

Digital media: Is it good for us?

p.5 **Facebook:
What's not to like?**

p.8 **Life outside the bubble**

Matthew Gentzkow
(Stanford)


p.11 **State policies and
online sympathy for
religious extremism**

Elizabeth Dekeyser
(IAST)

p.14 **Data literacy
has become
essential**

Jonathan Levin
(Stanford)





New technologies, new promises and new problems

Jacques Crémer

SAVE THE DATE
14th Digital Economics
Conference in Toulouse
January 7-8, 2021



Dear friends,
We are sending this newsletter in the middle of the COVID-19 crisis. This crisis has brought home the importance and benefits of the digital age. Like many of you, all of TSE is locked out of our wonderful new building and this is coming from our computers at home - 20 years ago it would have been so much more difficult!

The extraordinary technological progress of the digital era has transformed the way we interact with each other. We have only begun to embrace the benefits and to face the challenges. Is digital media bringing us together or pulling us apart? How will we adjust to a world dominated by AI, machine learning, and data science?

To answer such questions, we will need more evidence-based analysis by TNIT researchers like Stanford's Matthew Gentzkow. In this issue of TNIT News, he presents his findings that deactivating Facebook leaves people less informed but happier. He also reviews two recent studies on the promise and pitfalls of diversifying news consumption.

Meanwhile, IAST political scientist Elizabeth Dekeyser uses a machine-learning tool to analyze Twitter responses to the 2015 Paris killings, showing that state policies can sway sympathy for terror attacks. And in an exceptional interview Jonathan Levin, who is both a TNIT member and the Dean of the Stanford Graduate School of Business, talks about his work on a life-saving economic mechanism to promote vaccines, and the challenges of preparing leaders for our fast-changing future.

Take care of yourselves and of others: respect the isolation measures recommended or mandated by your local health authorities.

Is digital media pulling us apart?



Matthew Gentzkow is Professor of Economics at Stanford University. Previously at the University of Chicago Booth School of Business, he was awarded the John Bates Clark Medal in 2014.

For this issue of TNIT News, he presents two white papers exploring the impact of digital media. In the first, he presents his own findings that deactivating Facebook leaves people less informed but happier. In the second, he reviews two recent studies on the promise and pitfalls of diversifying news consumption.

“ Facebook has 2.3 billion monthly active users, and by 2016 the average user was spending nearly an hour per day on Facebook and its sister platforms. There may be no technology since television that has so dramatically reshaped the way we get information and spend our time.

Facebook: What's not to like?

Within the span of a decade, social media has woven its way deep into our lives. Facebook has 2.3 billion monthly active users, and by 2016 the average user was spending nearly an hour per day on Facebook and its sister platforms. There may be no technology since television that has so dramatically reshaped the way we get information and spend our time.

Early on, platforms like Facebook, Twitter and Instagram were hailed for their potential to make communication and the sharing of information easier. Now, the conversation is dominated by potential harms, from addiction to depression to political polarization. Despite the abundance of speculation about the potential effects of social media, hard evidence remains scarce.

In a recent paper, we provide a large-scale randomized evaluation of the welfare impacts of Facebook, the largest social media platform. This provides the largest-scale experimental evidence to date on Facebook's impact on a range of outcomes. We find that deactivating Facebook for one month leads people to spend more time with friends and family.

It also leaves them less informed about the news, less polarized in their political opinions, and a little happier and more satisfied with their lives. We find that after the time off Facebook, users want it back, but they use it significantly less than before their one month “detox”. Our findings are in line with other important work on the same topic (see, for example, [here](#) and [here](#)).

Study design

We recruited 1,600 US Facebook users online and randomized them into a “deactivation” and a “control” group. The deactivation group received US\$102 in exchange for staying off Facebook for the four weeks leading up to the US midterm election in November 2018; the control group kept using Facebook as usual.

