How new Technologies Change Politics

Nikolas Bahati Schöll

TESI DOCTORAL UPF / Year 2021

THESIS SUPERVISOR
Ruben Enikolopov & Aina Gallego
Department Departament d'Economia i Empresa

Universitat
Pompeu Fabra
Barcelona

"Technology is a useful servant but a dangerous master."

– CHRISTIAN LANGE

Nobel Prize for Peace Lecture
December 13, 1921

ACKNOWLEDGEMENT

Looking back, I can say that writing a Ph.D. thesis is much harder than I expected. I can also say that I would have long given up if it was not for a lot of professional and personal support from many great people around me. Far from complete, I want to express my gratitude to you.

First of all, I count myself extremely lucky that I found in Aina, Gaël, and Thomas three amazing coauthors. Working with you and learning from you helped me tremendously on my journey. You gave me the confidence to now say I feel like an actual researcher. Furthermore, without you, this thesis would have never reached the quality and depth it has now. Thanks to our collaboration, I am very proud of the result of my dissertation.

Next, I want to thank my supervisors, Ruben and Aina. Both of you dedicated a lot of time and energy to guiding me through the ups and downs of writing this thesis, and never stopped encouraging and supporting me.

Also very important to me was the wider university network. The students, I had the privilege to teach, colleagues inside and outside my department, and all those inspiring encounters with academics at conferences and elsewhere. Furthermore, I want to express my deepest gratitude to Marta and Laura from the departments' administration. You did an amazing job at smoothing my way through the official steps of writing a Ph.D. and at sending me countless reminders whenever needed. More in the background but may be of greatest importance is the support and love I received from my family. My parents, Mechthild and Volker, and my siblings, Elisa and Clemens. You have always been there for me and never stopped believing in me. Thank you!

Last but not least, I want to thank are all my friends at UPF, my friends, and beloved ones in Barcelona and around the world who had a beer with me and listened to my sorrows whenever needed. To name but a few: Chloé, Milan & Silvia, Laurenz, Ruben, Miriam, Alex & Corinna, Julia, Laura, Simone & Machu, Luca, Ece, Karolis, Sebastién, Cyrill and many more I forgot to mention... All of you helped me feel at home in Barcelona and gave me the energy for this endeavor!

NIKO SCHÖLL
Barcelona/Berlin
November 2021

Abstract

This dissertation investigates how new technologies such as social media, digitalizing workplaces, and robotization change society and politics. The first chapter focuses on new technologies affecting the democratic process itself, more specifically we study political communication on social media. Analyzing the behavior of Spanish politicians on Twitter, we describe how the structure of the social network lets politicians focus more on issues relevant to their own social group. The second and third chapter analyze how political preferences are shaped by new technologies through changing work environments and diverging job opportunities. We show that digitalization and robotization tranform labor markets by affecting existing jobs, but also by changing the type of jobs available. As a second step, we link these changes in career prospects to changes in political preferences and voting behavior. In particular, chapter two studies the economic and political trajectories of British workers directly affected by digitalization. Chapter three studies the indirect effects of digitalization and robotization by comparing the trajectories of different regions in West Germany.

Zusammenfassung

Diese Dissertation unersucht wie neue Technologien wie soziale Medien, Digitalisierung von Arbeitsplätzen und Robotisierung die Gesellschaft und Politik verändern. Das erste Kapitel befasst sich mit neuen Technologien, die den demokratischen Prozess selbst beeinflussen, genauer gesagt untersuchen wir die politische Kommunikation in den sozialen Medien. Wir analysieren das Verhalten spanischer Politiker auf Twitter und beschreiben, wie die Struktur des sozialen Netzwerks Politiker sich stärker auf Themen konzentrieren lässt, die für ihre eigene soziale Gruppe relevant sind. Im zweiten und dritten Kapitel analysieren wir, wie politische Präferenzen durch neue Technologien, veränderte Arbeitsumfelder und auseinanderstrebende Beschäftigungsmöglichkeiten geprägt werden. Wir zeigen, dass Digitalisierung und Robotisierung den Arbeitsmarkt beeinflussen, indem sie bestehende Jobs verändern, aber auch indem sie die Art neugeschaffener Jobs verändern. In einem zweiten Schritt verknüpfen wir diese Veränderungen der Berufsaussichten mit Veränderungen der politischen Präferenzen und des Wahlverhaltens. In Kapitel zwei werden insbesondere die wirtschaftlichen und politischen Entwicklungen britischer Arbeitnehmer untersucht, die direkt von der Digitalisierung betroffen sind. In Kapitel drei werden die indirekten Auswirkungen von Digitalisierung und Robotisierung untersucht, indem die Entwicklung westdeutscher Landkreise verglichen wird.

Resum

Aquesta tesi investiga com les noves tecnologies, com ara les xarxes socials, la digitalització dels llocs de treball i la robotització, canvien la societat i la política. El primer capítol se centra en les noves tecnologies que afecten el propi procés democràtic, més concretament estudiem la comunicació política a les xarxes socials. Analitzant el comportament dels polítics espanyols a Twitter, descrivim com l'estructura de la xarxa social fa que els polítics se centrin més en temes rellevants per al grup social. Al segon i tercer capítol analitzem com les preferències polítiques es veuen modelades per les noves tecnologies a través dels canvis als entorns laborals i la divergència d'oportunitats de treball. Mostrem que la digitalització i la robotització alteren els mercats laborals en afectar els llocs de treball existents, però també en canviar el tipus de feina disponible. Com a segon pas, vinculem aquests canvis en les oportunitats laborals amb els canvis en les preferències polítiques i el comportament de vot. En particular, el capítol dos estudia les trajectòries econòmiques i polítiques dels treballadors britànics directament afectats per la digitalització. El capítol tres estudia els efectes indirectes de la digitalització i de la robotització comparant les trajectòries de diferents regions d'Alemanya Occidental.

Resumen

Esta tesis investiga cómo las nuevas tecnologías, como las redes sociales, la digitalización de los lugares de trabajo y la robotización, cambian la sociedad y la política. El primer capítulo se centra en las nuevas tecnologías que afectan al propio proceso democrático, más concretamente estudiamos la comunicación política en las redes sociales. Analizando el comportamiento de los políticos españoles en Twitter, describimos cómo la estructura de la red social hace que los políticos se centren más en temas relevantes para su propio grupo social. En el segundo y tercer capítulo analizamos cómo las preferencias políticas se ven moldeadas por las nuevas tecnologías a través de los cambios en los entornos laborales y la divergencia de oportunidades de trabajo. Mostramos que la digitalización y la robotización alteran los mercados laborales al afectar a los puestos de trabajo existentes, pero también al cambiar el tipo de empleo disponible. Como segundo paso, vinculamos estos cambios en las oportunidades laborales con los cambios en las preferencias políticas y el comportamiento de voto. En particular, el capítulo dos estudia las trayectorias económicas y políticas de los trabajadores británicos directamente afectados por la digitalización. El capítulo tres estudia los efectos indirectos de la digitalización y la robotización comparando las trayectorias de diferentes regiones de Alemania Occidental.

PREFACE

This dissertation does not contain a grand unifying theory on how new technologies affect politics - but rather a relatively loose collection of efforts to push the research frontier on various sides. Yet, I do not see this as a short-coming. In my understanding, science is an incremental process, and I count myself very lucky that I had the privilege to make these contributions to the scientific community.

Nevertheless, I will take advantage of these introductory words to reflect upon my research and if there is a unifying element. The three chapters of my dissertation seem mainly connected by the ambition to test specific narratives currently debated by academics and the public at large - narratives about how new technologies shape society and politics. In practice, I touch upon social media, digitalization and robotization as new technologies. These technologies have a tremendous impact on modern society. Still, it has to be admitted that quantitative research (and this includes my work) is sometimes biased towards technologies with high-quality data.

Now, before I start discussing the content of this dissertation, let me briefly reflect on the role of technology in human history more generally. It is widely accepted that technology is the foundation of the unprecedented prosperity we enjoy today, especially in the global north.

At the same time, technology became so powerful that it now threatens to destroy the very basis of our lives. Military technology such as nuclear weapons, but also anthropogenic climate change caused by a fossil fuel-based economy, acutely endanger the future of our planet. What is more, new technologies do not only threaten the natural environment around us, but also pose risks to the functioning of our society as such. On the one hand, communication technology shapes how we relate to each other and consequentially how societal debates and politics play out. Some see in it a tool that allows everybody to participate in political decision-making processes. Others point to the risk of alienation in a hyper-individualized society, and the threat that filter bubbles and fake news pose to our democratic system. In particular, the advent of social media is often seen as a threat to Western democracy, as it allegedly polarizes and manipulates public opinion. Social media abuses have been blamed for a wide range of events such as the rise of far-right populist forces, the election of Donald Trump, the Brexit vote, and conspiracy theories around Covid-19.

On the other hand, how we earn our living is determined by the way the economy works, which in turn is heavily shaped by technology. While fears of a jobless future (with robots and algorithms doing all the work) have not materialized yet, the relative importance of different societal groups did change tremendously. The focal group of society changed from the archetypal male bread winner working in manufacturing to a diverse group of people working in the knowledge economy. Furthermore, technological change is considered to be one of the major drivers of ever-increasing economic inequality in the Western world, as new technologies lead to a 'hollowing out' of the middle class. Here again, doubts have been raised on how we can adjust to the ever-increasing speed of technological change.

Even though I think that all these processes deserve our attention, it might not come as a surprise that I, as a political economist, focused more on the latter, the societal aspects of new technologies. In particular, I wanted to understand if and how social cohesion and our democracy might be threatened by new technologies. The remainder of this preface summarizes the three chapters of my dissertation.

The first chapter, "How politicians learn from citizens' feedback: the case of gender on Twitter" is joint work with Aina Gallego and Gaël Le Mens. Each of us brought in the perspective of a different discipline. The intersection of political theory of representation, a learning model from psychology to describe how politicians process information from social media and econometric techniques to empirically test our hypotheses (and a considerable amount of data science) allowed us to study how politicians are affected by social media in an innovative way. As Aristotle already stated: "The whole is greater than the sum of the parts." We asked the following questions: Does feedback from citizens on social media affect the issues that politicians choose to discuss? Are politicians of different social categories exposed to different feedback? Does this affect their issue attention? To answer those questions, we use a reinforcement learning framework from psychology to model how politicians choose which policy issues to address. The model predicts that politicians respond to citizen feedback by increasing attention to issues that received more feedback. Furthermore, we hypothesized that citizens provide more positive feedback to female politicians for writing about gender and that this contributes to their specialization in gender issues. These predictions were confirmed in analyses of 1.5 million tweets published by Spanish MPs over three years. We identified gender issue tweets using a deep learning algorithm (BERT), and measured citizen feedback using the number of retweets and likes. To conclude, we discuss how reinforcement learning generates responsiveness, but can also be the cause of unequal representation, misperceptions, and polarization.

