will be a long-term displacement of face-to-face retail jobs in favour of online commerce. The persistence of high levels of remote working will also have an impact on city-based commercial services such as dry cleaners and sandwich shops. Given the novelty of these phenomena compared to the Great Recession, the great unknown is how much of this job displacement there will be and how difficult it will be for workers to reallocate.

For policy makers, the challenge

is to find the thin line between supporting struggling activities to prevent the loss of jobs and human capital, and accommodating the necessary mutations of the economy to encourage the reallocation of workers to sectors and occupations that will need them. The record of economic policy in dealing with previous mass displacements of manufacturing workers in earlier recessions has been a very mixed one. Hence, a focus on reallocation needs to be a priority for analysis as we enter 2021. ◀

We thank the ESRC for funding our project ‘Rescuing a “Sick” Labour Market: Using Online Vacancy Data to Track COVID-19’s Economic Impact’. See the CAGE website for future briefings and working papers on the post-COVID labour market.

About the authors

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

Emma Duchini is Research Fellow at the University of Warwick and a CAGE Associate.

Roland Rathelot is Associate Professor of Economics at the University of Warwick and a CAGE Theme Leader.


Giulia Vattuone is a PhD Student at the University of Warwick.

Footnotes

1. As the ONS Vacancy Survey provides only the industrial breakdown of vacancies, we computed the occupational breakdown by multiplying total vacancies for the LFS shares of newly employed workers in each occupation.

2. Note that these are aggregate categories of industries. Therefore, these figures could mask heterogeneous effects across subsectors (e.g. passenger vs freight transport subsectors which are both included under transport and communication).

Figure 3: 2020 vs 2019 % change in trend by job occupation



Source: BGT 2019-2020. Note: The figure shows the percentage changes between the 2020 and 2019 trends in job postings across occupations. Each yearly trend for a given occupation is calculated using 42-day moving averages of job vacancies, normalised to the 12/02 of the corresponding year.

Putting administrative data to work for public health: The case of health screening

By Ludovica Gazze

Travel costs can reduce take-up of important government programmes, including healthcare screening. High-risk individuals appear disproportionately affected by travel costs, hindering timely detection of healthcare issues and treatment access.

Travel costs can reduce take-up of important government programmes, including healthcare screening. High-risk individuals appear disproportionately affected by travel costs, hindering timely detection of healthcare issues and treatment access.

Effective health care delivery often requires timely screening to detect issues at early stages, when consequences can be more easily mitigated. The importance of screening has become especially apparent in the current COVID-19 pandemic, with policy discussions focusing on the screening of both symptomatic and asymptomatic individuals. However, individual motives to seek or refuse screening may not align with public health guidelines. For example, when test supply shortages are a concern, it may be best to target resources to high-risk individuals. Yet, some low-risk people may insist on getting tested, perhaps because they have read about the gravity of the pandemic in the media. Conversely, some high-risk individuals may avoid getting tested if they feel that they could not access life-saving treatment, or could not afford the prescribed quarantine following a positive test result, or simply if they cannot take time off work to get to a screening location. To design effective screening policies and improve public health it is crucial to understand what determines an individual's decision to seek screening (Einav et al, forthcoming).

My research examines one potential barrier to screening – travel costs – in the context of screening for lead poisoning in the United States.

Do travel costs deter people from accessing screening or do they help target screening resources to the highest risk children? The answer to these questions can help policymakers make important choices about the location and targeting of health screening provision.

Barriers to screening: The case of lead poisoning in Illinois, the United States

Lead exposure has extremely severe consequences for children's cognitive and non-cognitive abilities and it imposes large costs on society. Nonetheless, sources of lead exposure are still pervasive in the United States: two-thirds of the Illinois housing stock, almost 3.6 million homes, was built prior to the residential lead paint ban in 1978 and may have lead paint. As children start moving autonomously around the house, usually between 9 and 24 months of age, they may ingest or inhale lead dust from deteriorating paint. However, getting the lead out of every home indiscriminately is very costly. Current programmes rely on children's visits to the doctor at ages 1 and 2 to identify 'poisonous' homes. This system potentially introduces a barrier if some families have difficulties in getting to the doctor's office at the right ages. And perhaps because of this barrier, over a third of high-risk children in Illinois, who are required by the state to be screened, do not receive testing.

Are families deterred from screening due to travel costs to the doctor? To answer this question, I assemble a large set of address-level administrative data detailing children's screening outcomes, travel costs, and lead exposure risk. First, I link almost

2 million birth records of all children born in Illinois between 2001 and 2014 to 2.9 million blood lead test records to construct children's screening histories. Second, I map children's addresses at birth and lead-screening providers' addresses to measure the distance a child must travel to get screening. Third, I link these individual-level data to address-level housing age to construct a measure of exposure risk at birth addresses, as housing age is a good predictor of both lead paint use and lead content in paint.

Travel costs decrease lead screening for both high-risk and low-risk children

Because health care providers open and close for reasons largely unrelated to lead screening policies, I can compare screening outcomes for children born in the same neighbourhood, or even at the same address, but in different years. These children face different travel costs, as some providers may have opened or closed in the meantime. Figure 1 shows that children who live closer to providers are more likely to be screened: a difference of 10 kilometres, approximately a 15-minute round trip by car, decreases screening by 9% relative to the mean screening rate of 46%.

But who are the children that forego screening at higher distances? Consider two children, one in an old house and one in an adjacent new house. There is a clinic 250 metres away, and both get screened. Years later, two new children move in; the clinic is closed, and the closest provider is now a kilometre away. Knowing that the new house is lead-free, only the parents in the old house take their child to get screened. Among the screened children in this example, the probability that a child lives in an old home increases with distance: it is 0.5 at 250 metres and 1 at one kilometre. This would suggest that travel costs help target screening resources to the highest risk children. My findings, as shown in Figure 2, do not support this hypothesis. Children who attend

“Lead exposure has extremely severe consequences for children’s cognitive and non-cognitive abilities and it imposes large costs on society.”

screening from farther distances do not have higher risk: if anything, they have lower blood lead levels.

