TNIT News
TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

Is AI creating or destroying jobs?

Daron Acemoglu
MIT

Did Covid-19 bring Americans together?

Matthew Gentzkow
Stanford

Political polarization: Which countries are the most divided?

Matthew Gentzkow
Stanford

➔ Issue 24 ➔ September 2021

TSE Digital Center

Microsoft

A scientist's view: Taking the right measures

Jacques Crémer
TNIT Coordinator

Dear friends,

Scientific activity is often described as a two-step process. First, one develops theories; second, one gathers data to test these theories. There are many ways in which the scientific enterprise functions differently. One of them is that, in practice, large efforts are done to simply measure "what is out there". Astronomers have spent considerable amounts of time and effort measuring the number of galaxies and their repartition in the universe. An important component of climate science is the measurement of the effects that global warming has already had on temperatures, water salinity, precipitation, etc. Although observation would seem to be the simplest of tasks, in many cases it requires a great deal of technical sophistication and imagination.

This issue of *TNIT News* presents three measurement exercises which exemplify how this process plays out in economics and other social sciences.

The first article by Daron Acemoglu summarizes research that he has conducted with his MIT colleague David Autor, Jonathan Hazell from London School of Economics, and Pascual Restrepo from Boston University. They asked: Can we observe whether artificial intelligence (AI) is destroying jobs, by, for instance, replacing humans for some tasks, or creating jobs, perhaps by increasing productivity as it provides humans with the tools to do more tasks? Dear reader,

I recommend that before looking at Daron's piece you spend a few minutes asking yourself how you would go about such a task. Not obvious, is it? There are two difficulties: how to observe where AI is used in the economy; and how to correlate this use with changes in employment. I will let you dive into Daron's fascinating piece and admire how he and his co-authors tackled these issues. (And, of course, as a supremely gifted economist, Daron is able to venture a little beyond measurement and speculate about the reasons for their conclusions and the consequences for public policy.)

The other two examples of measurement exercises are presented by Matthew Gentzkow and concern political polarization. As recent news events have sadly shown, polarization has increased substantially in the US. Given the importance of what happens there and the global accessibility of US news, I must admit that I had a tendency to believe that this increased polarization is a world-wide phenomenon. In his study of international trends in affective polarization, Matt shows that the US is an outlier, and that this phenomenon has been much less prevalent in other countries. Here again, notice the difficulty of the measurement issue, from finding an operational definition of polarization to exploiting hundreds of polls conducted in different countries with different methodologies.

In the second piece presented by Matt, he discusses the following issue: "Did the beginning of the pandemic prompt an increase in polarization in the US?" Just as climate scientists look at different pieces of data to estimate changing temperatures over the past 20 years, Matt harnesses different studies to show that this has not been the case. It would be fascinating to know whether this trend has continued into 2021 and whether it is international.

As always, thank you for reading!

SAVE THE DATE
15th Digital Economics Conference, January 13-14, 2022
More information

Is AI creating or destroying jobs?

Daron Acemoglu
MIT

Daron Acemoglu is an Institute Professor at MIT and is the author of five books. His academic work covers a wide range of areas, including political economy, economic development, economic growth, technological change, inequality, labor economics, and economics of networks. He was awarded the John Bates Clark Medal in 2005, the Erwin Plein Nemmers Prize in 2012, and the 2016 BBVA Frontiers of Knowledge Award.

Daron has received the inaugural T. W. Shultz Prize from the University of Chicago in 2004, and the inaugural Sherwin Rosen Award for outstanding contribution to labor economics in 2004, Distinguished Science Award from the Turkish Sciences Association in 2006, the John von Neumann Award, Rajk College, Budapest in 2007, the Carnegie Fellowship in 2017, the Jean-Jacques Laffont Prize in 2018, the Global Economy Prize in 2019, and the CME Mathematical and Statistical Research Institute Prize in 2021.

Artificial intelligence (AI) is one of the most promising technologies currently being developed and deployed. There is a lot of excitement, some hype, and a fair bit of apprehension about what AI will mean for our security, society, and economy.

Central to many of the debates is whether AI is creating or destroying jobs. Despite the huge and growing interest in this question, we know relatively little about the answers.

Some commentators are convinced that AI is the harbinger of a jobless future (e.g., Ford, 2015; West, 2018; Susskind, 2020). Yet others are equally adamant that AI will enrich work experiences and increase human productivity, contributing rather than detracting from job growth (e.g., McKinsey Global Institute, 2017). These contrasting visions persist in part because there is very little evidence on what AI is doing to work and workers. There are currently no representative data sets of AI, so we lack representative evidence on whether there has been a major increase in AI adoption (as opposed to just talk of AI). It is possible to find examples of AI technologies either replacing work or complementing workers, precisely because AI, as a broad technological platform, is capable of doing both. The level of job displacement that AI will create is thus partly a matter of societal and business choice (Acemoglu and Restrepo, 2019).

Tracking the rise of AI activity

In recent work, "AI and Jobs: Evidence from Online Vacancies", David Autor, Jonathon Hazel, Pascual Restrepo and I have studied AI adoption in the US and its labor market implications. AI adoption can be partially identified from the footprint it leaves at adopting establishments as they hire workers specializing in AI-related activities, such as supervised and unsupervised learning, natural language processing, machine translation, or image recognition. To put this idea into practice, we built an establishment-level data set of AI activity based on the near-universe of US online job vacancy postings from Burning Glass Technologies for the years 2007 and 2010 through 2018. This data set, which has been used in several recent papers, contains detailed information on occupations and the skills required for each posted vacancy.