This constitutes a case where a new technology (social media) affected the democratic process as such. In this specific case, it is about changes in the behavior of politician, which can be interpreted as the supply side of politics.

The second and third chapter instead focus on the demand side, i.e. the political preferences of citizens and how they are affected by new technologies. This happens through technology-induced changes in work environments and diverging career perspectives. We show that computers and robots have the potential to disrupt labor markets and to change the composition of the workforce, especially by replacing routine work oftentimes performed by members of the middle class. As a second step, we then analyze how these technology-induced changes in career prospects translate into changes in political preferences and voting behavior.

The second chapter called "Neither Left-Behind nor Superstar: Ordinary Winners of Digitalization at the Ballot Box" is joint work with Aina Gallego and Thomas Kurer. It departs from the observation that the previous literature on the political consequences of technological change studies either left-behind voters or extremely successful technology entrepreneurs ("superstars"). However, a large share of skilled workers who benefit from limited but steady economic improvements in the knowledge economy had been ignored in previous analyses. This chapter fills a lacuna by examining how workplace digitalization affects political preferences of given individual workers among the entire active labor force. To do so, we combined individual-level panel data from the United Kingdom with industry-level data on ICT capital stocks between 1997-2017. We first demonstrate that digitalization was economically beneficial for workers with middle and high levels of education. We then show that growth in digitalization increased support for the Conservative Party, the incumbent party, and voter turnout among beneficiaries of economic change. Our results hold in an instrumental variable analysis and multiple robustness checks. While digitalization undoubtedly produces losers (along with some superstars), ordinary winners of digitalization are an important stabilizing force content with the political status quo.

These results go against a popular narrative that technological change first and foremost result in political disruption. Furthermore, while writing the second chapter, we became aware of the pros and cons of studying the economic and political trajectories of individual workers. On the positive side, studying how changing careers translated into changing political values and voting behavior as within-individual changes has huge benefits as it allows us to abstract from outside factors such as the social milieu that the individual belongs to. This is extremely important as socialization plays a huge role for political preferences. On the other hand, it necessarily made us ignore changes across individuals related to labor market turnover. Technology affects the type of jobs available to new labor market entrants, and therefore affects educational and occupational choices. As a consequence, new generations will hold different political values than their predecessors.

These considerations made us ask ourselves if there was an additional mechanism at play that we could not grasp with our approach? Was it rather labor market outsiders that fueled political disruption? To answer these questions, my coauthor Thomas Kurer and I decided to conduct a complementary study which considered the electoral behavior of entire regions rather than individuals. This way, we would make sure that we do capture generational turnover as well as those not attached to the labor market.

The third chapter called "How technological change affects regional electorates" is the result. Drawing on fine-grained labor market data from Germany, we first show that the well-known decline in manufacturing and routine jobs in regions with higher robot adoption or investment in information and communication technology (ICT) was more than compensated by parallel employment growth in the service sector and cognitive non-routine occupations. This change in the regional composition of the workforce has important political implications: Workers trained for these new sectors typically hold progressive political values and support progressive pro-system parties. Overall, this composition effect dominates the politically perilous direct effect of automation-induced substitution. As a result, we conclude that technology adopting regions are unlikely to turn into populist-authoritarian strongholds.

Taken together, chapter two and three offer a relatively optimistic account of how new technologies affect society and politics. It complements (but not necessarily contradicts) more gloomy perspectives oftentimes voiced in both academic and public debates.

To finish, I hope this dissertation as a whole contributes to creating a more evidence-based, and more nuanced debate about how new technologies affect society and politics and ultimately how we decide how to face the societal challenge of new technologies.

Contents

1

HOW POLITICIANS LEARN FROM CITIZENS' FEEDBACK: THE CASE OF GENDER ON TWITTER

Joint with Aina Gallego (Universitat de Barcelona) and Gaël Le Mens (Universitat Pompeu Fabra)

1.1 Introduction

Does feedback from citizens on social media affect the issues that politicians choose to discuss? Recent research on issue responsiveness finds that when an issue becomes salient among citizens on social media, politicians quickly follow and become more likely to discuss it over the next days (Barberá et al., 2019). This finding raises the question of how politicians can learn and respond so rapidly to changes in public mood. In traditional dynamic representation models, (Stimson et al., 1995; Soroka and Wlezien, 2010; Wlezien, 1995) public policy adjusts to shifts in aggregate public opinion over much longer periods of time, typically years, and politicians learn about changes in public opinion through tools that require careful analysis, such as opinion polls (Druckman and Jacobs, 2006), expert consensus (Stimson et al., 1995), or by recording and analyzing information (Henderson et al., 2021). These approaches to detecting changes in public opinion do not seem applicable to the social media setting because they are impractical in settings in which new information is highly decentralized and spreads in minutes (Cagé et al., 2020). While online information is abundant, unbiased and up-to-date summaries about which issues are relevant for citizens are not available. The strategies that politicians use in order to be responsive to citizens on social media thus remain unspecified to date.

This article focuses on how politicians use one source of information about the preferences of citizens that is continuously available on social media: feedback from citizens. We study how feedback affects subsequent decisions about which political issues to discuss. Contact with politicians has long been considered a relevant way in which citizens can influence politicians' issue agendas (Miller and Stokes, 1963; Fenno, 1977) but has been difficult to measure. Studying the impact of feedback from citizens is timely as the interactive features of social media have reduced the cost of two-way communication between politicians and citizens and made it more abundant (Jungherr et al., 2020). In this article, we specify the process through which politicians respond to citizen feedback in terms of a 'reinforcement learning' model grounded in research about how people learn from feedback (Holland, 1992; Sutton and Barto, 2018). We propose that after talking about an issue, politicians observe the amount of positive feedback from citizens, update their perceptions about the popularity of the issue, and respond by increasing attention to popular issues and decreasing attention to unpopular issues.

This simple strategy allows politicians to be continuously responsive, but only to the self-selected citizens who interact with them. A relevant characteristic of social media is that users, including politicians, are exposed to information environments that tend to match and possibly reinforce their preexisting views (Sunstein, 2018; Zhuravskaya et al., 2020). To study the implications of exposure to fragmented audiences, our reinforcement learning model allows politicians of different social categories to be exposed to systematically different feedback from citizens. Specifically, we focus on the social category of gender – whether a politician is a female or a male – and the extent to which female and male politicians attend to gender issues. The model shows that, if female politicians receive more positive feedback for talking about gender as compared to male politicians, reinforcement learning creates a difference in attention to gender issues between female and male politicians. The model is general and can apply to other issues and to other social categories such as race or partisanship. It can also apply to offline settings.

We test the theory with rich social media data that record politician-citizen interactions over time and allow longitudinal analysis. We collected 1.5 million tweets published by elected representatives in national and regional assemblies, active during the 2016 to 2019 election cycle in Spain. We measured the reception of each tweet in terms of 'retweets' and 'likes' and use these data to estimate politicians' responsiveness to feedback. To code gender issues, we rely on 'BERT' (Devlin et al., 2018), a deep learning language model which is sensitive to word dependencies, vastly outperforms standard bag-of-word models, and works well in multi-lingual contexts. We estimate the effect of citizen feedback on attention to gender issues by female and male politicians using two-way fixed effect panel models, which allows us to control for all factors that are constant for a given politician or for a given point in time.

We find that politicians are responsive to citizen feedback on social media: after receiving more retweets for tweeting on gender issues, they increase attention to this issue. This is also the case with 'likes.' Moreover, we find that female and male politicians are exposed to systematically different feedback environments: female politicians receive relatively more retweets and likes for tweeting about gender issues. This leads them to talk more about gender issues. Our analyses of mechanisms also reveal that female politicians obtain more feedback because they are treated differently by citizens, and not because their messages are more engaging.

Our study advances research on how politicians respond to changes in public opinion. It is most related to Barberá et al. (2019), who document issue responsiveness on social media, but do not study the underlying mechanisms. More generally, theoretical models of dynamic representation remain unclear about how politicians learn about public opinion (e.g. Stimson et al., 1995). We propose and test one specific learning process that allows politicians to be continuously responsive to the citizens with whom they interact. Methodologically, we develop an empirical approach that allows the analysis of actual interactions between politicians and citizens on social media, instead of relying on inferences from population-wide averages. Substantively, we document, for the first time, that the direct interactions between politicians and citizens influence the issues that politicians choose to discuss on social media and show that differential treatment from citizens leads politicians with different characteristics to diverge in issue attention.

We also contribute to the large literature on the political representation of women by connecting the gender-specific experiences of women in office to the rise of attention to gender issues. Theoretical work about descriptive representation argues that female representatives are more likely to talk about issues relevant to women because they have different experiences both in life and in office (Mansbridge, 1999; Phillips, 1995). Empirical research supports the claim that descriptive representation increases substantive representation (Lowande et al., 2019; Lawless, 2015; Wängnerud, 2009; Clayton, 2021), but has more difficulties at examining specific mechanisms that link both. In particular, existing empirical studies have not been able to isolate the effects of specific gendered experiences in office on politicians' attention to gender issues. We demonstrate that exposure to systematically different feedback environments contributes to differences in attention to gender issues between female and male politicians beyond what can be explained by differences in intrinsic motivation or pre-existing preferences. Furthermore, our analyses of mechanisms shed light on why this happens.

1.2 Learning from feedback on social media

In order to be responsive to citizens, politicians first need to learn about citizens' preferences both in terms of issue salience and issue position.[1] How do politicians learn about the preferences of the public? Dynamic representation theory (Stimson et al., 1995) provides one answer to this question. While recognizing that politicians cannot directly know the preferences of the public, this theory proposes that all politicians have access to a "consensus view" about the direction of change in preferences which is produced by a community of opinion leaders, including politicians, journalists, and academics. In a similar spirit, thermostatic models of public opinion (Wlezien, 1995, 2004) assume that politicians are aware of directional changes in aggregate public opinion.