What are the lessons for public policy?

So, travel costs do not improve targeting, and in fact decrease timely detection of lead poisoning. Can different policies do better? Figure 3 shows that travel subsidies, pay-for-performance incentives and targeted screening mandates can all increase detection.

Currently, only children living in zip codes defined as high-risk for lead poisoning are required to be screened. Based on this research, the state of Illinois decided to pause the extension of universal screening to the whole state to focus resources on removing the barriers still preventing these high-risk children from being screened. For example, in follow-up work we are using machine learning techniques to identify under-screened high-risk clusters, demonstrating the potential for coupling granular administrative data with frontier methods to improve public policy design.

The findings of this research, and the methods used, could provide significant lessons for the screening of COVID-19. As the UK strives to keep schools and workplaces open for the long term, understanding what motivates or hinders people to seek testing could be crucial to identifying and containing outbreaks of the virus. ◀

About the author

Ludovica Gaze is Assistant Professor of Economics at the University of Warwick and a CAGE Deputy Theme Leader.

Further reading

Einav, L., Finkelstein, A., Oostrom, T., Ostriker, A. J. & Williams, H. L. (forthcoming). Screening and Selection: The case of Mammograms. *American Economic Review*.
 Gaze, L. (2020). Hassles and Environmental Health Screenings: Evidence from Lead Tests in Illinois. *CAGE working papers* (no. 509).

Figure 1: Children who live closer to screening providers are more likely to be screened




Figure 2: Children who attend screening from further distances do not have higher risk





Figure 3: Travel subsidies, pay-for-performance incentives and targeted screening mandates can all increase lead poisoning detection



Should the UK lower its VAT threshold?

By Ben Lockwood

In the UK, businesses are required to pay VAT if their annual turnover is over £85,000. In 2018, the Government considered lowering this threshold to increase tax revenue. This move could be welcome in the wake of the economic shocks of Brexit and the COVID-19 pandemic. But would such a policy stifle the growth of small businesses?



Most countries use the value-added tax (VAT) as their primary indirect tax. It is standard to set a minimum registration threshold, usually based on annual turnover, below which businesses do not need to register for VAT. In the EU, a large majority of countries currently have a registration threshold while a few require all businesses to register for VAT. The UK also sets a registration threshold – the highest in Europe – at £85,000.

Until recently, policy in the UK has been to raise the VAT threshold roughly in line with inflation. However, after Brexit the UK Government investigated the possibility of lowering the threshold in order to raise more tax revenue. Following an in-depth investigation into the current design of the threshold by the Office for Tax Simplification, in the November 2018 Budget the VAT threshold was frozen at £85,000 until April 2022.

This outcome reflects a compromise between the need to raise revenue following the negative economic shocks of Brexit (and now the COVID-19 pandemic) on the one hand, and a concern that some

businesses were purposely keeping their turnover below the £85,000 threshold in order to avoid becoming eligible for VAT registration on the other. Such turnover restriction, or ‘bunching’ as it is sometimes called, clearly constrains the growth of very small businesses – many of them micro-businesses with fewer than ten employees – and thus affects the dynamism of the economy.

In our work, we use rich administrative data on UK corporations to investigate these important issues. We begin by documenting two stylised facts in the data: bunching and voluntary registration. We find substantial bunching below the registration threshold. Our precise definition of bunching is that a firm ‘bunches’ when it restricts its reported turnover in its corporate income tax (CIT) return below the VAT threshold to avoid having to register for VAT.

Figure 1 shows the extent of bunching for all firms in the UK with turnover between £40,000 below and £100,000 above the turnover threshold between financial years 2004/05 and 2014/15, about 3.5 million firms. We normalise turnover by subtracting the threshold value

from the turnover value so that the threshold is located at zero in any year. The blue line shows the actual distribution of firms around the threshold, and for clarity, the orange line shows the *counterfactual* distribution that we would observe if there were no bunching.

This figure shows quite clearly that some firms at least, restrict their reported turnover to avoid crossing the VAT threshold.

Next, we ask the question: do all or most firms do this? Our second finding suggests that this is not the case. Specifically, in any given year, a significant number of firms are registered for VAT even though their turnover is *below* the threshold in the current year. On average, over our sample period, 43% of firms below the threshold are in this position. Possibly, part of this may be due to rules of registration. In the UK, a business must register for VAT if its taxable turnover is likely to go over the threshold in the next 30 days, or if its taxable turnover in the previous 12 months was above the threshold. So, for example, a firm may register on the basis of the previous year’s turnover, and then its turnover may fall below the threshold in the current year.

However, there is considerable persistence in registration below the threshold. This is evident in Figure 2, which shows what happens to firms initially registered and below the threshold during fiscal year 2004/05. Almost half are still registered three years later, and over one third are still registered five years later. So, it is likely that registration below the threshold is a conscious decision by firms, rather than just due to inability to forecast turnover one year in advance, or inertia. This is known as ‘voluntary registration’. On the face of it, voluntary registration is quite a surprising finding, given that the costs of registering for VAT, and filing VAT annual or quarterly returns, are not trivial for micro-businesses.