> **There are still relatively few vacancies in core AI areas, such as machine learning and natural language processing, though the rate of growth since 2015-2016 has been staggering**

We then linked the adoption of AI and its possible implications to the task structure of an establishment. Put simply, the idea is to look for a set of "AI-suitable" tasks, which may be targeted by AI applications, then investigate whether establishments with a high fraction of such tasks are more likely to show rapid AI adoption, as measured by the hiring of AI workers. There is no consensus on which tasks are AI-suitable. Nevertheless, a number of recent studies have started developing systematic ways of measuring which tasks can be performed or augmented by current AI technologies. For example, Felten, Raj and Seamans (2018) construct an index of the effect of AI on various occupations, meant to capture both the ability of AI algorithms to substitute for humans and their complementarity to humans. They build on experts' assessments of areas in which AI has made important advances then map these areas to the set of tasks performed by different occupations. Alternatively, Brynjolfsson, Mitchell and Rock (2018) build a measure of the suitability of an occupation's task to be performed by machine learning. Webb (2000), on the other hand, uses natural language processing on the text of patents to map them to specific tasks performed within various occupations. Each of these measures captures a different aspect of AI suitability (and we show they are quite distinct). There is information in each of them and our work uses all three of them.

> The fact that AI is a broad technological platform suggests there are important decisions for both corporations and public policymakers. What type of AI do we want? Can we create more jobs than we destroy?

The data paint a clear picture about AI activity, regardless of which specific measure one looks at. There is a notable takeoff in AI vacancy postings starting in 2010, but these postings remain very low until around 2015-2016, then undergo an inflection, trending up strongly thereafter. There are still relatively few vacancies in core AI areas (such as machine learning, natural language processing, etc.), though the rate of growth since 2015-2016 has been staggering.

We also find that AI-adopting establishments start demanding different skills than before, and in fact there is some evidence of increased "skills churn" associated with AI. This bolsters the case that there are significant changes in the organization of production and thus the skills demanded by businesses at the forefront of AI adoption.

Where does this growth come from?

We show that there is a strong association between the baseline task structure of an establishment and AI activity. This relationship is present with all of the measures mentioned above. As important, it remains even when we focus on establishments within a narrow industry or, more notably, when we compare two establishments belonging to the same multi-establishment firm that still differ in terms of their baseline task structure. This is evidence that AI adoption is being at least partially targeted to a specific set of AI-suitable tasks. This correlation, however, does not answer the key question we started with …

Is AI creating or destroying jobs?

The answer seems to be: Mostly, it's too soon to tell. Despite a remarkable takeoff, there is still very little AI activity at the moment, and AI-impacted job changes may be a small drop in a big bucket. The number of AI-suitable tasks may grow and lead to the hiring of more workers than before because of the rollout of AI technologies. Or, conversely, the workers who previously performed these tasks may be replaced by AI algorithms.

Nevertheless, we see some evidence of fewer vacancies for non-AI positions in the more heavily impacted establishments (those with a high fraction of AI-suitable tasks). For example, establishments with a high share of AI-suitable tasks in 2010 subsequently show significantly slower growth in vacancies. Yet, confirming our conclusion that AI activity is still too small, this establishment-level result does not translate into slower growth in the more AI-exposed occupations or industries.

So where do we go from here?

The evidence we gathered – coupled with advances in machine learning, big data and other areas of AI – suggests that the rapid takeoff in AI activity will continue in the years to come. This may imply more displacement (similar to the negative hiring effects we may be seeing already at some establishments), but AI is a broad technological platform and can be used in many different ways. The fact that early AI is targeted to specific tasks does not mean that, as the technology matures, it will not have other applications. There is already evidence that AI technologies are being used for new product development and reorganization (Bresnahan, 2019), and these uses may intensify in the years to come.

The fact that AI is a broad technological platform also suggests that there are important decisions for both corporations and public policymakers. What type of AI do we want? If AI can create and destroy jobs at the same time, can we make sure that we create more jobs than we destroy?

We sometimes hear a narrative suggesting that there is a clear path of future technology. For AI, a broad technological platform with many applications, this may be particularly untrue. The disagreement about the effects of AI for workers is rooted in the fact that AI can destroy as well as create jobs. But this also implies there is a lot of room for public policy and corporate strategies in shepherding AI in a direction that is more beneficial for society.

For more research by Daron on this subject, see our previous edition of *TNIT News (Issue 17)*:
www.tse-fr.eu/sites/default/files/TSE/documents/ChaireJJL/TNIT/issue_17.pdf

KEY FINDINGS

- AI vacancy postings took off in 2010, but remained very low until around 2015-2016, trending up strongly thereafter.
- There are still relatively few vacancies in core AI areas but the rate of growth has been staggering and looks set to continue.
- We see some evidence of fewer non-AI vacancies in establishments heavily impacted by AI. Yet this does not yet translate into slower growth in the more AI-exposed occupations or industries.
- We also detect the beginning of an AI-driven "skills churn" in which the use of AI technologies is associated with changes in the skills demanded.