The assumption that all politicians have access to a shared view about the preferences of the public may be well-suited to homogeneous information environments, as was the case when a few broadcast television channels were dominant and thus everyone was exposed to similar information (Prior, 2007). However, the low barriers to entry and the reliance on user-generated content have made online information environments much more fragmented than traditional media environments (Zhuravskaya et al., 2020). Politicians, like other users, are not exposed to content that is centrally produced by gate-keepers and similar for all users, but to content that depends on whom they choose to follow, which users choose to interact with them, and on how algorithms prioritize information. An additional challenge is that new topics appear and disseminate online at a very high speed (Cagé et al., 2020). This reduces the usefulness of tools such as traditional opinion polls to track changes in public opinion. Since social media platforms do not provide systematic information about the average preferences of citizens on political issues, politicians must find other approaches to learn about them.

Research on how representation operates in practice finds that when politicians (or their aides) try to learn about citizens' preferences, they rarely use tools like surveys, which are often not available. Instead, they keep track of their communication with interest groups and regular citizens and make inferences based on this information (Miller and Stokes, 1963; Fenno, 1977; Henderson et al., 2021).

[1] We focus on issue salience in this paper because the decision to talk about an issue or not is binary, and this facilitates empirical analysis. However, the logic applies to issue position as well.

Information obtained through direct interactions with citizens, and in particular the feedback they provide, is particularly relevant in social media contexts because it is abundant, immediately available, and easy to use. Before the advent of social media, citizens communicated their opinions to politicians through actions such as writing letters or talking in public meetings which require civic skills and are relatively costly (Verba et al., 1995). The built-in interactive features of social media, such as the ability to provide feedback to other users through easily clickable buttons, have made two-way communication between citizens and politicians easier and more convenient (Jungherr et al., 2020). Moreover, feedback is obtained in real time. As Zhuravskaya et al. (2020, p. 417,) note, "Social media allows politicians to receive immediate feedback on policy actions, to discuss policy proposals, and to measure political discontent." Finally, feedback in social media is more easily usable than traditional communication with constituents because it comes in a highly standardized quantitative form (such as the number of retweets, likes or hearts), which makes it easy to compare how different statements fare. Thus, we expect that politicians use the feedback they obtain on social media to make inferences about citizens' preferences.

How do politicians use feedback? To address this question, we assume that when making decisions about which issues to discuss and which positions to take, politicians aim to choose popular topics and positions. This could be because they believe that consistently doing so will increase support for themselves or their parties or because they see themselves as delegates of the public.[2] However, politicians are uncertain about the popularity of the issues they might discuss.

We propose that politicians learn about the popularity of issues by observing how their messages are received by the public and that they increase attention to issues that obtain more positive feedback than expected and reduce attention to those that obtain less positive feedback than expected. In short, issues that obtain relatively more positive feedback are 'reinforced.' Prior research has shown that people frequently behave this way when they make repeated choices between options with uncertain payoffs and aim to obtain positive payoffs (Denrell, 2005; Thorndike, 1927) and that this behavior is often reasonable (Holland, 1992; Le Mens and Denrell, 2011; Sutton and Barto, 2018). In the context of politicians writing on Twitter, the options consist in different political issues which they can choose to discuss in their next tweet. Feedback is the reaction of the citizens to

[2] In some conceptions of representation, such as gyroscopic or trustee representation (Mansbridge, 2003), politicians do not need to be responsive to represent the public. We recognize that politicians sometimes deviate from public opinion, but we assume that in general they are motivated to be responsive to citizens, as suggested by recent research which demonstrates that politicians change their votes when they receive information about the preferences of voters (see Butler et al., 2011; Pereira, 2021).

the politicians' tweets, which can be more positive or negative than expected. Politicians are responsive to feedback if they tend to choose issues that obtained positive feedback in the past and hence are perceived as more popular. We analyze a formal model of this learning-from-feedback process in Section 1.3.1 and provide empirical estimates of the model parameters in Section 1.5.2.

A key drawback of relying on feedback as a source of information is that citizens who provide it are self-selected and politicians cannot know in which way the preferences of their followers differ from the preferences of the population at large (or of other relevant groups, such as copartisans or voters in their districts). While politicians and their staff are aware that their online followers are not representative of the public (Henderson et al., 2021), they have no way to fully correct the ensuing biases.[3]

If politicians of different social categories, such as gender or race, are exposed to more positive feedback from the public when they talk about issues related to their social categories, they will form different perceptions of what the public wants and will ultimately be more likely to talk about issues related to their social category. Our study focuses on gender, which is a more politically relevant characteristic than race in the Spanish context. We expect that female politicians receive relatively more positive feedback from citizens when they talk about gender issues rather than on other issues – a difference in feedback that, from now on, we call the 'gender issue feedback advantage'. There are several reasons why the gender issue feedback advantage would be larger for female politicians than for male politicians.

First, female politicians may communicate more engagingly about gender issues because they are more knowledgeable and interested in these issues (Dolan, 2011; Lowande et al., 2019; Lawless, 2015), and this more engaging style may in turn generate more positive reactions from citizens. There exists abundant evidence that female representatives have different positions on gender issues than male representatives (e.g. Lovenduski and Norris, 2003), although whether female politicians communicate more engagingly about gender has not been rigorously assessed. We call this mechanism the 'engagingness channel.'

Second, female citizens may interact more with female politicians. This argument has been advanced most clearly by Mansbridge (1999, p. 641) who argues that politicians of traditionally marginalized groups provide better representation to in-group members because they have 'enhanced communication' with them. Empirical research finds that citizens are more likely to contact politicians of their race (Broockman, 2014; Gay, 2007), although there

[3]There exists evidence that when producing population estimates, people go beyond the information they obtain from their immediate social environments, yet they do not fully correct for the biases already present in their information sample (Galesic et al. (2018), see also Fiedler (2012)).

is less direct evidence about gender (for null results see for instance Bush and Prather, 2020; Haynes, 1997). If female citizens ‘self-select’ into interacting more with female politicians and female users are more likely to give feedback to tweets on gender issues, this could potentially explain the gender issue feedback advantage. We call this mechanism the ‘self-selection channel’.

Third, citizens may believe that female politicians are more competent to talk about gender issues (Huddy and Terkildsen, 1993; Dolan, 2010) and, for this reason, may provide them with more positive feedback for tweeting on the topic even if there is no difference in the content of the gender issue tweets written by female and male politicians. Recent research finds that partisanship or incumbency dominate gender stereotypes when citizens decide for which candidate to vote (Dolan, 2014; Lawless, 2015). But this does not rule out that voters reward female politicians for behaving according to stereotypes in social media contexts, where voters are not restricted in the amount of feedback they can provide and thus do not need to prioritize one consideration over others. Research in social psychology and sociology in general supports the claim that people tend to evaluate the behavior of others more positively if it is congruent with expectations related to their social categories (Eagly et al., 2000; Hannan et al., 2019). If Twitter users expect female politicians to talk more about gender issues they may react more positively when they do, because this is *congruent* with their expectations regarding the issues female politicians should attend to. These arguments imply that citizens are more likely to retweet tweets on gender issue when they are published by female politicians, rather than male politicians, even if there is no difference in tweet content. We call this mechanism the ‘congruity channel.’

We empirically test for the differences in gender issue feedback advantage feedback in Section 1.5.1 and test for the three potential mechanisms in Section 1.5.3.

1.3 Model

1.3.1 Reinforcement learning by an individual politician

Consider a politician i who publishes a series of messages on policy issues. Without loss of generality, we assume that there are only ‘gender issues’ and ‘other issues’ and denote them by GI and $other$. We refer to the first message by $m = 1$, the second message by $m = 2$, etc. In reinforcement learning models, agents have latent ‘valuations’ of each option, which they update based on feedback. The valuation of different policy issues can be interpreted as politicians’ perception of the popularity of that issue. Politician i’s valuation of the 'gender issues' option at the time they decide on the issue of message m is $V_{i,m}^{GI}$ and the valuation of the 'other issues' option is $V_{i,m}^{other}$. The politician is more likely to choose 'gender issues' if the difference in valuations favors this issue, i.e. they perceive it as more popular. We specify the probability that the politician chooses issue k as a logistic function of the difference in valuations of the two issues. We call this quantity the 'attention to the gender issue':

$$A_{i,m}^{GI} = Logit(\pi_i^{GI} + r\Delta V_{i,m}), \tag{1.1}$$

where $\Delta V_{i,m} = V_{i,m}^{GI} - V_{i,m}^{other}$ is the valuation difference, r denotes the responsiveness of issue attention to perceived popularity, and π_i^{GI} characterizes the baseline tendency to write about gender issues. This latter construct can be thought of as the intrinsic motivation to address the issue.

We denote by $V_{i,1}^{GI}$ and $V_{i,1}^{other}$ the initial valuations of the two issues. After every message m, the politician observes the feedback $FB_{i,m}^{k}$ and updates their valuation of the issue of the message. Following research on how people update valuation based on experience (see Denrell (2005) for a review), we assume that the new valuation of an issue is a weighted average of the previous valuation of that issue and the last feedback instance on that issue (see Appendix A.1.2 for a discussion of this assumption). Formally, if message m is on issue k, then

$$V_{i,m+1}^{k} = (1-\gamma)V_{i,m}^{k} + \gamma FB_{i,m}^{k}. \tag{1.2}$$

If message m is not on issue k, the valuation of issue k does not change: $V_{i,m+1}^{k} = V_{i,m}^{k}$.

We assume that feedback is normally distributed, with common standard deviation σ, and with means μ_i^{GI} and μ_i^{other} that differ between issues:

$$FB_{i,m}^{GI} \sim N(\mu_i^{GI}, \sigma); \qquad FB_{i,m}^{other} \sim N(\mu_i^{other}, \sigma)$$

It is possible to derive a formula for the long-run share of attention to gender issues, A_{∞}^{GI} (see proof in Appendix A.1.1).

$$A_{\infty}^{GI} = Logit(\pi_i^{GI} + r\Delta\mu_i), \tag{1.3}$$

where $\Delta\mu_i = \mu_i^{GI} - \mu_i^{other}$ is the difference between the means of the feedback distributions for the two issues ('gender' and 'other'). This corresponds to what we call the 'gender issue feedback advantage'. Unsurprisingly, the long-run attention to gender issues increases with the gender issue feedback advantage. This feedback effect is stronger when the issue responsiveness parameter, r, is larger. It is noteworthy that the long-run attention to gender issue does not depend on the initial valuations. This means that our main result holds whether the politician initially believes average feedback for the two issues to be the same or different (see Appendix A.1.2 for further discussion of this).