Given that there is evidence of both bunching and voluntary

Figure 1: Evidence of bunching by small firms


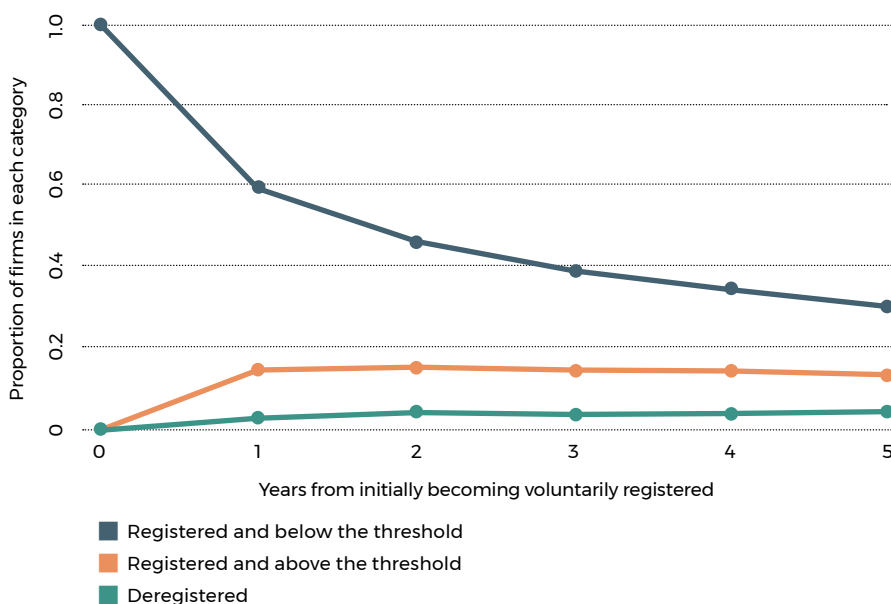


Figure 2: Evidence of voluntary registration by small firms



Around the threshold, firms are growing at an average rate of 8%, so the slowdown is quite substantial. However, statistical analysis shows also that on average, the acceleration after the threshold completely offsets the slowdown. We also find that the slowdown in growth for firms that are voluntarily registered below the threshold is much smaller, as we might expect. Finally, the same pattern of slowdown and acceleration is seen in the turnover of small businesses that are not incorporated i.e. sole traders. So, the policy take-away is that the threshold does not have any permanent effect on small firm growth, and there may be some scope for lowering the threshold in the UK. ◀

registration, one might then ask; why do some firms bunch, whereas other firms actually choose to register voluntarily, given the compliance costs of doing so? Basic economic analysis suggests three possible reasons. First, if the cost of inputs relative to sales is high, registration allows the small firm to claim back a considerable amount of input VAT. Second, business customers of a small firm can usually claim back any VAT charged, so a small firm that registers can pass on the burden of VAT charged on outputs to their customers in this case. Finally, if the market in which the small firm operates is less competitive, it is easier to pass on the burden of output VAT onto buyers even if they are final consumers, because prices are less likely to be undercut by other firms. In our empirical work, we find evidence that all three of these channels affect the probability of voluntary registration, and the amount of bunching, in the expected way.

Finally, we turn to the issue that is probably of most interest to the UK government: does the threshold affect small firm growth? In Figure 3, we plot the rate of growth of turnover against the distance of last year's turnover from the VAT threshold for that year.

We plot the average turnover growth rate of firms whose normalised lagged turnover (i.e. last year's turnover minus the threshold) is in a £1000 interval. We see that there is a noticeable slow-down in turnover growth as firms approach the threshold, and an equally noticeable catch-up when firms are about to pass the threshold. More formal statistical analysis shows that the slowdown in the growth rate is between 1 and 2 percentage points.

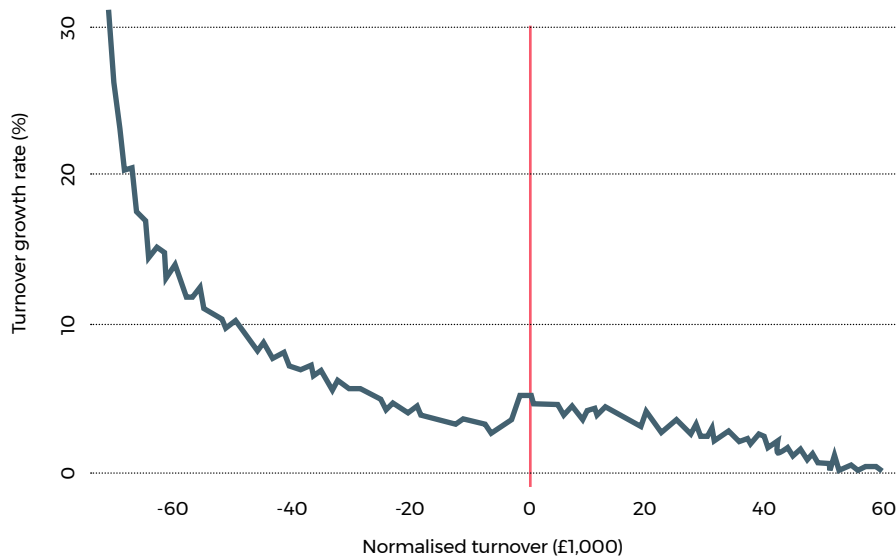
About the author

Ben Lockwood is Professor of Economics at the University of Warwick and a CAGE Associate.

Publication details

This article is based on the paper Liu, L., Lockwood, B., Almunia, M., and Tam, E. (2019). VAT Notches, Voluntary Registration, and Bunching: Theory and UK Evidence. *Review of Economics and Statistics*, forthcoming. https://doi.org/10.1162/rest_a_00884

Figure 3: Evidence of slow-down in VAT growth as firms approach the VAT threshold



Exposing the glass ceiling: hiring practices and gender equality in the UK labour market

By Emma Duchini and Stefania Simion



Kind

together

compassionate

commitment

loyalty

responsibility

Reading a job advertisement is the first step when looking for a job. But is there a relationship between a firm's hiring practices (i.e. how it advertises a job) and its gender equality? We conduct a text analysis of UK job listings, focusing on implicit gender biases, the offer of flexible work arrangements and the decision to post wages. We find that firms' hiring practices are correlated with indicators of their gender equality.