FURTHER READING

- **Daron Acemoglu, David Autor, Jonathon Hazel, and Pascual Restrepo** (2020), "AI and Jobs: Evidence from Online Vacancies." *MIT mimeo*
- **Daron Acemoglu and Pascual Restrepo** (2019), "The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand." *Cambridge Journal of Regions, Economy and Society 13(1): 25-35.*
- **Timothy Bresnahan** (2019), "Artificial Intelligence Technologies and Aggregate Growth Prospects." *Unpublished manuscript, Stanford.*
- **Erik Brynjolfsson, Tom Mitchell, and Daniel Rock** (2018), "What can Machines Learn, and What Does it Mean for Occupations and the Economy?" *AEA Papers and Proceedings. Vol. 108.*
- **Edward Felten, Manav Raj, and Robert Seamans** (2018), "A Method to Link Advances in Artificial Intelligence to Occupational Abilities." *AEA Papers and Proceedings. Vol. 108.*
- **Martin Ford** (2015), "The Rise of the Robots." *Basic Books, New York.*
- **McKinsey Global Institute** (2017), "Artificial Intelligence: The Next Digital Frontier?" *Discussion Paper.*
- **Daniel Susskind** (2020), "A World Without Work: Technology, Automation and how We Should Respond." *Penguin, UK.*
- **Michael Webb** (2020), "The Impact of Artificial Intelligence on the Labor Market." *Unpublished manuscript, Stanford.*
- **Darrell West** (2018), "The Future of Work: Robots, AI, and Automation." *Brookings Institution Press.*

Did Covid-19 bring Americans together?

Matthew Gentzkow

Stanford

Matthew Gentzkow is the Landau Professor of Technology and the Economy at Stanford University. He studies applied microeconomics with a focus on media industries. He received the 2014 John Bates Clark Medal, given by the American Economic Association to the American economist under the age of 40 who has made the most significant contribution to economic thought and knowledge. Other awards include the 2016 Calvó-Armengol International Prize, the Alfred P. Sloan Research Fellowship, grants from the National Science Foundation, National Institutes for Health, Sloan Foundation, and a Faculty Excellence Award for teaching. He is a fellow of the American Academy of Arts and Sciences and the Econometric Society, a senior fellow at the Stanford Institute for Economic Policy Research, and a former co-editor of *American Economic Journal: Applied Economics*. He earned his PhD from Harvard in 2004.

The scientific, economic, and social challenges of responding to the coronavirus pandemic have been compounded in the US by political divisions. Studies from the early days of the pandemic show that partisan divisions were among the most significant drivers of health behaviors, concern about the virus, support for specific policies, attributions of responsibility, and even beliefs about basic facts. This echoes similar divisions among politicians and the media. It seems possible that the pandemic has been yet another force pushing toward greater polarization.

This is not the only possible narrative, however. A different possibility is that the health crisis might have pulled Americans together (at least temporarily) — due to a "rally around the flag" effect as is often seen in times of war or natural disaster, or perhaps simply by giving Americans something to focus on other than politics.

A plague on all our houses

In a recent paper[1], we turn to new data sources to see how polarization evolved during the pandemic. We focus on affective polarization — the extent to which partisans feel more negatively toward the opposing political party than toward their own[2].

Affective polarization in the US has been steadily increasing in recent decades (see Figure 1), and this has generated widespread concern about its impact on democratic institutions and representation, legislative gridlock, and partisan violence. This trend has been a source of alarm to both policymakers and academics, and one recent paper describes the study of its causes and consequences as "one of the most influential literatures in contemporary American politics scholarship"[3].

We find no evidence that affective polarization rose during the health crisis. Our main measure suggests that affective polarization in fact fell significantly with the onset of the pandemic. Three of five other data sources display a similar downward trend, with the other two showing neither a decline nor an increase. A survey experiment adds further evidence, showing that priming respondents to think about the pandemic significantly reduces affective polarization. Taken together, our results suggest a cautiously optimistic conclusion that the coronavirus may have brought partisans together in the face of a common threat.

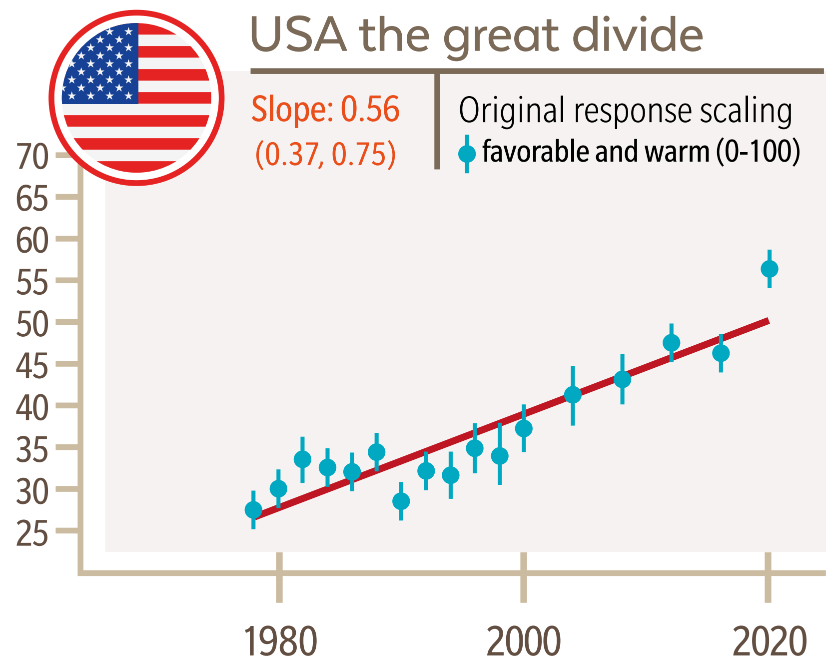

Figure 1: Long-run trends in affective polarization

Using data from American National Election Study surveys going back to the late 1970s, Figure 1 shows the long-term trend in polarization in partisan affect, defined as the difference between respondents' feelings toward their own party vs. the opposite party on a 1-100 scale where higher numbers indicate "warmer" feelings. The magnitude of this gap has increased from roughly 27 at the beginning of the series to roughly 56 in the most recent years.