1.3.2 Differences between female and male politicians

Now consider two hypothetical politicians, F and M who behave according to the reinforcement learning model but are exposed to different feedback environments such that the gender issue feedback advantage differs between the two politicians ($\Delta\mu_F \neq \Delta\mu_M$). Using equation 1.3, we can derive a necessary and sufficient condition for a difference in long-run issue attention such that attention to the gender issue is larger for F than for M:

$$A^{GI}_{F,\infty} > A^{GI}_{M,\infty} \iff \pi^{GI}_F + r_F\Delta\mu_F > \pi^{GI}_M + r_M\Delta\mu_M. \tag{1.4}$$

This difference in issue attention can emerge as the result of a difference in the feedback received by F and M.

A feedback-driven difference in valuations and issue attention can emerge even if F and M have identical baseline propensities for publishing tweets on gender issues ($\pi^{GI}_F = \pi^{GI}_M$) and are equally responsive to changes in issue valuations ($r_F = r_M$). In this case, politician F will devote a larger attention to the gender issue whenever the gender issue feedback advantage is stronger for F than for M ($\Delta\mu_F > \Delta\mu_M$). We discuss model dynamics for different values of the initial valuations in Appendix A.1.3.

In the general case, feedback contributes to the difference in issue attention between politicians F and M beyond what could be explained just by a difference in baseline propensities to write about gender issues when the following condition holds:

$$r_F\Delta\mu_F > r_M\Delta\mu_M. \tag{1.5}$$

We test whether the condition in equation 1.5 holds in Section 1.5.2.

1.4 Case, data and measurement

To analyze whether and how citizen feedback affects politicians' issue attention, we collected the tweets published by all politicians who served in the national parliament of Spain or any of its regional parliaments between the start and the end of the national legislature (from July 2016 to March, 2019).

Spain is a relevant case to study the rise of gender issues. Gender evolved from being a relatively niche issue into a major topic during the time covered by our study, culminating in a general strike in March 2018, which was probably the largest women's strike in history (Campillo, 2019). Spain is a fairly typical consolidated democracy. It has a proportional representation system and closed party lists. It is also a decentralized state, with regional governments holding significant powers. Therefore, both national and regional representatives are relevant for the political process. Social media use is high. We collected the

Twitter user names of 1530 parliamentarians. More than 80% of the politicians who were in office for some time during this period had a Twitter account. They posted more than 1.5 million original tweets in this period.

The set of 'original' tweets consists of tweets politicians posted on their own wall and replies to other users' tweets. We included all tweets with at least two words published by politicians who were active Twitter users (writing on average at least one original tweet per month). We only consider the first tweet of a thread of tweets. The resulting data contains the tweets of 1265 politicians (554 females and 711 males).

In comparison to male politicians, female politicians were less active and their tweets received fewer retweets and likes (Table 1.1). Additional summary statistics are reported in Appendix A.3.

Table 1.1: Female politicians post fewer tweets than male politicians and receive fewer retweets, likes or replies.

	Female politicians	Male politicians
Number of politicians	554	711
Number of tweets (mean)	1087.9	1380.1
Number of tweets (median)	568	697
Average number of retweets (mean)	22.4	45.8
Average number of retweets (median)	6.3	7.4
Average number of likes (mean)	38.6	80.2
Average number of likes (median)	8.4	10.4
Average number of replies (mean)	3.4	7.5
Average number of replies (median)	0.5	0.7
Standard deviation retweets (mean)	52.1	95.9
Standard deviation retweets (median)	10.2	12.7
Standard deviation likes (mean)	38.60	80.20
Standard deviation likes (median)	8.40	10.40
Standard deviation replies (mean)	10.9	22.0
Standard deviation replies (median)	1.5	2.0

Note: We first calculate average values per politician and then the mean or median value of those averages for female and male politicians.

1.4.1 Measuring attention to gender issues

The main empirical challenge consisted of identifying tweets related to gender issues. We used human-coded data to train and validate a text classifier based on a state-of-the-art deep learning language model, BERT (Devlin et al., 2018). This consists of an artificial neural network with many layers (a 'deep neural network') that takes the text of a tweet as an input and labels it as being about gender or not. We chose this model, because it has been shown to perform much better than 'bag-of-words' classifiers which are most often used in the social sciences (Grimmer and Stewart, 2013). We recruited research assistants to code about twenty thousand tweets as being on gender issues or not and we fine-tuned our BERT classifier to optimize its classification performance on our data. We used 10-fold cross validation to identify the optimal training parameters.

Our model achieved an excellent classification performance on our validation data: 90% of the tweets the model classified as gender issue tweets are actually on gender issues and 79% of gender issue tweets are classified as such. For comparison with the more traditional 'bag-of-words' approach, we trained a naïve Bayes classifier. It produced three times more mistakes than our BERT classifier. We discuss the advantages of BERT in Appendix A.2.1, coding details in A.2.2, how we fine-tuned the model in A.2.3, and model accuracy in A.2.4.

We define politician i's attention to gender issues in period p as the proportion of gender issue tweets posted by this politician over that period: $A_{ip}^{GI} = \frac{n_{ip}^{GI}}{N_{ip}}$.

There exists a large difference in attention to gender issues by female and male politicians. Over the entire sample period, female politicians devoted, on average, 11.2% of their tweets to gender issues whereas male politicians only devoted 3.4% of their tweets to gender issues. Figure 1.1 depicts the average attention to gender issues by female and male politicians over the period studied.

Figure 1.1: Attention to gender issues by female and male politicians

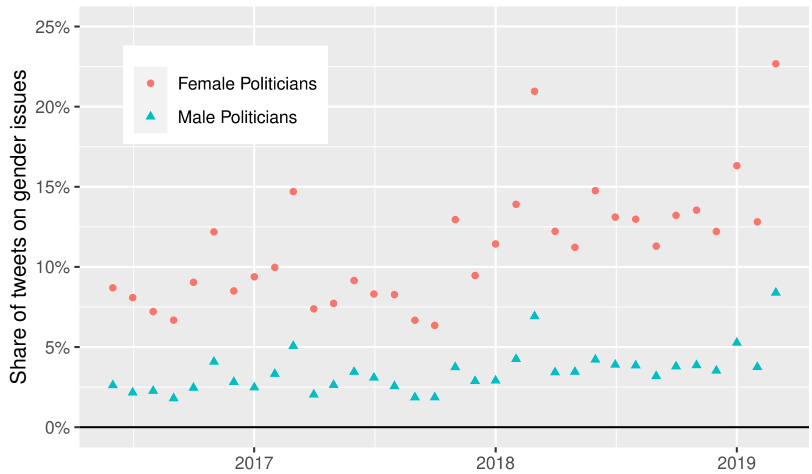

Note: Points represent monthly averages.

Comparing the mean number of raw retweets that each politician received for tweets on gender issues and other issues reveals the existence of gender issue retweet advantage for female politicians (see Table 1.2). Tweets on gender issues written by female politicians receive on average 18% more retweets in absolute terms than tweets on other issues. By contrast, male politicians receive about the same number of retweets for tweeting about gender issues and other issues. A similar asymmetry between female and male politicians holds for likes.

Table 1.2: Summary Statistics: Retweets, likes and replies.

	Female Politicians (N= 554)			Male Politicians (N=711)		
	GI	other	$\Delta^{GI/other}$	*GI*	other	$\Delta^{GI/other}$
Number of tweets (mean)	123.5	966.4		49.6	1332.5	
Number of tweets (median)	46	506		23	683	
Average number of retweets (mean)	25.5	21.6	18%	45.5	45.7	-0%
Average number of retweets (median)	7.4	6.2	19%	7.4	7.4	0%
Average number of likes (mean)	45.4	37.2	22%	83.9	79.9	5%
Average number of likes (median)	9.4	8.3	13%	10.1	10.3	-2%
Average number of replies (mean)	3.9	3.3	18%	7.4	7.5	-1%
Average number of replies (median)	0.44	0.50	-12%	0.55	0.73	-25%

Note: To aggregate the data, we first calculate average values per politician and then the mean or median value of those averages for female and male politicians.

1.4.2 Measuring issue-specific feedback

We construct our main measure of citizen feedback based on the number of *retweets*. Prior research has shown that a higher number of retweets implies approval (Metaxas et al., 2015). Consistent with the view that most retweets are instances of positive feedback, we observe in our data that most of the retweets between politicians happen within parties, (see Appendix A.4). Rather than using the raw number of retweets as the measure of feedback to politician i about the tweet message m they published, we construct a feedback measure grounded in behavioral research on how past experience affects future decisions. We proceed in several steps.

First, we take the natural logarithm of the number of retweets. This transformation is motivated by research that shows that payoffs have declining marginal effects (Tversky and Kahneman, 1992). Taking the logarithm also reduces the weight of instances of extremely large numbers of retweets which have the potential to drive the model estimation results.[4] Differences in logs express scale-invariant ratios of feedback, implying that the added utility of receiving 10% more *retweets* would be the same for a politician who usually receives 10 or 10,000 retweets.

[4]The number of retweets is strongly skewed and approximately follows a power-law distribution for each politician: The median tweet received 4 retweets, the mean is 58 retweets and the maximum is almost 43.000 retweets.

Second, we take out a politician-specific time trend.[5] This step is motivated by research on learning from feedback that has shown that agents tend to evaluate outcomes with respect to a time-dependent 'aspiration level' or reference point (Cyert and March, 1963; March and Shapira, 1992). In our context, the average number of retweets increases over time for most politicians, probably because the politicians' followership is growing. Thus, comparing the number of retweets received by tweets published many months or several years apart is not meaningful.

Finally, we proceed to within-politician z-score standardization. The relevant comparison for a given politician to learn about issue popularity is to compare the number of retweets they received for tweeting on a specific issue with the average level of retweets they receive themselves, rather than the number of retweets that other politicians received.[6] By construction, the distribution of feedback for each politician now has mean zero ($E[FB_{i,m}] = 0$) and standard deviation one ($\sigma_{FB_{i,m}} = 1$).

Our feedback measure can be interpreted as follows: a one-unit increase in feedback means that the tweet received one standard deviation more in 'feedback utility units' relative to other tweets published by the same politician around the same point in time.