We measured a suite of outcomes using text messages, surveys, emails, and administrative voting records. We recorded key measures twice – once in October, before the beginning of the deactivation

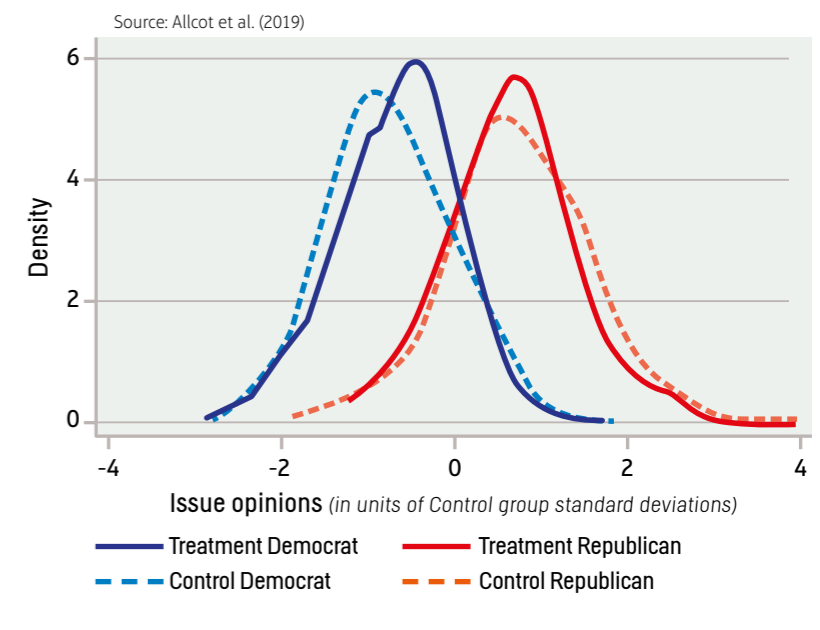


Figure 1 shows how Facebook deactivation reduced polarization on a range of measures, including views on policy issues such as immigration and policing. The dashed lines show the distribution of these views among control-group Democrats (blue) and Republicans (red).

The solid lines represent the treatment group. In both groups Democrats' views are well to the left of Republicans' views, but the inter-party differences are visibly smaller in the treatment group, suggesting that deactivation moderated views in both parties.

Figure 1: Polarization



Deactivating Facebook for one month leads people to spend more time with friends and family. It also leaves them less informed about the news, less polarized in their political opinions, and a little happier and more satisfied with their lives.

period, and once in November, after the deactivation period had concluded. We then compared the changes in those outcomes in the deactivation group to those in the control group.

To verify deactivation, we repeatedly pinged subjects' public Facebook URLs. These return a valid page when an account is active but return an error message when an account is deactivated. Overall, 90 percent of users in the treatment group followed our instructions and deactivated their accounts.

Key findings

Being off Facebook freed up an average of one hour to spend on other activities. How people use this extra time helps us understand which activities Facebook is crowding out, and this in turn tells us something about Facebook's effects. If Facebook time just replaces other social media or similar digital activities, the effects of deactivation might be small. If it replaces high-quality social interaction with family and friends, we might worry more about outcomes like (un)happiness, loneliness, and depression. If it replaces consumption of high-quality news sources, we might worry more about impacts on political knowledge and polarization.

Our findings show that Facebook does not substitute for other digital activities - if anything, people spend less time on other social media and digital platforms when their Facebook accounts are deactivated. Instead, Facebook time comes entirely from offline activities including face-to-face socializing and solitary activities like watching TV.

Our next set of findings focuses on news knowledge and political outcomes. Deactivating Facebook caused a significant reduction in total news consumption, and significantly reduced news knowledge and political engagement. Among other things, we find that subjects in the deactivation group were much worse at answering quiz questions about current issues in the news. At the same time, the deactivation group ended up significantly less polarized by a range of measures, including their views on policy issues such as immigration and policing.

In terms of well-being, we find that Facebook deactivation causes small but significant increases in self-reported individual life satisfaction and happiness, and significant decreases in self-reported levels of anxiety. We also elicited self-reported well-being using daily text messages, and find positive, though less precise effects of Facebook deactivation on this outcome. As shown in Figure 2, an index of all measures together shows that deactivation caused significant improvements in overall well-being.

Finally, we measured whether deactivation affected people's demand for Facebook after the study was over, as well as their opinions about Facebook's role in society. As the experiment ended, participants assigned to the deactivation group reported planning to use Facebook much less in the future. Several weeks later, the treatment group's reported usage of the Facebook mobile app was about 11 minutes (22 percent) lower than in control.