Source: Levi Boxell, Matthew Gentzkow, and Jesse M. Shapiro, 2021, "Cross-country trends in affective polarization," Stanford University working paper

Party politics

Measures of affective polarization vary in the type of attitudes elicited (e.g., feelings, trust, or behaviors) and the subject of those attitudes (e.g., voters, parties, or candidates). To avoid relying on a single measure or data source, we report trends across six different measures and data source combinations.

Figure 2 reports trends in affective polarization towards partisans or parties. Panel A uses our preferred large-scale survey data to show that, from July 2019 until the start of the pandemic in the US, the trend of affective polarization was relatively flat. After the first publicized coronavirus-related death in the US, however, affective pola-

(1) Levi Boxell, Jacob Conway, James N. Druckman and Matthew Gentzkow, 2021, "Affective polarization did not increase during the coronavirus pandemic". *Quarterly Journal of Political Science* (forthcoming). This article for *TNIT News* draws on material from the authors' Stanford University working paper.

(2) Shanto Iyengar, Gaurav Sood, and Yphtach Lelkes, 2012, "Affect, not ideology: A social identity perspective on polarization". *Public Opinion Quarterly, 76(3): 405-431.*

(3) David E. Broockman, Joshua L. Kalla and Sean J. Westwood, 2020, "Does affective polarization undermine democratic norms or accountability? Maybe not." *UC Berkeley working paper.*

Our main measure suggests affective polarization fell significantly with the onset of the pandemic. An additional survey experiment indicates that thinking about the pandemic reduces polarization

rization exhibits a significant decline before ticking back upward following the death of George Floyd.

Using questions about feelings towards partisan members of Congress, Panel B shows a smaller decrease in affective polarization during the onset of the pandemic. Using a separate panel of respondents from Druckman et al. (2020), Panel C shows little change in affective polarization between July 2019 and April 2020.

The Trump effect

Figure 3 turns from feelings toward parties to feelings toward Donald Trump, focusing on the difference between his favorability as reported by Democrats and Republicans.

Panel A reports a slight upward trend in this difference in feelings prior to the onset of the pandemic, but there is a significant subsequent decline. Panel B reports a similar decline in partisan differences in presidential approval ratings. Panel C uses a separate data source to show that the Trump approval gap between Republicans and Democrats is relatively constant in December 2018 and December 2019, but significantly smaller in April 2020, before returning close to pre-pandemic levels in August-November 2020.

Pandemic priming

To supplement these results, we also conduct an experimental analysis to see if priming people to think about coronavirus leads people to express more or less polarized attitudes. This priming strategy follows previous work by our coauthor Jamie Druckman, among others.

The experiment has three conditions. First, in the coronavirus treatment, respondents are asked to read two news article excerpts that cover the initial phases of the pandemic and to reflect on their own experiences and faith in the United States' ability to respond. Second, in the placebo treatment, subjects are asked to perform an analogous exercise where they reflect on news articles unrelated to coronavirus — specifically, articles about Prince Harry and Meghan Markle's separation from the UK royal family (which occurred just before the onset of the pandemic). Finally, in the control group, subjects are not asked to perform any exercise.

We then ask subjects standard questions about their feelings toward political parties and groups. The experimental results show that, relative to the control group, the coronavirus group displays significantly lower affective polarization. This is not true for the placebo group. The effect mainly reflects less negative sentiment toward the opposing party.

A double-edged sword

Overall, combined with existing evidence on the polarized response to the pandemic, our results show that a crisis may at once decrease affective polarization while simultaneously exacerbating its consequences (for example, partisan divisions in behavioral responses to the pandemic).

Group attitudinal changes and related behavioral changes need not align. Scholars and practitioners who examine interventions to mitigate polarization may need to separately distinguish the size of attitudinal divisions from their consequences

Our results show that a crisis may at once decrease affective polarization while simultaneously exacerbating its consequences, such as partisan divisions in behavioral responses