We focus on retweets over likes because information about retweeters is more easily available on Twitter than information about those who gave likes, and we use this information in some analyses. Results are similar with likes (Appendix A.5.3). In ancillary analyses we also analyzed replies (Appendix A.5.4).

1.5 Results

1.5.1 Gender issue feedback advantage for female and male politicians

We estimate by OLS a set of linear models with feedback as the dependent variable. In our baseline specification, the feedback received by politician i for tweet message m, $FB_{i,m}$, is regressed on politician gender and the issue of the

[5]We regress $log\,\text{retweets}_{i,m}$ on the time t the tweet was posted using OLS and then take the residual:

$$\widehat{u_{i,m}} = log\,\text{retweets}_{i,m} - \widehat{trend}(log\,\text{retweets}_i) * t \tag{1.6}$$

[6]As a robustness check, we replicate our main analyses by omitting within-politician normalization in A.5.2 in the SI. Our main results remain.

tweet:

$$FB_{i,m} = \beta_{GI} GI_{i,m} + \beta_M M_i + \beta_{GI*M} GI_{i,m} * M_i + \epsilon_{i,m}, \quad (1.7)$$

where $GI_{i,m}$ is a dummy variable equal to 1 if tweet m published by politician i is on gender issue, M_i is a dummy equal to 1 if politician i is male and $\epsilon_{i,m}$ is an error term.

We are most interested in the coefficient of the interaction term, β_{GI*M}, which measures how the gender issue feedback advantage differs between female and male politicians. If it is negative, the gender issue feedback advantage is stronger for female politicians. In most specifications, we include politician fixed effects to absorb the effect of politician characteristics which remain constant over time such as their gender, specialization of policy area, or political party. We also add day and hour of the day fixed effects to absorb the effect of temporal variations affecting all politicians such as general shifts in issue salience.

Estimation results are reported in Table 1.3. In all specifications, the gender issue feedback advantage is stronger for female politicians ($\beta_{GI*M} < 0$, $p < .01$ in Model 1-3, $p < .05$ in Model 4). Model 1 is a basic specification without controls or fixed-effects. We find that the gender issue feedback advantage is larger for female politicians (+0.23 standard deviation) than for male politicians (+0.14 standard deviations). The pattern remains similar when politician and day fixed effects are included (Model 2) as well as when additional time-varying control variables are included, such as the hour of the day the tweet was published, the number of tweets published by the politician on that day, and the length of the thread of the tweet (Model 3). Model 4 shows that the effect is similar for left-wing and right-wing politicians (see Appendix A.5.1 for details on coding).

Appendix A.6.1 reports the robustness checks.

Table 1.3: Linear regressions of tweet feedback on politicians' gender and issue of the tweet (eq. 1.7)

Dependent variable:	Tweet-level standardized feedback, FB_{im}			
Model:	(1)	(2)	(3)	(4)
GI	0.2359***	0.2860***	0.2834***	0.2682***
	(0.0041)	(0.0203)	(0.0200)	(0.0513)
GI × Male politician	-0.0904***	-0.1221***	-0.1250***	-0.1228**
	(0.0068)	(0.0235)	(0.0232)	(0.0550)
Part of thread			0.2920*	0.2921*
			(0.1514)	(0.1513)
Tweets on day by politician			-0.0072***	-0.0072***
			(0.0015)	(0.0015)
GI × Left				0.0199
				(0.0551)
GI × Male politician × Left				0.0009
				(0.0606)
Male politician	0.0213***			
	(0.0017)			
(Intercept)	-0.0264***			
	(0.0014)			
Fixed-effects				
Politician		Yes	Yes	Yes
Day		Yes	Yes	Yes
Hour of day			Yes	Yes
Fit statistics				
Squared Correlation	0.003	0.008	0.018	0.018
Observations	1,583,917	1,583,917	1,583,917	1,583,917

Note: Estimations of variations of equation 1.7. Standard error are clustered by politician in specifications with fixed effects: *p<0.1; **p<0.05; ***p<0.01.

1.5.2 Responsiveness to issue-specific feedback

Do politicians increase attention to gender issues after obtaining relatively more positive feedback? To address this question, we estimate the parameters of the reinforcement learning model described in Section 1.3 using two-way fixed-effect logistic panel models.

To render the data amenable to analysis using panel models, we discretize it into fixed-length time periods p. We use the calendar month as the time period since this provides a compromise between two goals: having a precise estimate of the attention given to gender issues (longer time intervals) and having more observations (shorter time intervals).

To estimate the latent issue valuations, we update valuations with every tweet m and then 'freeze' the valuations at the beginning of each period to make them conform to our panel data structure, i.e.

$V_{ip}^{k} = V_{i,m}^{k}[m = \text{first message in period}\, p]$. We take the valuation at the beginning of the month (rather than the average valuation, for example) to avoid feedback endogeneity issues.

Some politicians have breaks in their Twitter activity. Hence, assuming that feedback still affects issue attention after several months does not seem realistic. Therefore, we restrict our analysis to politician-month cells where the valuation of each issue was updated at least once during the previous month. Furthermore, we use the number of tweets published by the politician in the respective month (N_{ip}) as regression weights. Each tweet thus receives the same weight in our estimations.

In accordance with the reinforcement learning model, we estimate a logistic regression of issue attention, A_{ip}^{GI}, on the difference in valuations of gender issues and other issues, ΔV_{ip}, and politician fixed effects, π_i^G. To account for factors that affect issue attention in our empirical setting but that, for parsimony, were left out of the formal model, we augment the equation with month fixed effects, τ_p, politician fixed-effects, π_i, and time-varying control variables. Global shifts in issue attention over time are captured by the month fixed effect. For example, around March 8^{th}, the International Women's Day, politicians tweet more on gender issues. Beyond accounting for differences in baseline attention to gender issues, the politician fixed-effects, capture other time-invariant confounds such as their gender, party, region, policy focus, etc., as well as time-invariant characteristics of their followers (e.g. level of interest in gender issues).

Issue valuations are not directly observable in our data. They are latent variables constructed based on the feedback received by tweets on the issues. Therefore, the valuation updating equations have to be estimated jointly with the issue attention equation. The full model thus consists of two equations, jointly estimated as a generalized linear model using GLS:

$$\begin{cases} V_{i,m}^{k} = (1-\gamma)V_{i,m-1}^{k} + \gamma FB_{i,m-1}^{k} \\ A_{ip}^{GI} = Logit(\pi_i^{GI} + r * \Delta V_{ip} + \tau_p + \epsilon_{ip}). \end{cases} \tag{1.8}$$

Because standard software packages do not include readily available commands for the estimation of such models, we performed a grid search for the updating parameter γ. For each possible value of $\gamma \in (0,1]$ (step size = 0.01), we construct the issue valuations and the valuation difference ΔV_{it}, estimate the parameters of the responsiveness model and select the updating parameter γ with best model fit (lowest mean squared error). The exact value of γ depends on the model specification but estimates are close to 0.07 in all cases, meaning that the issue valuation is revised by approximately 7% with each tweet on the issue.

Estimation results are reported in Table 1.4. Model 1 corresponds to equation 1.8. The combination of a positive coefficient for the valuation difference ΔV_{it} and the positive valuation updating weight γ reveals that an increase in feedback to gender issue tweets is associated with an increase in attention to gender issues. A one unit increase in the difference in valuation between gender issues and other issues is associated with an average marginal increase in attention to gender issues of 7.8% (0.55 percentage points).[7] We interpret this as a substantial effect given that our fixed effect specification likely leads to conservative estimates since it focuses on within-politician, within-month variation.

In Model 2, we examine the difference in how female and male politicians learn from feedback by introducing separate valuation difference coefficients for female and male politicians. We denote by ΔV_{ip_F} the valuation difference if politician i is female and ΔV_{ip_M} if i is male. Estimates reveal that politicians of both genders are responsive to valuation differences. The weighted average marginal effect implies that an additional standard deviation in valuation difference $(+1\Delta V)$ increases female politicians' attention to gender issues by 8.5% (1.02 percentage points) whereas male politicians' issue attention increases by 6.5% (0.24 percentage points). The difference between these two estimates is not statistically significant ($p > 0.1$).

Two mechanisms could explain why the valuation difference might affect issue attention. An increase in feedback for addressing gender issues could motivate politicians to talk more about them or an increase in the feedback for addressing other issues, diminishing ΔV, could crowd out attention to gender issues. We separate these two mechanisms in Model 3. We find evidence for both mechanisms, but effect sizes differ: the positive effect size for the valuation of gender issues is larger than the negative effect size for the valuation of other issues. This suggests that crowding out is of secondary importance. Again, we do not find significant differences between female and male politicians ($p > 0.1$).

We report robustness checks in Appendix A.6. Our main results persist when controlling for politician specific trajectories in issue attention, serial correlation, or peer effects (Model 4, 5, & 6 in Table 1.4, see Appendix A.6.1). They are also robust to alternative specifications that employ different feedback measures (based on likes, A.6.2, replies, A.6.3, or a retweet-based feedback measure with no within-politician normalization, A.6.4), a different time period to compute issue attention (weeks instead of month, A.6.5), a different weighting scheme of

[7] To account for differences in the number of tweets across months, we weight for the number of tweets written in a month (N_{ip}) when calculating the average marginal effect (AME):

$$\widehat{AME} = \frac{1}{N} \sum_{i=1}^{I} \sum_{p=1}^{P} N_{ip} \big(\text{Logit}(\widehat{\pi_i^{GI}} + \hat{r} * 1 + \hat{\tau}_p) - \text{Logit}(\widehat{\pi_i^{GI}} + \hat{r} * 0 + \hat{\tau}_p)\big)$$

Table 1.4: Reinforcement learning model: female and male politicians are both responsive to feedback (eq. 1.8)

Dependent Variable:	Monthly share of tweets written on GI, $A_{ip}^{GI} = \frac{n_{ip}^{GI}}{N_{ip}}$					
Model:	(1)	(2)	(3)	(4)	(5)	(6)
ΔV	0.0936***					
	(0.0201)					
ΔV_F		0.1108***				
		(0.0287)				
ΔV_M		0.0677***				
		(0.0243)				
V_F^{GI}			0.1386***	0.1215***	0.1293***	0.1386***
			(0.0436)	(0.0376)	(0.0399)	(0.0437)
V_M^{GI}			0.1247***	0.1297***	0.1249***	0.1247***
			(0.0335)	(0.0304)	(0.0321)	(0.0334)
V_F^{other}			-0.0864***	-0.0363	-0.0698**	-0.0863***
			(0.0296)	(0.0279)	(0.0276)	(0.0296)
V_M^{other}			-0.0235	-0.0210	-0.0235	-0.0236
			(0.0337)	(0.0339)	(0.0333)	(0.0338)
Indiv. trend				5.518***		
				(0.4425)		
Lagged DV					1.011***	
					(0.1013)	
Social Influence						-0.0449
						(0.8642)
$\hat{\gamma}$ (to calc. valuation)	0.07	0.07	0.07	0.07	0.07	0.07
Fixed-effects						
Politician	Yes	Yes	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes	Yes	Yes
Fit statistics						
Squared Correlation	0.573	0.573	0.573	0.601	0.578	0.573
Observations	18,482	18,482	18,482	18,482	18,482	18,482

Note: Estimation of the model in equation 1.8. All regressions use cell-size regression weights, ie. number of tweets published by politician i in period p (N_{ip}). Standard errors are clustered by politician: *p<0.1; **p<0.05; ***p<0.01.

observations (A.6.6) or a more substantive reference category instead of 'other' (A.6.7). We do not find statistically significant differences in responsiveness between female and male politicians in any of the specification (always $p > 0.05$). Finally, we conduct a placebo test by randomly swapping politicians' issue valuations with the feedback-based valuation of another politician of the same gender (male or female), for the same issue, and in the same month (A.6.8). Using another politician's valuation leads to null results across all specifications. Thus, the robustness checks confirm that politicians are responsive to feedback.