While gender-targeted job advertisements have been banned in most countries, the wording of a job listing is not necessarily gender neutral.

Implicit Association Tests (IAT) show that adjectives frequently encountered in job advertisements such as 'cooperative' and 'committed' tend to be associated with female figures, while traits such as 'competitive', or 'ambitious' are mostly related to male figures. We explore whether the vocabulary employed in job listings helps predict the degree of gender equality in a firm. Research has also shown that women have a stronger preference for flexible working than men, and are less likely to bargain about their pay (Goldin, 2014; Liebbrandt and List, 2015). We want to understand whether the offer of flexible work arrangements correlates with a firm's gender composition, and study if upfront wage information predicts gender pay differentials.

In 2018 and 2019, more than 10,000 UK firms were required to publicly disclose their gender pay gap and gender composition – which we will call GPG indicators. With an average of 40% of women in the top quartile of the wage distribution, and the median gender pay gap at around 12%, data from this pay transparency regime suggest that women still face a (in-)visible glass ceiling in the UK labour market.


We combine the GPG indicators with Burning Glass Technologies (BGT) job advertisement data. BGT offers

more than 50 million individual job vacancies over the period 2012-2019.¹ These are collected from a wide range of online job listing sites and crucially provide the text of the job advertisement and occupational identifiers. Around one third of vacancies also include the name of the employer, which allows us to merge this dataset with firms' GPG indicators.² We explore three key dimensions in this matched data set: gendered wording, flexible working arrangements and wage posting.

Gendered wording

Job advertisements usually describe the ideal candidate, listing a series of desirable behaviours and traits. Psychology literature shows that women are less likely to apply for vacancies citing attributes that are stereotypically associated to men (Gaucher et al, 2011). To study whether these implicit biases predict the degree of gender equality in a firm, we exploit a list of so-called gendered words developed by psychologists through IAT. We assign a gender score to each vacancy, defined as the difference between the number of 'feminine' and 'masculine' words over the total number of words in the job advertisement. A positive (negative) gender score means the vacancy is female-oriented (male-oriented), while a neutral advertisement has a zero score. Figure 1 presents the relative frequency of gendered words (figures a and c), and shows that only a minority of job ads can be classified as gender neutral (graph b).

Figure 1: Gendered wording

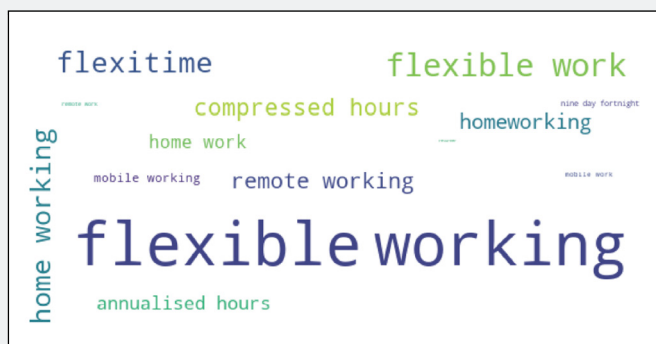


Note: Figures a) and c) present the relative frequency of male-oriented (red bar, graph b) and female-oriented (green bar, graph b) words. The sample includes firms publishing GPG indicators both in 2017/18 and 2018/19, with non-missing registration numbers, and perfectly matched with BGT. N. observations = 4,722.

Flexible work arrangements

Gender differences in preferences for flexible working play a key role in predicting gender occupational segregation, which in turn contributes to explaining the gender pay gap (Wiswall and Zafar, 2018). Firms may reduce gender occupational segregation by offering those forms of flexible working that help women maintain their work-life balance. Yet, this could also have ambiguous effects on within-occupation gender pay differentials whenever flexibility entails a wage penalty. To study these dynamics, we construct a vocabulary of full-time flexible work arrangements (FWA) using job listings from Timewise, a website specialised in flexible working, and the Labour Force Survey definition of FWA.³ We then classify a vacancy as 'flexible' if it includes at least one of the expressions shown in Figure 2. Overall, in our data only around 5% of job listings offer FWA, with flexible working being the most frequent option offered.

Figure 2: Flexible work arrangements




Source: BGT 2014-2019, and Gender Pay Gap Reporting. Note: The figure shows the relative frequency of expressions representing flexible work arrangements (FWA) in BGT job listings. The sample includes firms publishing GPG indicators both in 2017/18 and 2018/19, with non-missing registration numbers, and perfectly matched with BGT. N. observations = 4,722.

Wage posting

Many studies document that women are less likely to ask for wage increases (Babcock et al., 2003) and tend to avoid negotiations when not explicitly encouraged (Leibbrandt and List, 2015). While offering upfront wage information may improve the transparency of the hiring

Figure 3: Hiring practices by occupation and industry



Source: BGT 2014-2019, and Gender Pay Gap Reporting. Note: The figure reports the industry and occupational distribution of hiring practices. The sample includes firms publishing GPG indicators both in 2017/18 and 2018/19, with non-missing registration numbers, and perfectly matched with BGT. N. observations = 4,722.

process, we find that only 22% of job vacancies post automatically identifiable information on wages.


Considering all three dimensions together, Figure 3 shows that there is considerable variation across occupations and industries in these hiring practices, with some positive correlations across them. In particular, occupations with a higher share of female-oriented vacancies tend to have a higher proportion both of flexible positions and of vacancies posting wage information.

How do hiring practices correlate with gender equality within a firm?