Figure 2: Recent trends in affective polarization

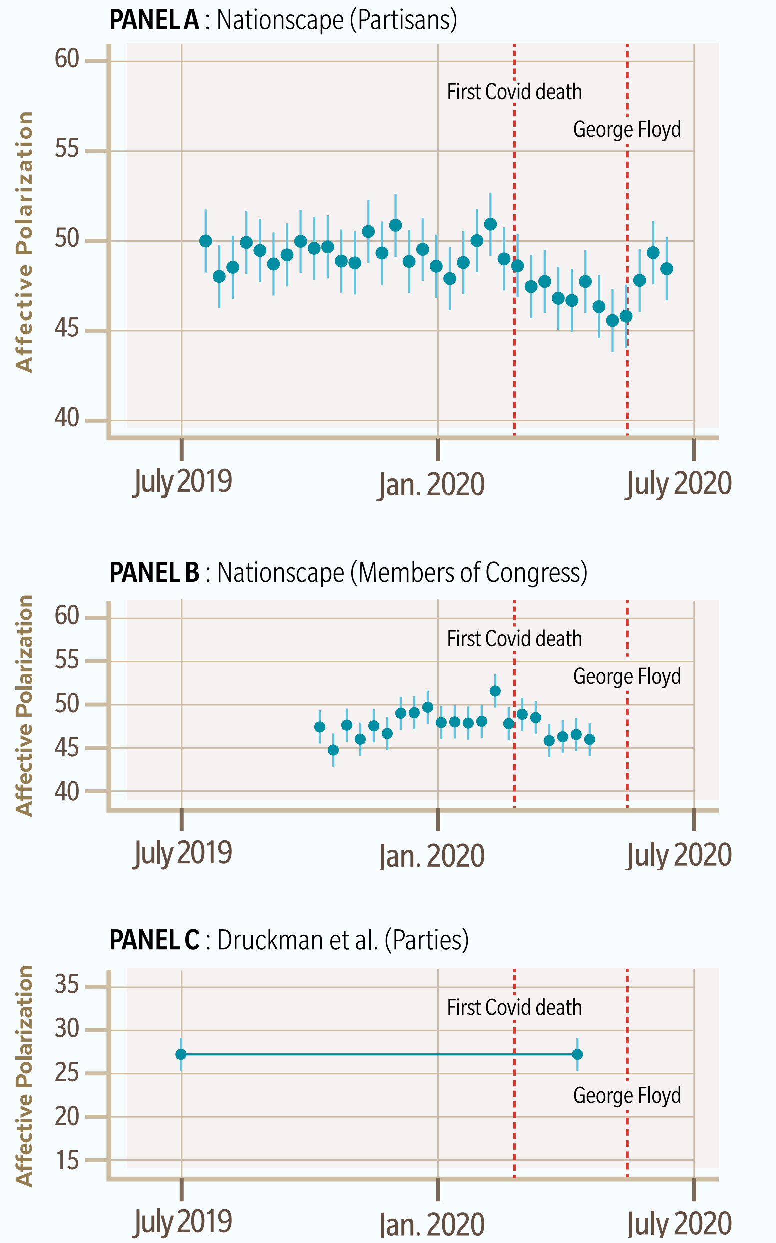

Figure 2 shows the extent to which members of each party feel more negative toward the opposite party than their own. **Panel A** uses a large-scale survey dataset from Nationscape that covers more than 300,000 interviews between July 2019 and July 2020. This is our preferred measure due to the size, frequency, and consistent methodology of the Nationscape data. Affective polarization was relatively flat from July 2019 until the first publicized coronavirus-related death in the US (on February 29th, 2020), after which we see a significant decline until the death of George Floyd.

Using Nationscape data about feelings towards partisan members of Congress, **Panel B** shows a smaller decrease during the onset of the pandemic. **Panel C** uses separate data from Druckman et al. (2020) and indicates little change in affective polarization between July 2019 and April 2020.

Source: Boxell, Conway, Druckman, and Gentzkow, 2021.