Big picture

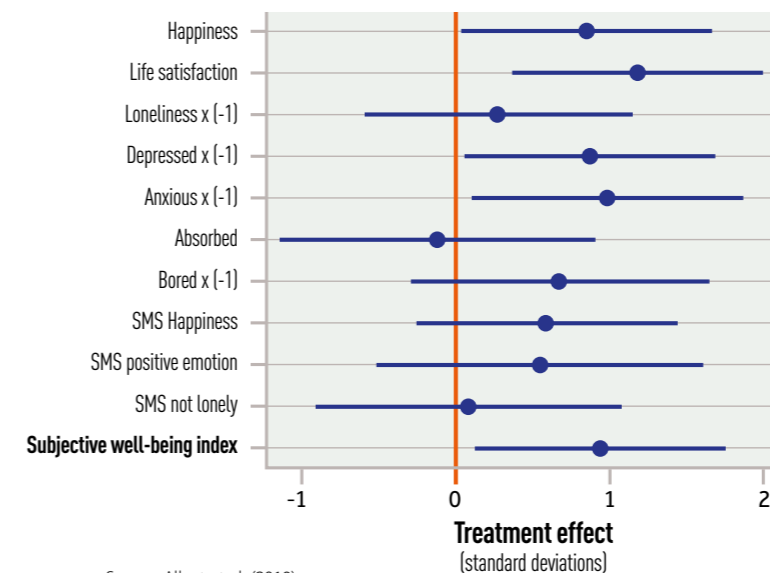
There is no question that many users perceive the benefits of Facebook to be large. A majority of participants would require a payment of \$100 or more to deactivate Facebook for



a month. Even after a four-week "detox," these valuations remained high and our participants continued to spend substantial time on Facebook every day. The results on news consumption and knowledge suggest that Facebook is an important source of news and information. Our participants' answers in free-response questions and follow-up interviews make clear the diverse ways in which Facebook can improve people's lives, whether as a source of entertainment, a means to organize a charity or an activist group, or a vital social lifeline for those who are otherwise isolated. Any discussion of social media's downsides should not obscure the basic fact that it fulfills deep and widespread needs.

At the same time, our results also make clear that the downsides are real. We find that four weeks without Facebook improves subjective well-being and substantially reduces post-experiment demand, suggesting that forces such as addiction may cause people to use Facebook more than they otherwise would. We find that while deactivation makes people less informed, it also makes them less polarized, consistent with the concern that social media have played some role in the recent rise of polarization in the US.

The trajectory of views on social media - with early optimism about great benefits giving way to alarm about possible harms - is a familiar one. Innovations from novels to TV to nuclear energy have had similar trajectories. Along with the important existing work by other researchers, we hope that our analysis can help move the discussion from simplistic caricatures to hard evidence, and provide a sober assessment of the way a new technology affects both individual people and larger social institutions.



Source: Allcott et al. (2019)

Figure 2: Effect on subjective well-being

Each point in Figure 2 measures the effect of Facebook deactivation on well-being outcomes, measured in standard deviations. The lines to the left and right of each point indicate the 95 percent confidence interval. All outcomes are scaled so that the right of the figure indicates more positive outcomes. (Thus, measures of loneliness, depression, anxiety, and boredom are inverted.) The final row shows an index of all measures together, showing that deactivation caused significant improvements in overall well-being.

FIND OUT MORE:

Matthew's own paper 'The Welfare Effects of Social Media,' written with Hunt Allcott, Luca Braghieri, and Sarah Eichmeyer, is due to be published in *American Economic Review*. Read about his research at: gentzkow.people.stanford.edu

Life outside the bubble

Matthew Gentzkow - (Stanford)



Digital media have been accused of causing many social ills. One of the most serious charges is deepening political polarization.

Early commentators like Cass Sunstein and Eli Pariser described how a world with vast selection of content and algorithmically tailored filters could trap individuals in bubbles of like-minded information where they would hear their own views and prejudices echoed endlessly back and rarely if ever encounter a conflicting view¹.

Examples such as the *“Blue Feed, Red Feed”* graphic produced by the *Wall Street Journal* showed how close these predictions might be to reality. Although recent evidence suggests that the worst fears are likely overstated², it also shows that political divisions are indeed deepening, and that content accessed through digital media can be an important contributor³.

The solution to this problem might seem obvious: change the algorithms to break the filter bubble. In the wake of the 2016 election, a raft of startups formed to offer tools to diversify peoples’ news diets, including *Read Across the Aisle*, *allsides.com*, *echochamber.club*, and *Escape your Bubble*. Many commentators have called for platforms like Facebook and Twitter to change their algorithms to give less priority to like-minded content⁴. The “contact hypothesis” - the idea that simply exposing one group

to another can reduce divisions and hostility - is supported by a large body of literature.