In Figure 4, we correlate these hiring practices with the GPG indicators to show several novel stylised facts. All three blue bars show that firms adopting these hiring practices over the observed period have a higher proportion of women at the top of the wage distribution. As for the relationship with the gender pay gap (green bars), unexpectedly, a more female-oriented job ad wording is associated with a larger gender pay gap (leftmost bar), pointing to the complexity of these relationships. However, firms offering more flexible working have a smaller gender pay gap (central bar), suggesting that full-time FWA may help employees manage work and family duties, without incurring wage penalties. Similarly, the rightmost bar shows that wage posting is also associated with a lower gender pay gap, highlighting its potentially beneficial effects for gender equality.

These are only correlations and should not be interpreted as causal relationships. Recruitment practices may be a proxy for broader management strategies. Also, hiring practices may be female-oriented because a firm has a higher proportion of women at the top, and not vice versa. Said differently, we do not know to what extent changing these dimensions of the recruitment process

Figure 4: Hiring strategies and gender pay gap indicators conditional correlations



Source: BGT 2014-2019, and Gender Pay Gap Reporting. Note: The Bar graph reports estimated coefficients from regressions of GPG indicators (averaged across 2017/18 and 2018/19) on hiring strategies (averaged over the period 201/15-2018/19), the occupational composition of firms' vacancies and their sector. The sample includes firms publishing GPG indicators both in 2017/18 and 2018/19, with non-missing registration numbers, and perfectly matched with BGT. N. observations = 4,722.

may crack the glass ceiling. Yet, this analysis shows that hiring practices are related to gender equality within a firm, something that HR recruiters may want to know when designing a new position and its job advertisement. ◀

About the authors

Emma Duchini is a Postdoctoral Research Fellow in the Economics Department at the University of Warwick and a CAGE Associate.

Stefania Simion is a Lecturer in the School of Economics at the University of Bristol.

Publication details

This article is based on Duchini, E., Simion, S., and Turrell, A. (2020). Pay Transparency and Cracks in the Glass Ceiling, *CAGE Working Paper* (no. 482).

Further reading

Adams A., Balgova M., and Qian, M. (2020). Flexible work arrangements in low wage jobs: evidence from job vacancy data. *CEPR discussion paper* (no. 15263).

Babcock, L., Laschever, S., Gelfand, M., and Small, D. (2003). Nice girls don't ask. *Harvard Business Review*, 81 (10), 14-16.

Gaucher, D., Friesen, J., and Kay, A. (2011). Evidence that gendered wording in job advertisements exists and sustains gender inequality. *Journal of Personality and Social Psychology*, 101 (1), 109.

Goldin, C. (2014). A Grand Gender Convergence: Its Last Chapter. *American Economic Review*, 104 (4), 1091-1119.

Leibbrandt, A., and List, A. (2015). Do women avoid salary negotiations? Evidence from a large-scale natural field experiment. *Management Science*, 61 (9), 2016-2024.

Rudman, Laurie A., and Kilianski, S. E. (2000). Implicit and explicit attitudes toward female authority. *Personality and Social Psychology Bulletin* 26.11: 1315-1328.


Wiswall, M., and Zafar, B. (2018). Preference for the Workplace, Investment in Human Capital, and Gender. *The Quarterly Journal of Economics*, 2018, 133(1), 457-507.

Footnotes

1. As BGT expressed concern over the quality of data at the beginning of the sample, we exclude the first two years.

2. In Duchini et al. (2020), we show that this restricted sample represents well the occupational distribution of the UK labour market. To merge it with GPG data, we use a cosine similarity name-matching algorithm for company names, and retain only firms with an exact match, representing two third of the GPG sample with non-missing company registration numbers.

3. Note that here we do not consider flexible work arrangements that give the employer discretion over scheduling, such as shift work or on-call work (Adams et al., 2020).



Who polices the police? Body-worn cameras do work after all

By Pedro Souza

‘To protect and to serve’ is a familiar motto adopted by police forces in the United States and across the world. Some mottoes go further: Minnesotan police aspire ‘To Protect with Courage, To Serve with Compassion!’, and officers in New York are ‘Faithful unto Death’ (*Fidelis ad mortem*). Across the Atlantic, law enforcement is relatively more circumspect, for example, ‘Working together for a safer London’.

These unabashed aspirations often clash with the reality of police and community relationships, which are regularly characterised by a lack of trust between the police and citizens and are sometimes plagued by claims of excessive use of force and racial discrimination. The riots that erupted in the wake of the killing of George Floyd in Minnesota, and growing calls to defund or abolish the police, highlight the importance of improving the relations between the police and the community they serve. In this context, adoption of police body worn cameras (BWC) have been over and over again proposed as a way to prevent excessive use of force, build trust and promote a sense of accountability. As Travis Easter from the San Diego Police Department stated, ‘If officers and citizens are being watched, we are both more liable to do the right thing’ (Siegel, 2020).

Existing work on BWC mostly suggests that they have limited or no effect on police behaviour. We provide experimental evidence – from Santa Catarina in Brazil – that BWC in fact trigger a significant change in officer and citizen interactions. Our findings suggest that the lack of high-quality evidence on the effectiveness of BWC from the many existing studies may be due to the research design, not due to an absence of a treatment effect.

In collaboration with the state Military Police – which is responsible for day-to-day policing activities and is highly visible to the public – we randomised 150 officers to receive the camera as part of their standard equipment and 303 officers were consigned to the control group, who were not allowed to use the camera. Using administrative data, we tracked the outcomes of the universe of the 9,259 dispatches they attended to, covering five police precincts over a three-month period starting in September 2018.