Clearly, politicians adjust their attention to gender issues in response to the feedback they receive for their tweets on this issue.

In Section 1.3, we specified conditions for feedback to contribute to a difference in attention to gender issues between female and male politicians (eq. 1.5). The model most apt to this comparison is Model 2 in Table 1.4 because it relies on differences in issue valuations and includes separate responsiveness coefficients for female and male politicians. Combining these estimates with the estimates of the gender issue feedback advantage for female and male politicians (see Model 3 in Table 1.3), we obtain:

$$r_F \Delta\mu_F = .11 * .29 > r_M \Delta\mu_M = .07 * .16 \tag{1.9}$$

The empirical evidence thus supports the claim that the difference in the feedback female and male politicians obtain in their interactions with citizens contributes to a difference in attention to gender issues.

1.5.3 Mechanisms for the difference in gender issue feedback advantage between female and male politicians

In this section, we report empirical tests of the three potential mechanisms for the difference in gender issue feedback advantage discussed in Section 1.3.2.

Female politicians write more engagingly on gender issues

The ‘engagingness channel’ posits that female politicians write relatively more engaging tweets on gender issues compared to male politicians.

We measured how engaging is a tweet by predicting the retweet-based feedback solely on the text of the tweet. For this, we trained a BERT model to predict the feedback received by a tweet just based on its content. Importantly, the model does not take any information about the identity of the politician who published a tweet as input. As such, the *predicted feedback* is a measure of how engaging is a tweet, independent of the social category of the politician who published it (including their gender – see Appendix A.2.5).

Consider the difference in *predicted feedback* for gender issue tweets and tweets on other issues. We call it the ‘gender issue *predicted feedback* advantage’. If the ‘engagingness channel’ operates, we expect this difference to be larger for female politicians than for male politicians.

Furthermore, we expect the gap in gender issue feedback advantage between female and male politicians would disappear once we control for how engagingly tweets are written.

Figure 1.2 describes the key results based on model estimations reported in Table A.6, Models 2 and 3. Panel (a) reports the gender issue feedback advantage for female and male politicians according to the baseline model (Model 3 in Table 1.3). Panel (b) shows that whether the politician is female or male hardly affects how engaging is a tweet: the average predicted feedback is almost the same for tweets of female and male politicians. Finally, panel (c) shows that the difference in gender issue feedback advantage is almost the same when controlling for *predicted feedback* as with the baseline model. These two findings imply that *predicted feedback* does not explain the difference in gender issue advantage.

Figure 1.2: Do female politicians write more engaging tweets on gender issues?

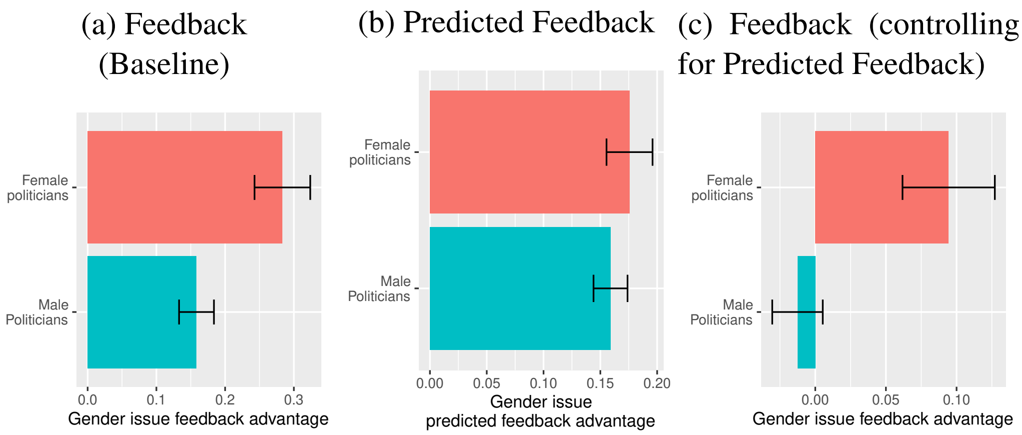

Note: Black bar represent 95% confidence interval. Differences in gender issue feedback advantage: (a) $\Delta^{F-M} = 0.125$***, (b) $\Delta^{F-M} = 0.0017$, (c) $\Delta^{F-M} = 0.107$***.

In ancillary analyses, we use stylistic features as another measure of how engaging is a tweet. We code for sentiment (from negative to positive), the number of words (tokens), hashtags, mentions, emojis, and if a tweet contains a link or a graphic element (picture or video) as alternative mediators. Model 4 in Table A.6 shows that the coefficient of the interaction term, β_{GI*M}, remains similar to that obtained in the baseline model when controlling for stylistic features. Hence, stylistic features do not contribute much to the difference in gender issue advantage (see Appendix A.7).

In conclusion, we do not find support for the 'engagingness channel'.

Self-selection of Twitter users with politicians of the same gender

The 'self-selection channel' posits that Twitter users are more likely to provide feedback to politicians of the same gender as them and that female Twitter users are more likely to provide feedback on gender issue tweets as compared to tweets on other issues.

Figure 1.3: Self-selection channel

(a) Female users react more to tweets from female politicians

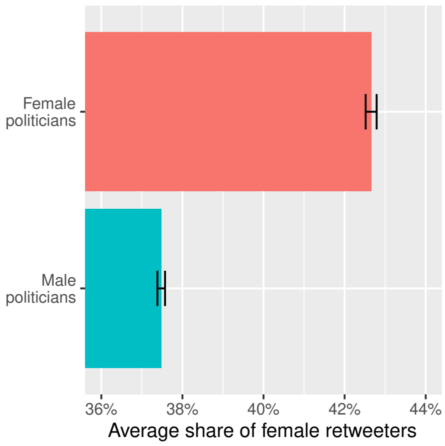

(b) Female users react more to gender issue tweets

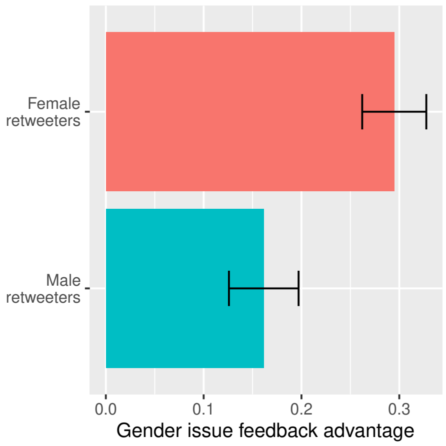

Note: Black bars represent 95% confidence intervals.

We first provide evidence for the hypothesis that Twitter users are more likely to provide feedback to politicians of the same gender. For this, we classified retweeters as female or male by applying a name recognition algorithm to their Twitter username (see Appendix A.7.2). We find differences in the gender composition of Twitter users who provide feedback to female or male politicians. The average share of female retweeters is 6 percentage points larger for female politicians (see Figure 1.3a). This difference holds for gender issue tweets (+11 percentage points) and for tweets on other issues (+5 percentage points). Hence, we find evidence that female citizens self-select into interaction more with female politicians.

To show that female Twitter users are relatively more likely to retweet tweets on gender issues, we construct the same reference-dependent standardized measure of feedback as described in Section 1.4.2, separately for female and male retweeters. This allows us to analyze how female and male retweeters react to gender issue tweets versus tweets on other issues. We find that the gender issue feedback advantage is almost twice as strong among female Twitter users compared to male Twitter users (see Figure 1.3b, and Models 6 and 7 in Table A.6).

In conclusion, we find evidence in support of the self-selection channel.

Congruity-driven feedback

The 'congruity channel' posits that, when Twitter users decide whether to retweet a tweet, their decision is affected by the interaction of the gender of the politician who published the tweet and the issue of the tweet, such that users are more likely to retweet a gender issue tweet if it was published by a female politician rather than a male politician, even after controlling for tweet content.

This mechanism differs from the 'self-selection channel' in that the latter focuses on the composition of the audience of a politician whereas the 'congruity channel' focuses on the behavior of the audience members. Accordingly, to test the hypothesis that the 'congruity channel' operates, we change the perspective from the politicians to the Twitter user as the unit of analysis. We assembled a sample of individual Twitter users and their retweeting behavior. For each user u, we take the set of tweets published by all politicians whom the user followed and test if a given user is more likely to retweet a tweet about gender if it was published by a female politician. We include a series of fixed effects to control for the general propensity of the user to retweet gender issue tweets and the user's general propensity to retweet a given politician - independently of the policy issue of the tweet. To be able to include these fixed effects, we focus on users who follow multiple politicians. For computational reasons, we focus on a sub-sample of the most active retweeters.[8] More specifically, we estimate the following logistic regression:

$$\begin{aligned} \text{retweet}_{i,u,m} = \quad & Logit(\beta_{GI*M} * GI_m * M_i \\ & + GI_m \times user_u FE + politician_i \times user_u FE + \epsilon_{i,u,m}) \end{aligned} \tag{1.10}$$

The dependent variable $\text{retweet}_{i,u,m}$ is a dummy equal to 1 if tweet message m published by politician i was retweeted by user u. The main coefficient of interest is the interaction between the politician being male and the the tweet being on gender issues, β_{GI*M}. Under the hypothesis that the congruity channel operates, we expect a negative coefficient.