As appealing as breaking our bubbles may sound, there are some reasons to question whether it will be an effective solution to polarization. The most obvious problem is that people read like-minded content for a reason. If Facebook or Twitter start filling peoples’ feeds with content from the other side, they may just ignore it. Demand for like-minded news may be driven by a bias toward confirmation, a genuine belief that like-minded sources are more trustworthy, or just greater interest in the stories those sources choose to highlight. Whatever combination of these factors is at play, it means that exposing people to diverse content is very different from getting them to actively and thoughtfully engage with it. Perhaps for this reason, none of the startups mentioned above seems to have gained large-scale traction and at least two of the four appears to be defunct.

Even more troubling, prior work suggests that exposure to cross-ideological experiences can sometimes produce backfire effects that deepen divisions⁵. It is not hard to imagine that forcing a committed liberal to sit and listen to Donald Trump speeches for an hour might increase rather than decrease the intensity of the liberal’s partisan ire. Handing a committed conservative a packet of liberal social-media memes might well have a similar effect.

What happens when digital media users are pushed outside their bubbles? High-quality evidence remains limited, but we are beginning to accumulate valuable data points. Here, I review two of the most compelling recent studies.

5 - See references cited in [Bail et al. \(2018\)](#).

1- See [Republic.com](#) by Sunstein and [The Filter Bubble](#) by Pariser.
2- [Gentzkow and Shapiro \(2011\)](#); [Allcott & Gentzkow \(2017\)](#); [Boxell et al. \(2017\)](#).
3- [Boxell et al. \(2017\)](#); [Allcott et al. \(forthcoming\)](#).
4- See, for example, [here](#) and [here](#).

Red Bot, Blue Bot on Twitter

Among the largest published studies on cross-partisan social-media exposure to date is *“Exposure to Opposing Views can Increase Political Polarization: Evidence from a Large-Scale Field Experiment on Social Media”* by Christopher Bail and co-authors (PNAS 2018).

The authors began by building two customized Twitter bots, one conservative and one liberal. The conservative bot was programmed to retweet a random selection of content from accounts the authors had identified as among the most influential conservative “opinion leaders” on Twitter (including politicians, organizations, and commentators). The liberal bot was programmed to do the same for content from liberal accounts. The bots retweeted 24 posts per day.

Next, the authors recruited a sample of self-identified Republicans and Democrats who reported using Twitter at least three times per week. They surveyed these subjects about their views on a set of political issues, then randomly assigned them to either a control group or a treatment group. Treatment subjects were paid

to follow the bot opposite to their ideology (conservative bot for Democrats, liberal bot for Republicans), and given additional incentives to pay close attention to the content of that bot’s tweets. They also completed a series of follow-up surveys that measured compliance as well as changes in the index of political-issue views. Control users completed the follow-up surveys but were not asked to change anything about their social-media behavior.

The results are not encouraging for supporters of bubble-bursting interventions. Not only did subscribing to the Twitter bot from the other side not significantly reduce the polarization of political views, it appeared to produce a backlash effect. This effect was small but highly significant for Republicans, with those exposed to the liberal bot shifting their views on the 7-item issue scale to be between 0.1 and 0.5 points more conservative depending on the estimation method. For Democrats, following the conservative bot produced a small and insignificant shift of views in the liberal direction.

Several points are important to note in interpreting this study. First, the sample consisted of heavy Twitter users with clear party attachments - hardly representative of US voters. It may be that backlash effects are especially likely among such engaged and committed partisans, and that effects would be different among those who are more moderate or who use social media less. That said, highly engaged partisans are those we would worry most about being affected by filter bubbles, so if cross-ideological content does not work for them it is unlikely to be a solution to the overall problem. Second, the partisan content assigned in the experiment may have been relatively extreme. It could be that more moderate content designed to appeal to both sides would have had a different effect. Finally, the outcome measure is an index of fairly broad political views such as “Government is almost always wasteful and inefficient” and “The best way to ensure peace is through military strength.” It could be that cross-ideological content could reduce some kinds of polarization even if it did not moderate these kinds of ideological views.

Liking your enemy on Facebook

A more recent study in a similar vein is *“Social Media, News Consumption, and Polarization: Evidence from a Field Experiment,”* a new working paper by Ro’ee Levy.