We find statistically and economically sizable results which strongly suggest that accountability through the recording of interactions between the police and citizens de-escalate situations in which conflict or violence could otherwise occur. Our initial set of results focuses on the actual use of the camera. Using the data from the camera logs, we find that the cameras were indeed used among the police apparatus: on average, one quarter of the dispatches in the treatment group were recorded. The internal protocol for camera use mandated that cameras should only be activated if there was an interaction with the citizen, which does not occur in all dispatches.

We then focus on comparing dispatch reports with and without BWC present. We find that dispatches in the treatment group were more diligently recorded: they are 13% more likely to generate a formal referral to the Civil Police (the branch of the police responsible for conducting investigations), prepare criminal reports and eventually push charges through the judiciary. We also find that those reports are 26% more likely to include a victim. We interpret these effects as ensuing from the accountability of the actions of the police and citizens’ alike, promoted by the video recordings. What is more, the proclivity to report on cases of domestic violence increased substantially by 66% comparing the treated with controlled dispatches.

We next consider margins that specifically signal a negative interaction between the police and citizens: We measure the incidence of any discharge of lethal or non-lethal weapons, if handcuffs were deployed or arrests took place, and whether the dispatch resulted in charges of disobedience or contempt towards the police officers. These latter charges may have been used both for legitimate purposes or to justify unwarranted use of force by the police.

On average, we find strong reductions in the use of force. The discharge of lethal and non-lethal weapons is reduced by 56%, the use of handcuffs by 12%, and contempt charges by 48% (while only the latter is statistically significant). The effects seem much stronger among cases that were internally classed as low risk prior to dispatch. In those cases, the use of force is driven close to zero, and contempt falls by close to 70%. No effects are detected among more serious cases that have already escalated prior to dispatch, where the presence of a camera itself may not affect the situational dynamic. Taken together, these results suggest that cameras indeed serve as way to de-escalate conflicts, diffuse tensions, and ensure a better co-operative environment on both sides.

We further estimate that officers learn about the use of the camera and adjust their behaviour even when a camera is no longer being used. To draw out this insight, we leveraged on a second feature of our experiment: we surprised the officers with some randomly selected days in which no camera would be available, even for those in the treatment group. In this way, we can compare the behaviour of treated officers across days in which cameras are being used and days when they are not. We detect no behavioural change beyond the mechanical absence of the recordings themselves. This is a central distinguishing feature of our study from numerous other BWC studies that typically assign cameras to be used randomly across whole shifts. As officers typically rotate across shifts this implies that officers that were subject to a BWC treatment may also appear in control group shifts. We also find that most gains are observed when just a single camera is present, implying that the positive effects are observed even if only a subset of officers has BWC. Overall, these findings suggest that even low-resourced police forces are likely to be able to draw on the positive aspects of BWC by allocating devices in such a way that most, if not all, officers are exposed to their use.

Lastly, we aim to shed some light on the question of accountability and monitoring. While we cannot fully disentangle the implicit monitoring effect of the

police vis-à-vis behavioural changes by the citizens themselves, we find that the composition of the dispatch groups substantially matters. Effects are more pronounced for policing groups formed by heterogenous types, e.g. low- and high-ranked officials. This suggests a third channel for monitoring.

The cameras might not just be a way to promote the accountability of the police to the wider public, or to bring citizens' actions to light and justice, but also for effective monitoring of police officers by other fellow police officers.

Until now, existing empirical evidence of the use of BWC in select localities in the United States or the United Kingdom were not wholly supportive of their merits (Lum, 2019). Our experiment in Santa Catarina, Brazil – a much more challenging setting than previous studies, and in a developing country – shows strong behavioural

responses in accountability and monitoring when BWC are used. They do work, after all. ◀

“On average we find strong reductions in the use of force. The discharge of lethal and non-lethal weapons is reduced by 56%, the use of handcuffs by 12%, and contempt charges by 48%. In cases classed as low risk prior to dispatch, the use of force is driven close to zero, and contempt falls by close to 70%.”

About the author

Pedro Souza is Assistant Professor of Economics at the University of Warwick and a CAGE Associate.

Publication details

This article is based on Barbosa, Fetzer, Soto and Souza (2020). Monitoring and Trust: The Effects of Body-worn Cameras on the Police Activity. Working paper.

References

- Lum, C, Stoltz, M., Koper, C., and Scherer, J. (2019). Research on Body-worn Cameras: What we Know, What we Need to Know. *Criminology & Public Policy*, 18, pp.93-118.
- Siegel, J. (2020). Why Police Say Body Cameras Can Help Heal Divide with Public. *The Daily Signal*, (online)] (Accessed 28 August 2020).

How early adopters spread Twitter across the United States

By Karsten Müller and Carlo Schwarz

Twitter is used by around 330 million people worldwide. But how did it achieve such widespread adoption? New research finds that much of Twitter's uptake across the US can be pinpointed to one event: The South by Southwest festival in Austin, Texas, in 2007.