Figure 3: Recent trends in partisan feelings toward Trump

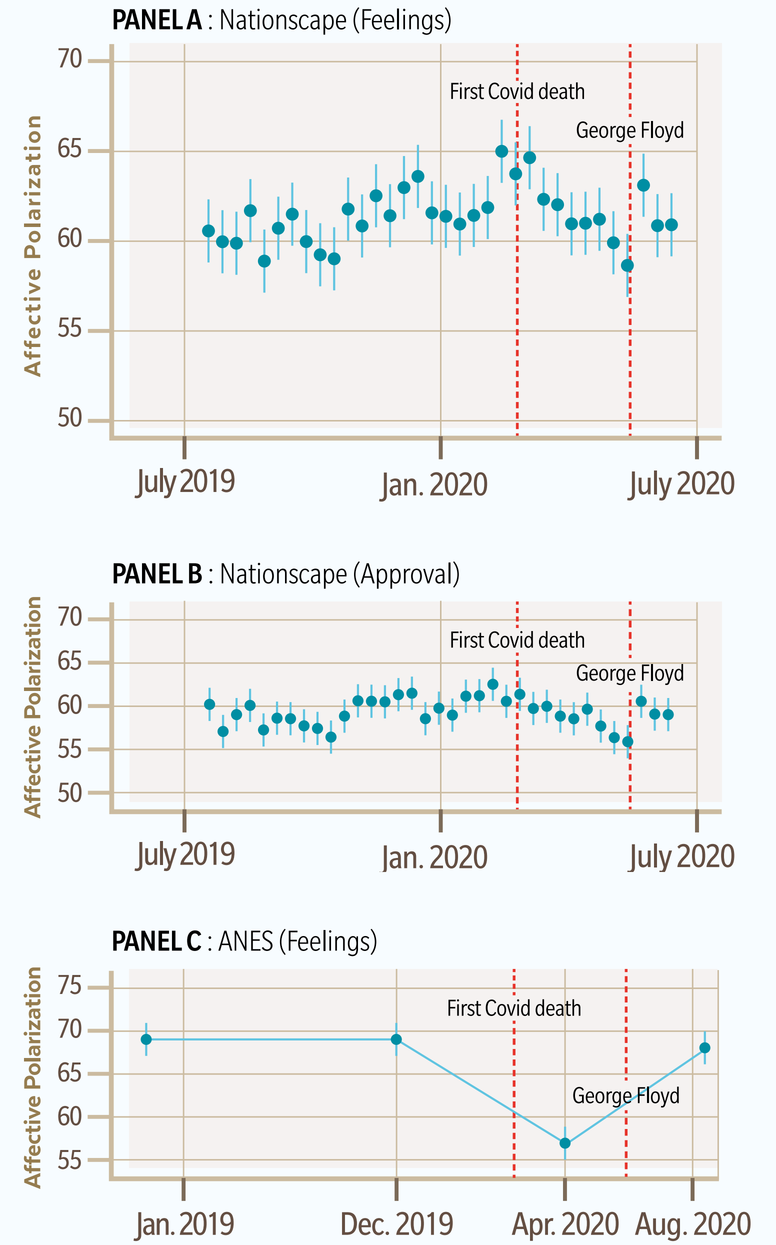

Figure 3 shows the difference between Democrats and Republicans in their feelings about Donald Trump. Using Nationscape data, **Panel A** reports a slight upward trend in this difference prior to the onset of the pandemic, followed by a significant decline. **Panel B** reports a similar decline in partisan differences in presidential approval ratings. **Panel C** uses a separate survey data source, the American National Election Study, to show that the Trump approval gap between Republicans and Democrats is relatively constant in December 2018 and December 2019, but significantly smaller in April 2020, before returning close to pre-pandemic levels in August-November 2020.

Source: Boxell, Conway, Druckman, and Gentzkow, 2021.

KEY TAKEAWAYS

- We find no evidence that affective polarization rose during the pandemic. Four of six data sources suggest it declined; the other two suggest neither a decline nor an increase.
- Our survey experiment shows that priming respondents to think about the pandemic significantly reduces affective polarization.
- Looking at this evidence together, we conclude that the pandemic is unlikely to have increased affective polarization and may well have decreased it.

Which countries are the most divided?

Matthew Gentzkow
Stanford

There are many ways to measure political polarization among voters. We can look at peoples' views on policy issues like taxes and immigration, or the intensity of their self-declared ideologies and partisan affiliations. We can look at how consistently they vote for one party or another. Many of these measures show trends toward increasing polarization in the US in recent decades.

For the clearest picture of America's deepening political divides, however, it helps to look not at traditional political choices but rather at measures of how people feel about those in the opposite party. Affective polarization is a standard measure of these feelings, defined as the extent to which people report feeling more negatively toward the opposite political party than to their own. It was popularized as a polarization measure by Shanto Iyengar and co-authors.[1]

Us and them

Affective polarization has risen substantially in the US in recent decades. In 1978, the average partisan rated in-party members 27.4 points higher than out-party members on a "feeling thermometer" ranging from 0 to 100. In 2020, the difference was 56.3 points. Scholars have argued that growing affective polarization may have important consequences, including a reduction in the efficacy of government, increasing the self-segregation of social groups, and altering economic decisions.

In a recent paper[2], we examine whether affective polarization has seen similar increases in other developed democracies over the past four decades. This kind of comparative evidence remains rare. It is interesting not only for its own sake but also because it tells us something about the likely drivers of polarization. In particular, if rising polarization is the result of factors like the growth of the internet and social media, globalization, and growing inequality, which have been present in all such countries, we might expect rising polarization to be relatively universal. We show that this is not the case. Among the 12 OECD countries for which we were able to get data for the past four decades, the US stands out with the largest increase in polarization. In five other countries — France, Denmark, Canada, New Zealand, and Switzerland — polarization also rose, but to a lesser extent. In the remaining six countries — Japan, Australia, Britain, Norway, Sweden, and (West) Germany — polarization fell. Focusing on the period after 2000, all countries except Britain and Germany exhibit a positive linear trend, with the US having the largest estimated trend among all sample countries. Overall, these results suggest rising polarization may be driven by factors that are more distinctive to the US.

> "Among the 12 OECD countries for which we were able to get data, the US has the largest increase in polarization. After 2000, all countries except Britain and Germany exhibit a positive linear trend

One important point to note is that a flat or declining trend over the 40 years of our sample does not rule out the possibility that countries have seen rising polarization in the most recent years. Britain, for example, shows a slight overall decline but a clear increasing trend post-2000 (and post-Brexit).

(1) Shanto Iyengar, Gaurav Sood and Yphtach Lelkes, 2012: "Affect, not ideology: A social identity perspective on polarization", *Public Opinion Quarterly, 76(3): 405–431.*
(2) This article for *TNIT News* draws on material from Levi Boxell, Matthew Gentzkow and Jesse M. Shapiro, 2021: "Cross-Country Trends in Affective Polarization", *Stanford University working paper.*