The author recruited a large sample of US Facebook users via Facebook ads. After completing a baseline survey, subjects were randomly assigned to a liberal or conservative treatment group, or a control group. The liberal treatment group was asked to “like” four liberal news outlets (e.g., MSNBC), an action which would lead more content from these outlets to appear in their Facebook feeds. The conservative treatment group was asked to “like” four conservative news outlets (e.g., Fox News). Neither group was



given any incentives to do follow through with this suggestion, but roughly half of both groups complied. The design was stratified so that both liberal and conservative subjects were included in both the liberal and conservative treatments – in other words, subjects could be treated with content either opposed to or consistent with their own ideology.

Outcomes are measured in three ways. First, participants had to log in to the baseline survey with their Facebook accounts and give permission to the author to observe the outlets they “liked” and the posts they shared. Second, some participants installed a Google Chrome extension that allowed the author to observe the content of their news feeds and the articles they read. Third, participants completed a final survey two months after the end of the experiment. The main question is how being assigned to like cross-ideological outlets affected polarization. 17,629 subjects completed the surveys, 8,080 subjects were offered the Chrome extension, and 1,838 installed the Chrome extension and kept it for the duration of the study.

There are several important differences between this study and the Bail et al. Twitter experiment. The population was a broad cross-section of Facebook users not directly screened on social media use or political affiliation. The cross-ideological content came from large news outlets rather than Twitter opinion leaders and may therefore have been more moderate or of wider appeal. Finally, in addition to measuring polarization on issue views the author also measured affective polarization – the extent to which respondents felt warmly or coldly toward those on the opposite side.

The first set of results show that the treatment indeed changed the mix of news that subjects saw in their news feed and also the mix of news they consumed. Those in the pro-attitudinal treatment group (i.e., assigned to like news outlets aligned with their ideology) saw 67 additional posts from the assigned outlets in their feeds. Those in the counter-attitudinal treatment group (i.e., assigned to like news outlets of the opposite ideology) saw 31 additional posts. The counter-attitudinal group was induced to make between 1 and 2 additional visits to the assigned outlets on average over the course of the study. The treatment also produced a detectable change in the composition of posts subjects shared.

The second set of results show the impact on polarization. Consistent with the Bail et al. study, there is no evidence that injecting content from the other side into subjects’ feeds reduced the polarization of their issue views. There is also no evidence of backlash effects; the result is a precisely estimated zero.

Most strikingly, exposure to the counter-attitudinal treatment significantly reduced affective polarization relative to the pro-attitudinal treatment. Subjects felt relatively less “cold” toward the other party and reported that they found it easier to see the

other side’s perspective. The magnitude of these effects is small in absolute terms (a few points on a 100-point “thermometer” scale) but moderately large relative to both the changes over time in affective polarization and the impact of other interventions.

Discussion

What do we learn from these studies taken together? One finding consistent across the studies is that relatively small interventions can meaningfully change the mix of content people are exposed to. The effects are small as a share of the total content flowing through peoples’ feeds, but they are sufficient to produced detectable effects on survey outcomes. This supports the view that even modest interventions have the potential to make a significant difference.

Both studies suggest that diversifying the content people are exposed to is unlikely to be sufficient to narrow polarization of issue views, and the Bail et al. study provides a significant note of caution that poorly conceived interventions may produce the opposite of the intended effect. While the study does not have enough power to unpack exactly what caused the backlash, one might infer that showing highly partisan opinion content from one side to strongly engaged partisans on the other side may be an especially risky approach.

The Levy findings on affective polarization, on the other hand, provide the most unambiguous piece of good news for the value of escaping bubbles. One possibility is that the interventions in both studies had this effect and it would have been detectable in the Bail et al. study had the authors measured affective polarization. Another possibility is that the news outlets in the Levy study were particularly conducive to helping those on each side see the others’ perspective. Either way, it is encouraging that a relatively low-cost and scalable intervention could produce a meaningful reduction in hostility.



State policies and online sympathy for religious extremism

by Elizabeth Dekeyser (IAST)



With a PhD from MIT, Elizabeth Dekeyser is a political scientist at the Institute for Advanced Study in Toulouse (IAST) who takes a data-driven approach to studying identity, immigration, and citizenship.

Through careful analysis of social media, her research shows that state policies can influence the responses of immigrant-origin communities to terrorism.