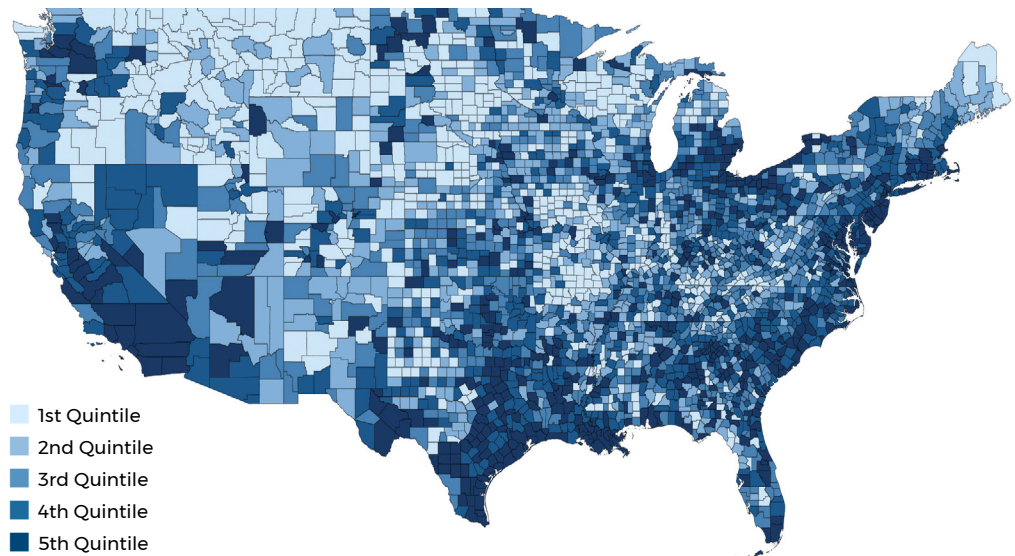
The adoption of new technologies is an important factor in explaining the large productivity differences across and within countries. When analysing the adoption of new technologies, it's widely agreed that small and even arbitrary variables can have a big effect on eventual uptake. We find evidence of this for one of the most popular social media apps, now used by around 330 million people worldwide: Twitter.

Twitter provides an interesting case to study technology adoption for two reasons. First, it is a technology with low adoption costs – everybody can open an account for free. Second, there are significant network effects, because the value of being on Twitter increases with the number of other people that use the social network.

We investigate the impact of the 2007 South by Southwest conference and festival (SXSW), which marked an early tipping point for the adoption of Twitter in the United States. We find that SXSW 2007 – which only had around 7,000 registered attendees at the 'interactive' part of the event – left its imprint on the frequency of Twitter usage across US counties today (Figure 1).

At the 2007 SXSW festival, Twitter held a launch event with a special option that allowed users to join Twitter by simply sending a text message, and screens in the main hallways showed tweets about the event. These measures proved to be extremely effective in spurring Twitter adoption. The daily volume of tweets increased from around 20,000 to 60,000. Figure 2 shows how SXSW influenced Twitter's success: we see a clear spike of tweets about the event during the SXSW conference in mid-March 2007, followed by an upward shift in the growth of the total number of tweets. While total tweets grew by 55% from February to March, this growth accelerated to over 190% from March to April. Twitter has been

Figure 1: Twitter usage per capita in the United States



on a successful growth path for the years since then.

Our analysis exploits that the home counties of SXSW attendees received a boost in the number of early-stage Twitter users around the time of the 2007 event. These inflows of additional early adopters put these counties on a differential growth path to a higher level of Twitter usage which is still evident today.

There are three pieces of evidence to support the idea that these early adopters were key to Twitter's rise. First, we compare counties with and without new SXSW followers in March 2007 and observe an uptick in Twitter adoption with the beginning of SXSW that persists until the end

of 2015. Figure 3 traces the impact of early adopters on Twitter usage per capita in their home counties, where the orange line marks the beginning of the SXSW festival. The data exhibit an S-shaped adoption curve typical for the diffusion of innovations (Arrow, 2000; Rogers, 2010), starting from the day of the festival, as more and more users signed up over time. The estimates imply that a one standard deviation increase in SXSW followers who signed up in March 2007 increased Twitter adoption by around 22% by the end of 2015.

Secondly, we find that early Twitter adopters were indeed largely connected to the SXSW festival. Figure 4 plots the share of Twitter

Figure 2: Twitter activity around SXSW 2007 and over time





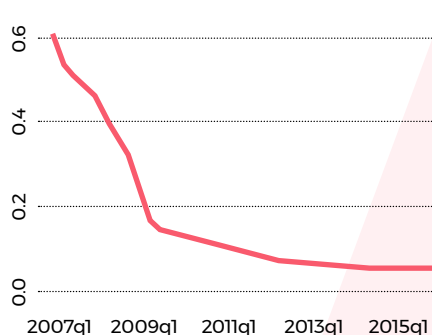
Figure 3: Twitter Adoption per Capita in SXSW Home Counties



users that followed either the SXSW festival or an SXSW follower who joined during the time of the event. In March 2007, as many as 60% of Twitter users had either a first or second degree connection to the SXSW festival. With the diffusion of Twitter over time, this decreased to around 5% today, as the platform's usage spread to wider and wider subsections of the population.

Finally, the idea of the diffusion of Twitter starting from the SXSW festival can also be seen in the predictive power of the locations of SXSW attendees. As shown in Figure 5, the predictive power is highest for the year of the festival and decreases from there onwards. This suggests that, over

Figure 4: Share of Twitter users with SXSW connection



time, Twitter was used by more and more people who did not have a direct connection to the SXSW festival.

Taken together, this evidence shows the 2007 SXSW festival led to increased adoption of Twitter in the home counties of attendees, and this pattern of technology adoption persists until today. This evidence could be crucial for both companies and policymakers. The success of new technologies such as the COVID-19 tracing app in the UK, depends on a relatively widespread adoption across the country. Policymakers should thus consider how encouraging a targeted group of early adopters could improve the overall use and success of such

technologies. If 7,000 people at SXSW can influence the geography of Twitter usage 10 years after the 2007 event, pushing early adopters to spread a technology might be more successful than trying to reach the general public. ◀

About the authors

Karsten Müller is a Postdoctoral Research Associate at the Julis-Rabinowitz Center for Public Policy and Finance at Princeton University and a CAGE Associate.

Carlo Schwarz is Assistant Professor of Economics at Bocconi University and a CAGE Associate.