We control for the average propensity of each user to retweet tweets on gender issues by including a set of user fixed effects interacted with the issue dummy, $GI_m \times user_u$, and we control for all time-invariant aspects of the politicians-user interaction (general propensity to retweet a given politician) by including a set of politician-by-user fixed, $politician_i \times user_u$.

Estimation results are reported in (Table 1.5). In models 1 and 2, the coefficient of the interaction term, β_{GI*M}, is negative and strongly significant.

[8] We selected the 1000 male and 1000 female most retweeting users, and drew a 10% random sample of the tweets of the politicians they follow. This yielded 4.4 million potential retweets.

The marginal effect implies that a given user is 10.5% less likely to retweet a tweet on gender issues if it was published by a male politician. Models 3 and 4 reveal that the effect is similar for female and male users (difference not statistically significant, $p > 0.05$).

In summary, we find clear evidence for the 'congruity channel.'

Table 1.5: Retweeting Probabilities by Gender of Politician (equation 1.10)

Dependent Variable:	Dummy = 1 if tweet is retweeted by user			
Model:	(1)	(2)	(3)	(4)
GI * Male politician	-0.1352***	-0.1396***		
	(0.0225)	(0.0224)		
Predicted feedback		1.031***		1.031***
		(0.0127)		(0.0127)
GI * Male politician * Female user			-0.1649***	-0.1752***
			(0.0322)	(0.0323)
GI * Male politician * Male user			-0.1080***	-0.1066***
			(0.0314)	(0.0312)
Fixed-effects				
Retweeter × GI	Yes	Yes	Yes	Yes
Retweeter × Politician	Yes	Yes	Yes	Yes
Fit statistics				
Squared Correlation	0.138	0.152	0.138	0.152
Observations	4,396,339	4,396,339	4,396,339	4,396,339

Note: Standard errors in parenthesis are clustered as the fixed effects. *p<0.1; **p<0.05; ***p<0.01

1.6 Discussion

In this article, we advance the understanding of issue responsiveness on social media by studying how feedback from citizens affects politicians' issue attention through the lens of a reinforcement learning model. We show that politicians respond to issue-specific feedback by adjusting issue attention. Using gender as an important case study, we demonstrate that female politicians receive systematically more positive feedback from the public when they address issues related to gender than male politicians. Our analyses suggest that this difference in feedback exists because citizens treat politicians differently depending on their gender ('self-selection' and 'congruity' channels), and not because female politicians approach the issue in a more engaging way. The difference in feedback environments to which female and male politicians are exposed leads them to focus on different issues.

Reinforcement learning allows politicians to be responsive, but only to the self-selected set of citizens who choose to interact with them. Being responsive to other entities, such as the median voter, may be more desirable from a normative perspective, but reinforcement learning is not conducive to responsiveness to such entities because politicians lack information about the preferences of citizens they do not see, and cannot perfectly adjust for biases in the feedback they receive. Of course, politicians do not learn about public opinion between elections only through interactions with the public via Twitter or in other settings. They also rely on other strategies such as opinion polls (Druckman and Jacobs, 2006). Yet, information about the average views of the public is not available continuously and for all issues, while the learning strategy we describe in this article is readily available to politicians who want to test the popularity of different issues. Uncovering when politicians rely on reinforcement learning versus public opinion polls or other tools to form perceptions of public opinion is an interesting avenue for future research.

Another relevant extension of this research would consist in applying our reinforcement learning approach to study whether the rise of Twitter and social media has increased polarization among politicians (Zhuravskaya et al., 2020). Our results imply that politicians shift attention to issues relevant to citizens they personally interact with. Hence, if politicians are frequently exposed to views from one extreme of the political spectrum on social media while seeing less moderate or opposing views, reinforcement learning could contribute to polarization of politician's discourse and behavior. Our approach could be combined with advances in text scaling methods to code the 'extremity' of tweets and study citizen-driven political polarization.

The study of politicians' behavior on Twitter is important in its own right since this behavior has real consequences (Jungherr, 2016). Still, an important next step would study the extent to which feedback on Twitter affects politicians' offline behavior. Furthermore, we suspect that the mechanism we study in this article generalizes to other settings such as other social networks, campaign meetings (applause is a clear source of feedback), or any setting in which a politicians interact with an audience, and this could be tested empirically.

Finally, more work is needed to clarify the implications of our findings for the political representation of historically under-represented groups. On the one hand, the stronger gender issue feedback advantage for female politicians strengthens the case for descriptive representation. Our findings imply that there would be less attention to gender issues if there were fewer female politicians. On the other hand, the mechanism we describe could perpetuate group-based specialization and the relegation of representatives from under-represented social categories to niche issues. Future work should aim to uncover if the differences in the feedback environments of politicians from different social categories affect their political careers.

Appendix A

A.1 Model - Proof and Additional Analyses

A.1.1 Proof of Equation 1.3

Lemma 1. *The reinforcement learning model described in Section 1.3 defines a stochastic process for* $\left(V_{i,m}^{GI}, V_{i,m}^{other}\right)_{m\geqslant 1}$ *that has a unique stationary distribution characterized by the following density:*

$$h\left(V^{GI}, V^{other}\right) = e^{-\frac{r^2\sigma^2\gamma}{2(2-\gamma)}} \frac{e^{-rV^{GI}-\pi_i^{GI}} + e^{-rV^{other}}}{e^{-r\mu_i^{GI}-\pi_i^{GI}} + e^{-r\mu_i^{other}}} g_i^{GI}(V^{GI}) g_i^{other}(V^{other}), \tag{A.1}$$

where, for $k \in \{GI, other\}$*,* $g_i^k(\cdot)$ *is a normal density with mean* μ_i^k *and variance* $\sigma^2\gamma/(2-\gamma)$*.*

Proof. This follows from Lemma 2 in Le Mens et al. (2019). □

Proof. The asymptotic probability of choosing the gender topic is obtained by integration of the choice probability (equation 1.1) with respect to the joint asymptotic density described in Lemma 1.

$$A_\infty^{GI} = \int_{V^{GI}, V^{other}} \frac{1}{1 + e^{-(\pi_i^{GI} + r(V^{GI} - V^{other}))}} dV^{GI} dV^{other} \tag{A.2}$$

$$= \frac{e^{-\frac{r^2\sigma^2\gamma}{2(2-\gamma)}}}{e^{-r\mu_i^{GI}-\pi_i^{GI}} + e^{-r\mu_i^{other}}} \int_{V^{GI}, V^{other}} e^{-rV^{other}} g_i^{GI}(V^{GI}) g_i^{other}(V^{other}) \tag{A.3}$$

$$dV^{GI} dV^{other}, \tag{A.4}$$

$$= \frac{e^{-\frac{r^2\sigma^2\gamma}{2(2-\gamma)}}}{e^{-r\mu_i^{GI}-\pi_i^{GI}} + e^{-r\mu_i^{other}}} \int_{V^{other}} e^{-rV^{other}} g_i^{other}(V^{other}) dV^{other}. \tag{A.5}$$

Noting that $\int_{V^{other}} e^{-rV^{other}} g_i^{other}(V^{other})dV^{other}$ is the moment generating function of the distribution $g_i^{other}(\cdot)$, evaluated at $-r$, we have:

$$\int_{V^{other}} e^{-rV^{other}} g_i^{other}(V^{other})dV^{other} = e^{-r\mu_i^{other} + \frac{r^2\sigma^2\gamma}{2(2-\gamma)}}. \tag{A.6}$$

We finally obtain:

$$A_\infty^{GI} = \frac{e^{-r\mu_i^{other}}}{e^{-r\mu_i^{GI} - \pi_i^{GI}} + e^{-r\mu_i^{other}}} \tag{A.7}$$

$$= \frac{1}{1 + e^{-(\pi^{GI} + r(\mu_i^{GI} - \mu_i^{other}))}}. \tag{A.8}$$

□

A.1.2 Alternative Specifications of the Valuation Updating Rule

The model analyzed in the body of the paper assumes that the feedback weight in the valuation updating equation (γ in eq. 1.2) is constant. This is as if earlier feedback instances receive a lower weight than the more recent feedback instances. Here, we relax this assumption and analyze what happens when all past feedback instances receive the same weight. In other words, we assume that the valuation of issue k is the arithmetic average of all feedback instances obtained about this issue (we assume the initial valuation to be a random draw from the feedback distribution – as such initial valuations generally differ for the two issues). In this case the law of large number implies that, in the long run, the valuation of issue k, $V_{i,m}^k$ almost surely converges to the mean of the payoff distribution μ_i^k. Therefore, issue attention almost surely converges to the same quantity as with the model analyzed in the paper (see equation 1.3).

It is also possible to analyze a 'rational' model in which the agent possesses correctly specified priors and valuations are updated according to Bayes' rule. More specifically, let f^{GI} denote the prior on the mean μ_i^{GI}. Similarly, let f^{other} denote the prior on the mean μ_i^{other}. We assume that means are realizations of the priors: $\mu_i^{GI} \sim f^{GI}$, $\mu_i^{other} \sim f^{other}$.

Let $B_{i,m}^{GI}$ denote the mean of the posterior distribution on the payoff of GI at the time of posting message m. We define the probability that agent i chooses issue k for message m in terms of $B_{i,m}^{GI}$ and $B_{i,m}^{other}$:

$$A_{i,m}^{GI} = Logit(\pi_i^{GI} + r\Delta B_{i,m}),$$

where $\Delta B_{i,m} = B_{i,m}^{GI} - B_{i,m}^{other}$.

$(B_{i,m}^{GI})_{m\geqslant 1}$ is a discrete-time martingale, and $E[B_{i,m}^{GI}] = E[\mu_i^{GI}]$ for all $m \geqslant 1$. Similarly, $E[B_{i,m}^{other}] = E[\mu_i^{other}]$ for all $m \geqslant 1$. Moreover, the optional sampling theorem implies these equalities hold asymptotically: $E[B_{i,\infty}^{GI}] = E[\mu_i^{GI}]$ and $E[B_{i,\infty}^{other}] = E[\mu_i^{other}]$.