In 2015, France was shaken by two terror attacks – at the offices of the satirical newspaper Charlie Hebdo in January, and at the Bataclan concert hall and other public spaces in November. The vast majority of France’s Muslim community condemned the attacks and expressed solidarity with the victims, often using their religious identity as a reason for their response. But a small minority expressed, if not outright support, sympathy for the attackers and their motivations, also citing their religious identity. These responses reflect broader variation among immigrant-origin communities about the compatibility of religious and national identities, and understanding the origins of this online rhetoric can provide important insights into broader questions of integration and social cohesion in multicultural societies.

While a range of factors unquestionably influence individuals’ choices to express attack-supporting online rhetoric, I focus on understanding the role of the state’s engagement with marginalized communities. This is both because of the relative weight of its influence – in many densely immigrant communities, the state is the primary contact between the community and

society more broadly, and thus has an outsize impact on behavior and beliefs – as well as its potential for policy implications.

Twitter rhetoric

To examine this, I scraped over 1.6 million geolocated tweets posted in the 48 hours following the Charlie Hebdo and Bataclan attacks. Using a support vector machine, a supervised machine-learning model, I classified attack-supporting and non-attack-supporting tweets. A map of the tweets can be seen in Figure 1.




Figure 1: Attack-Supporting Tweets
White points indicate all tweets, red points attack-supporting tweets



Politicians and pundits have sometimes argued that community-development policies have been a failure. Yet they have no counterfactual; in other words, it is unclear if outcomes would have been worse if no investment occurred.

The most common attack-supporting responses varied between the two attacks. Following the Charlie Hebdo attack, attack-sympathetic responses primarily discussed how the satirical newspaper provoked the attackers through their disrespectful depictions of the prophet. For example, *“But they provoked this through mocking Islam for years... they should have expected a massacre like this.”* Others merged this with concerns about the lack of attention paid to deaths in Muslim societies, such as, *“in Palestine, there are two hundred deaths a day and nobody hears, Charlie criticized the prophet man!”* Following the November attacks, attack-supporting tweets sometimes attributed the attacks to Western actions, such as, *“We should have as much disgust for the jihadists as for the leaders who provoked the Muslims of the world, #Iraq #Syria #Libya #Palestine,”* a response fueled by one terrorist allegedly saying *“this is for Syria”* during the attack.

Tweets do not provide much data about their authors. Yet by merging the geographic information in the tweets with administrative data, it is possible to gain a clearer understanding of the places the tweets are coming from. These towns were significantly poorer on a range of metrics, as well as slightly smaller. Using a measure I developed from the number of google searches in a town, they also had lower levels of connectivity. Using data from Google maps on religious structures in each of France’s 36,000 towns, I found that they were also more likely to have mosques and other community religious structures. These correlations, however, do not give any causal insight into what circumstances might encourage individuals to write more attack-supporting tweets.

Election of far-right parties

The first factor I examine in order to understand how anti-state policies influenced attack-supporting tweets is the election of far-right parties to local councils. Far-right parties have increased their electoral presence in France and throughout the West over the past decade. Their central policy tenets often include harsh stances on immigration and broader concern about the cultural threat posed by Islam and immigrant-origin communities. I hypothesize that these harsh stances

can lead to perceptions of rejection among immigrant-origin individuals and a resulting embrace of extreme responses, such as sympathy for terror attacks.

To test this, I use an electoral regression discontinuity design based on the election of a single far-right party member to a municipal council. This design compares towns where a far-right party member barely won or barely lost an election, based on the assumption that at the cutpoint at which victory occurred, this victory could be considered as-if random in these towns and so allow for examination of its causal effect.

As seen in Figure 2, I find that towns that elected far-right parties see a statistically and substantively significant increase in attack-supporting tweets – from about one in every 200 tweets, to one in every 100. This finding has important implications for far-right parties, which are often elected under the assumption that their hardline stances will lead to less, not more, social fracturing within immigrant-origin communities.