Publication details

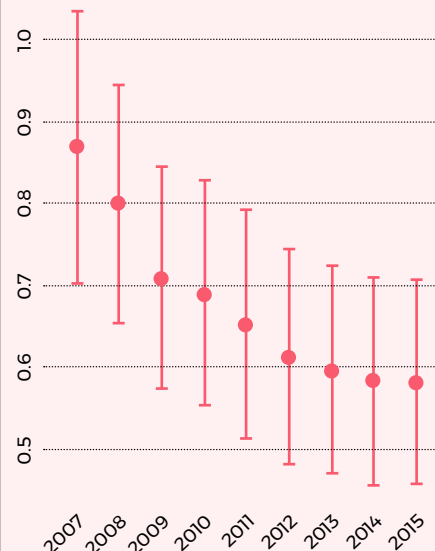
This article is based on the working paper Müller, K. and Schwarz, C. (2020). From Hashtag to Hate Crime: Twitter and Anti-minority Sentiment. Available at SSRN 3149103.

Further reading

Arrow, K. J. (2000). Increasing Returns: Historiographic Issues and Path Dependence. *The European Journal of the History of Economic Thought*, 7 (2), pp. 171-180.
Rogers, E. M. (2010), *Diffusion of Innovations*, New York: Simon and Schuster.

“Policymakers should consider how encouraging a targeted group of early adopters could improve the overall use and success of technologies. If 7,000 people at SXSW can influence the geography of Twitter usage 10 years after the original 2007 event, pushing early adopters to spread a technology might be more successful than trying to reach the general public.”

Figure 5: The effect of SXSW on Twitter adoption decreases over time



Parting Shot

We are all Nate Silver now

By Mirko Draca, Director, CAGE

The COVID pandemic will be noted for many things, but when it comes to the policy sphere there has never before been a time when the use of data and statistics has been so prominent. Political press conferences in last six months – including those featuring the centre-of-attention loving President Trump – have often had the added feature of graph-filled slide packs. This was very rare in the pre COVID-19 era. I’m also regularly reading wonkish and thoughtfully argued data-focused pieces in the popular press, which previously mainly served to inform me about the travails of celebrities.

But this data-focused debate over COVID-19 policy increasingly has a dueling quality to it. Statistical analysis is now being integrated into the disciplines of rhetoric and polemic. And then there is the simple matter that there is now *just so much* information and analysis available in 2020. ‘Data science’ boomed as a professional industry in the 2010s but now it is ever present in news debates, with lots and lots of ‘armchair’ contributions being fired off by the minute. To update a turn of phrase from Richard Nixon, ‘we are all Nate Silver now’.

Nixon coined the phrase ‘we are all Keynesians now’ when he embraced a more interventionist economic policy stance in the mid-1960s. It was meant to convey the idea that there was a consensus on certain aspects of economic policy. But just as Keynesian principles seemed to go wrong in the 1970s after the consensus was declared, I think that the consensus on the value of data analysis is clearly having some problems. Data is now an element in the overall contest for attention and influence.

And in one of those weird bits of symmetry, this year is the 50th anniversary of Alvin Toffler’s 1970 book *Future Shock* – a key text on the social challenges of technology

and one that prominently flagged the emerging problem of ‘information overload’. This is, as they say, an understudied topic but I think the technological changes of the late 2010s – mobile computing, social media, fast broadband – facilitated an expansion of information that tipped us into Toffler’s ‘information overload’ state.

So how do we navigate this challenge in the context of data, research and communication in general? While the model of ‘gatekeepers’ that regulated access to information and institutions up until the 2000s is gone, there is now a clear requirement for information curation and filtering. Simply put, there is now actually a *market* out there for curation and filtering, one that I don’t think is being specifically served.

This is where I think the ‘information economics’ literature that kicked off in the early 1970s is likely to be useful to us. For example, I definitely detect elements of George Akerlof’s (1970) *The Market for Lemons* in various corners of the internet and news media: that is, the overwhelming amount of bad (or just hard to understand) information drives out the good. It’s hard to think of a single tool that will get around this fundamental problem of asymmetric information. My best guess is that this



“I think the technological changes of the late 2010s – mobile computing, social media, fast broadband – facilitated an expansion of information that tipped us into Toffler’s ‘information overload’ state.”

is a long game that’s fundamentally about building up curatorial reputations and taking steps to increase the size of the market for high quality information. This latter set of steps related to market size boils down to education, specifically with respect to data literacy. One side effect of the surge of interest in data analysis in the COVID-19 era is that there has been a *de facto* increase in data literacy (or at least an interest in learning). So, even in the midst of this current information crisis there is a chance to capitalise and expand the market for good information.

Thinking strategically, the weakness of ‘bad’, polemical information is that it turns over fast – ad hoc modifications are needed to fit the polemic as new data comes in. Better quality information and analysis

is able to adapt and should be more persistent. This is what I’m hoping for. If my hopes are dashed, then the future media of my imagination looks like a website with loads of bizarre graphs on one side of the page and a big column of celebrity news on the other. At least I will have celebrities to read about. ◀

Further reading

Akerlof, C. (1970), ‘The Market for Lemons: Quality Uncertainty and the Market Mechanism’, *Quarterly Journal of Economics*, 84(3): 488-500.

Toffler, A (1970), *Future Shock*, Random House: New York.

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 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).

Director: **Mirko Draca**

Research Director: **Bishnupriya Gupta**

Impact Director: **Arun Advani**

Research theme leaders

Economic History: **Stephen Broadberry**
and **James Fenske**

Culture, Behaviour and Development: **Sharun Mukand**
and **Robbie Akerlof**

Political Economy: **Thiemo Fetzer** and **Christopher Roth**

Public Policy and Data: **Roland Rathelot**, **Lucie Gadenne**
and **Ludovica Gasse**

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](https://twitter.com/cage_warwick)



cage.centre@warwick.ac.uk