Figure 1: Trends in affective polarization by country

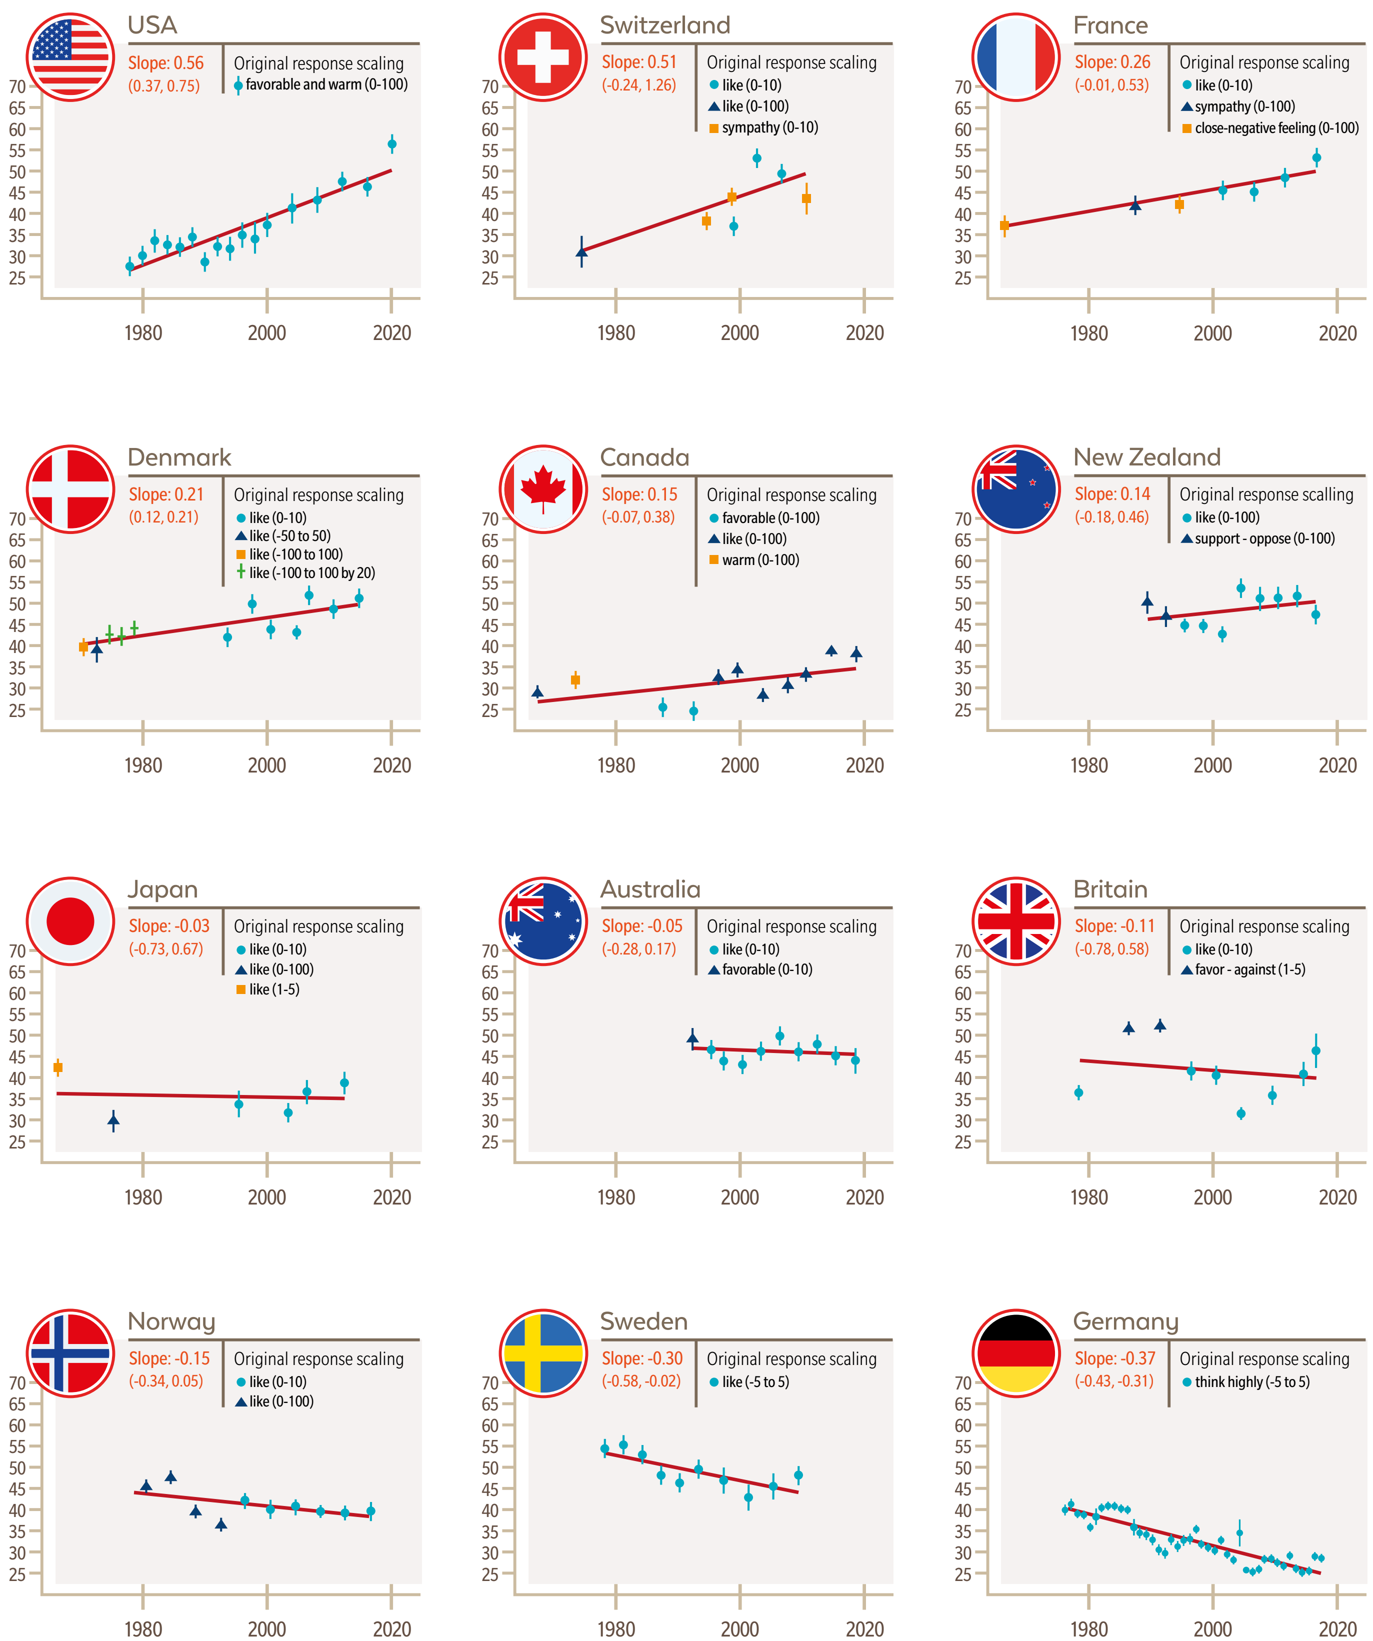

Figure 1 shows trends in affective polarization for the 12 countries in our sample. There is a strong rising trend in the US; smaller positive trends in France, Denmark, Canada, Switzerland, and New Zealand; and declines in Japan, Australia, Britain, Norway, Sweden, and Germany.

> One explanatory factor is increasing racial diversity. The non-white share of the population has risen sharply in the US; and in New Zealand and Canada, where we see rising polarization as well

Measuring the divide

To conduct our analysis, we constructed a new database from 149 different surveys, many of which had to be harmonized manually. These data permit a first look at long-term cross-country trends in affective polarization, but they also have important limitations. Question wording and response scales — as well as the set of years with available survey data — differ across countries and, in some cases, across survey years for a given country. Because the number and nature of political parties differ across countries and within countries over time, even identically structured survey questions may take on different meanings in different contexts. We analyze the sensitivity of our findings to restricting attention to the top two parties in each country and focusing on periods in which this pair of parties is stable.

From each survey, we extract each respondent's party identification as well as a measure of each respondent's affect towards the parties in their country. Questions about affect vary across surveys, commonly asking respondents how they feel towards a given party, how much they like the party, or to what extent they sympathize with the party. We transform the responses in each survey so that the minimum response is 0 and the maximum response is 100. We refer to the transformed response as the respondent's reported affect towards the given party.

We then define our main measure of polarization to be the difference between average affect toward one's own party and average affect toward all other parties, weighted by their shares in the population.

American exceptionalism

In the final part of the paper, we consider potential explanations for rising polarization. We look at the correlation between trends in affective polarization and trends in possible drivers. This analysis cannot conclusively establish what causes polarization, but it can suggest explanations that would be most promising to explore further.

The data do not provide much support for the hypothesis that digital technology is the central driver of affective polarization. The internet has diffused widely in all the countries we look at, and under simple stories where this is the key driver we would have expected polarization to have risen everywhere as well. In our data, neither diffusion of internet nor penetration of digital news are significantly correlated with increasing polarization. Similarly, we find little association with changes in inequality or trade.