Community engagement

I run a second design that examines the effect of community investment projects on attack-supporting rhetoric. France, like many countries, has for many years poured money into impoverished

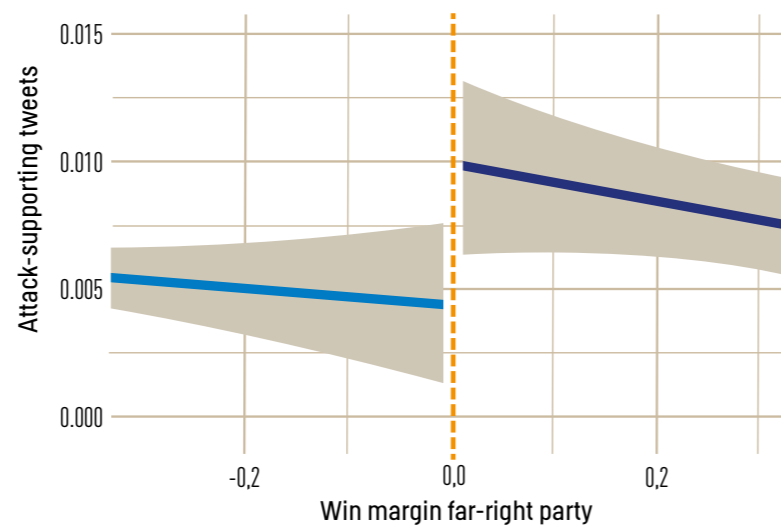


Figure 2: The effect of far-right council members on attack-supporting tweets

communities, many of which are densely populated with immigrant-origin individuals. With these communities still struggling with social and economic integration decades later, politicians and pundits have sometimes argued that these policies have been a failure. Yet they have no counterfactual; in other words, it is unclear if outcomes would have been worse if no investment occurred. I use a unique design to answer this question.

Specifically, I examine the effect of Quartiers Prioritaires, a community development policy that over the past five years has poured resources into impoverished communities for everything from associational life to public housing renovations. Uniquely, however, inclusion in this scheme was based on strict income and population thresholds and the policy was implemented between the two terror attacks. Using Twitter results as an outcome, this allows for the use of a “difference-in-differences” design to examine whether this community investment caused a differential change in towns that did and did not receive the funding.

I find that towns that received the funding had on average 20% fewer attack-supporting tweets than those that did not. This substantively significant result is even more impressive given the time frame - only 11 months of investment in these communities. I argue that this effect functions through a similar mechanism as the election of far-right parties. Physical investment in these communities signaled to individuals that their community did, indeed, matter, decreasing perceptions of rejection and support for hardline stances following terror attacks.

Taken together, these findings have important implications for how the state can influence the social and economic integration of marginalized communities. Harsh stances meant to encourage integration, such as those embraced by the far right, can have a significant backlash effect. Policies meant to encourage social and economic integration, however, do indeed have the desired effect.



Towns that received funding had on average 20% fewer attack-supporting tweets than those that did not – after only 11 months of investment.

'Data literacy has become essential'

An interview with Jonathan Levin (Stanford)



Jonathan Levin is Philip H Knight Professor and Dean of Stanford Graduate School of Business. Winner of the John Bates Clark Medal in 2011, his recent research has focused on internet platforms, the healthcare system, and ways to incorporate new datasets into economic research. Here he talks to TNIT News about his work on a life-saving economic mechanism to promote vaccines, and the challenges of preparing leaders for the digital age.

Please note that this interview was conducted in February 2020

Does the Dean of Stanford Graduate School of Business have any time to carry out research? If so, could you tell us about your most recent work?

My job puts me in a position to come across great research topics, but less time to execute on them! I did just write a short paper with Michael Kremer and Chris Snyder looking at a program we helped to design about a decade ago, the Advanced Market Commitment (AMC) for pneumococcal vaccine. It originated with an idea that Michael had to correct the market failures that have led to underinvestment in vaccines targeted at low-income countries.

Vaccines give some of the highest returns on investment in terms of saving lives, but we lack vaccines for diseases such as malaria, and the historical adoption of existing vaccines in developing countries has been terribly slow. With an AMC, donors pledge a fund to subsidize purchases of newly developed vaccines, creating an incentive for vaccine development and for investments in production capacity.

The pneumococcal AMC involved a donor fund of \$1.5 billion, and several billion dollars from GAVI (Global Alliance for Vaccine

Initiatives), to help low-income countries purchase the vaccine, with economically motivated rules to allocate the money. A decade on, 150 million children have been vaccinated, with an estimated 700,000 lives saved. It is hard to assess exactly what would have happened without the program, but adoption of pneumococcal vaccine has been much faster than historical norms, and quite a bit faster than adoption of rotavirus vaccine, which happened at a similar time without an AMC.

There are still a lot of interesting questions, primarily around pricing and the potential for ongoing competition in the market. It was a lot of fun to write

“ Skills that will be at a premium in organizations are ones that can't easily be automated: leadership, communication, teamwork, being able to pose good questions and deal with ambiguity. The last job to be automated will be the CEO's.

the paper and see what we could learn from the past 10 years. I give enormous credit to Michael for coming up with an economic mechanism that took hold of people's imaginations, generated so much support, and has had such a significant impact.

How do you prepare future business leaders for the huge changes that AI, machine learning, and data science are bringing to organizations?

Data literacy has become essential. At the GSB, we've rethought our core data-science curriculum, which we teach in a flipped classroom model, very hands-on, with students working on real datasets. Then there are opportunities to take advanced classes in the business school and across campus.

Even more important, we want students to understand the opportunities created by digital technology, but also to think responsibly about its consequences. One of the lessons of the past decade is that technology can move faster than the regulatory framework and public understanding of its effects. Business leaders need to be able to think ahead about the social and political effects of technology.

This brings up a final point that is sometimes lost in discussions of AI, and is especially relevant for MBA education. Over time, the skills that will be at a premium in organizations are ones that can't easily be automated: leadership, communication, teamwork, being able to pose good questions and deal with ambiguity. The last job to be automated will be the CEO's. An effective MBA education combines analytical training and humanistic training, and the rise of AI makes the latter even more important.

What role does economics play in that education process?

The view we've taken at Stanford is that the data revolution and the rise of AI inherently cross many disciplines. Last year, we created the Stanford Institute for Human-Centered Artificial Intelligence (Stanford HAI), to span the campus and connect technologists, social scientists, and humanists thinking about different aspects of AI. An example of how that works is a new class we started offering at the business school, co-taught by Fei-Fei Li, one of the leading computer scientists in AI, and Jennifer Aaker, a social psychologist who studies meaning and happiness. They enrol an equal mix of MBA students and CS graduate students, and the class is about designing AI systems that promote human well-being.

What about economics? I'm biased of course, but economics is incredibly helpful when it comes to thinking about the implications of new technology. Economics naturally leads you to think about the second- and third-order consequences of technological change; the

way that people and firms respond and adapt when there's a powerful new general-purpose technology, which is what we have today. There's also another valuable aspect of economics training, namely that it teaches you to use data to tell coherent stories; or conversely, to see if a story you'd like to tell is supported or disproved by data. We're starting to emphasize that aspect of economics more in the way we teach, and it's terrific.

What aspects of your job are you most passionate about?

First, something I love about Stanford is being right at the frontier of new things that are happening in business, in technology, in society. In my current job, I get to spend a lot of time thinking about those things, and how the GSB can stay ahead of them. Second, I'm a big believer in institutions, and great universities are some of the most important institutions on this planet. I find it hard to imagine doing something more meaningful than helping to lead one. Finally, it's hard to enjoy being a dean unless you take a lot of pleasure in the success of others, because a big part of the job is enabling faculty and students to realize their own aspirations. Fortunately, I'm surrounded by people who manage to do pretty amazing things on a regular basis.



“ Economics naturally leads you to think about the second- and third-order consequences of technological change; the way that people and firms respond and adapt when there's a powerful new general-purpose technology, which is what we have today.

TNITNews

TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

- ▶ *Scientific Director:* Jacques Crémer
- ▶ *Editorial contributions:* James Nash
- ▶ *Graphics:* Olivier Colombe
- ▶ *Illustrations:* I-Stock

TNIT

Toulouse School of Economics
1, Esplanade de l'Université
31080 Toulouse Cedex 06
Tel: +33 (0)5 67 73 27 68

www.tse-fr.eu/digital
tnit@tse-fr.eu

The Toulouse Network for Information Technology (TNIT) is a research network funded by Microsoft, and coordinated by TSE. It aims at stimulating world-class research in the economics of information technology, intellectual property, software security, liability, and related topics.

All the opinions expressed in this newsletter are the personal opinions of the persons who express them, and do not necessarily reflect the opinions of Microsoft, TSE or any other institution.

Designed by award-winning architects, TSE moved into its new home earlier this year.

At the time of writing, the building was closed due to Covid-19 restrictions but we hope to welcome everyone back very soon.

↻ Issue 21 ↻ April 2020

 Digital
Center

 Microsoft