One explanatory factor that looks more promising is increasing racial diversity. The non-white share of the population has increased faster in the US than in almost any other country in our sample, and other countries like New Zealand and Canada where it has risen sharply have seen rising polarization as well.

Another potential driver is the nature of the divisions between political parties. The period we study saw important changes in the composition of the parties in the US. Among both political elites and voters, party identification became increasingly aligned with both political ideology and social identities such as race and religion. This was due in part to the political realignment of the South, where conservative whites shifted from the Democratic to the Republican party. The period also saw sharp increases in the polarization of party elites in the US as measured by roll-call voting. Increases in the sorting of parties to ideologies and polarization of elites are associated with rising polarization in our data, the second significantly so.

KEY TAKEAWAYS

- There is a strong rising trend of affective polarization in the US; smaller positive trends in France, Denmark, Canada, Switzerland, and New Zealand; and declines in Japan, Australia, Britain, Norway, Sweden, and Germany.
- The evidence does not support simple stories in which digitalization, inequality, and globalization are the main drivers of polarization.
- Rising polarization appears more likely to be a result of factors that vary across countries, such as changing party coalitions and racial divisions.

The Toulouse Network for Information Technology (TNIT) was created in 2005 to stimulate high-quality economic research on the software industry, the development and role of the Internet, and intellectual property.

Sadly, in 2020, and after 15 very successful years, the network has come to an end. A last issue of *TNIT News* will be published in a few months retracing the history of TNIT and its achievements.

TNIT News

TOULOUSE NETWORK FOR INFORMATION TECHNOLOGY

- *Scientific director:* Jacques Crémer
- *Production editor:* Pascale Maréchal
- *Editorial contributions:* James Nash
- *Graphics:* Olivier Colombe
- *Photo D. Acemoglu:* Adam Glanzman
- *Photo M. Gentzkow:* Viceth Vong
- *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 aims to stimulate world-class research in the economics of information technology, intellectual property, software security, liability, and related topics.

➔ Issue 24 ➔ September 2021

TSE Digital Center

Microsoft