For simplicity, let us assume that the priors on the means are normally distributed and that the payoffs distributions are normal (as in the paper). In this case, the asymptotic posterior for agent i and issue k is a single value distribution equal to μ_i^k (the variance of the posterior converges to 0). The asymptotic attention to gender issues for agent i is thus the same as with model analyzed in the body of the paper (eq. 1.3). This implies that our main results about issue attention are not contingent on assuming that agents engage in biased processing of information or have mistaken priors. They are also produced by rational information processing of possibly un-representative samples of information.

A.1.3 Simulations of Model Dynamic

To emphasize the fact that the asymmetry in feedback for messages on gender issues can be a sufficient cause for the emergence of an asymmetry in issue attention between politicians F and M, Figure A.1a presents the results of simulations obtained by assuming that the initial valuations of the two issues are the same for the two politicians (set to 0) but there is a sizeable gender issue feedback advantage for the F politician $\Delta\mu_F = 1 > \Delta\mu_M = .5$. We also assumed $r = 1$, $\gamma = .1$, $\sigma = 1$. Initially, both politicians devote the same attention to gender issues, but as the number of periods grows, an asymmetry in issue attention emerges. In Figure A.1b, we consider the case in which the initial valuations of the options correspond to the means of the feedback distributions. This amounts to assuming that agents F and M know the existence and the strength of the gender issue feedback advantage. In this case, there is a difference in issue attention between F and M from the start and it does not change (on average). Finally, in Figure A.1,c we consider the case in which the initial valuation of the gender issue by politician F is larger than the mean of the feedback distribution. In other words, F initially overestimates the gender issue feedback advantage. Accordingly the difference in attention to gender issues between F and M is initially large and it does down as F responds to feedback and adjusts her evaluation of the gender issue downward.

Figure A.1: Simulations of the dynamic of issue attention for F and M.

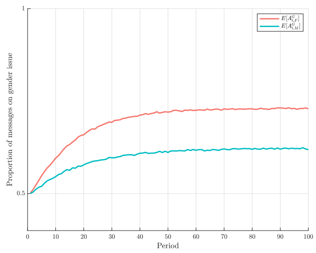

(a) Same initial valuations for politicians F and M, for tweets on both issues.
$V^{GI}_{F,1} = V^{GI}_{M,1} = V^{other}_{F,1} = V^{other}_{M,1} = 0$

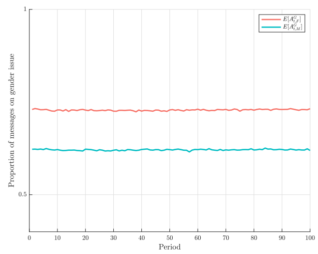

(b)
Initially correct valuations for politicians F and M (equal to the means of the feedback distributions).
$V^{GI}_{F,1} = \mu^{GI}_{F}$, $V^{GI}_{M,1} = \mu^{GI}_{M}$,
$V^{other}_{F,1} = \mu^{other}_{F}$, $V^{other}_{M,1} = \mu^{other}_{M}$

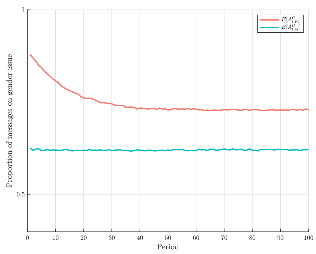

(c)
F initially overestimates the gender issue feedback advantage, M has correct initial valuations of the two issues.
$V^{GI}_{F,1} = 2\mu^{GI}_{F}$, $V^{GI}_{M,1} = \mu^{GI}_{M}$,
$V^{other}_{F,1} = \mu^{other}_{F}$,
$V^{other}_{M,1} = \mu^{other}_{M}$

Note: Figure based on 100,000 simulations with $\mu^{GI}_{F} = 1$, $\mu^{GI}_{M} = .5$, $\mu^{other}_{F} = \mu^{other}_{M} = 0$, $r_F = r_M = 1$, $\gamma = .1$ and $\sigma = 1$.

A.2 Deep Learning Tweet Issue Classifier

This section describes how we used the BERT language model to classify tweets and predict feedback.

A.2.1 Why using BERT to classify tweets?

BERT-based text classifiers offer three advantages over other machine learning classifiers. First, they perform better than 'bag-of-words' classifiers which are most often used in the social sciences (Grimmer and Stewart, 2013). By contrast to the latter, BERT is sensitive not only to word frequencies or word sequences but also to context effects. The mathematical representation of a word depends on the other words that come before and after in the text. BERT performs so well because of this sensitivity to bi-directional dependency in word meaning. Second, BERT is pre-trained on a vast amount of data (the text of all Wikipedia articles) to learn a rich language representation but can then be 'fine-tuned' for specific tasks such as classification. Most text classifiers based on machine-learning techniques are trained from scratch on a particular dataset. If the data is of limited size, performance suffers. Classifiers that are pre-trained on large amounts of text but cannot be fine-tuned are limited by the fact that word representations are not adapted to the particular task at hands (in our case, identifying tweets on gender issues). Our BERT-based model overcomes the limitations of these two earlier approaches. Third, there exists a multi-lingual version of BERT that can be used with text written in more than 100 languages. This implies that it is not necessary to translate the texts before inputting them into the model. This was vital for us, as Spanish politicians regularly tweet in Spanish (Castilian), Basque, Catalan and Galician.

A.2.2 Human Coding Stage

Supervised machine learning algorithms require a set of tweets which are correctly labeled as being on gender issues or not. We manually classified tweets as follows. First, we developed coding guidelines by creating a list of issues related to gender. Second, we selected a random sample of 19,377 tweets from that topic to be the training set and another 1975 tweets as a test set. To maximize the information contained in the training set, we over-sampled tweets on gender issues using an unsupervised topic model (LDA). We sampled tweets from one of the topics constructed by the model which contained many of words related to gender issues. The test set was sampled without over-sampling, to be representative of the whole sample. Third, we trained research assistants to code tweets independently, and resolved inter-coder disagreement or ambiguous cases by discussing with them the tweets on which such disagreement occurred. Based on a pilot study, we decided that each tweet was to be coded by two research assistants and in case of disagreement, we would search for a consensus solution. They reached an inter-coder reliability of 0.89 measured as Fleiss' Kappa which is considered a very high agreement (Landis and Koch, 1977). Disagreement occurred in only 5.2% of cases.

A.2.3 Fine-tuning BERT-based artificial neural network models

To fine-tune the algorithm we use a 10-fold cross validation (for an introduction see Hastie et al., 2009). This was implemented with Python relying on the Pytorch machine learning library by adapting publicly available code provided as part of the 'Transformers' library of language models (Wolf et al., 2019), available at https://github.com/huggingface/transformers. We created our main script by editing the provided 'run_glue.py'. We used all the default training parameters except for the following parameters which we found would lead to higher performance on the kind of data we are using: `per_gpu_train_batch_size=64, learning_rate= 2e-5, warmup_steps=.1, max_grad_norm=1.0, num_train_epochs=1.0`. The model was trained using a distributed training procedure on a GPU equipped workstation configured to perform `fp16` computations (NVidia RTX 3090).

A.2.4 Accuracy of Classification

Our model achieved an excellent classification performance. More precisely, it obtained a precision of .90 and a recall of .79. This means that 90% of tweets the model classified as being on gender issues are actually on gender issues and that 79% of gender issue tweets are classified as being on gender issues. For comparison, we also trained a model that adopts the 'bag-of-words' approach, the naïve Bayes classifier.[1] Our fine-tuned BERT classifier produces about one third of the mistakes produced by the naïve Bayes classifier (39 vs. 140). Table A.1 reports the confusion matrices for the predictions of our fine-tuned model and of the naive Bayes classifier on the test data.

To develop an intuition for the quality of the model predictions, we computed the coefficient of inter-rater reliability (Fleiss' kappa) by assuming the fine-tuned BERT model is a rater, and human categorization by the research assistants is another rater. The obtained coefficient is .83, which is a level generally considered as 'almost perfect agreement.' The same coefficient for the naïve Bayes classifier is .55, which is generally considered as 'moderate agreement.'

[1] We use the Multinomial Naive Bayes model of the scikit-learn machine learning Python package. For details see: https://scikit-learn.org/stable/modules/naive_bayes.html.

Table A.1: Confusion matrices for the BERT gender issue classifier and the naïve Bayes gender issue classifier on the validation dataset (N=1974).

(a) BERT Multilingual Cased Classifier

		Model Prediction		
		No	Yes	Total
Human Coding	No	1832	11	1843
	Yes	28	103	131
	Total	1860	114	

(b) Naïve Bayes Classifier

		Model Prediction		
		No	Yes	Total
Human Coding	No	1735	108	1843
	Yes	32	99	131
	Total	1767	207	

A.2.5 BERT-based regression model for tweet feedback prediction

To predict feedback, we relied on an artificial neural network based on BERT Multilingual-cased. Model training is performed following similar steps as in the model identifying tweets on gender issue, but the output layer in this case is not a classification layer, but a linear regression layer which takes as an input the 768 dimension vector output by BERT and outputs predicted feedback as a linear combination of the vector elements.

We split our dataset of tweets into two sets of approximately the same size, by creating a random split of politicians such that all the tweets published by a given politician would fall in one of the two sets (call them set A and set B). We adopted this politician-level split of the data to prevent the algorithm from learning about the communication style of individual politicians and the popularity of the topic of gender issues among their followers - which it could theoretically do even though no explicit pointers to politicians form part of the input data.

We constructed the measure of feedback for a given tweet by starting with the number of retweets, taking out the politician level time trend (eq. 1.6), taking out day fixed effects, and then normalizing within politicians. Unlike the measure used in the main analyses ($FB_{i,m}$), this measure includes day fixed effects. This step was not necessary when constructing the main measure because day fixed effects could be included in the regression analyses. Yet, such *post-hoc* inclusion of fixed effects is not possible in this case because we aim to use the trained model for out-of-sample predictions and thus need to remove the effects of day to day variations at the training stage.

We used all the tweets in set A to fine-tune the model and applied the resulting model to predict the success of tweets in set B. We then used the tweets in set B to fine-tune the model and applied the trained model to predict feedback for the tweets in set A. This procedure allowed us to produce out-of-sample predictions of the amount of feedback expected by a tweet, just based on its semantic content. We would like to emphasize that no features of the tweet author were included as inputs, just the text of the tweet. The correlation between out-of-sample prediction and true feedback was about .50 (.52 for the model fine-tuned on set A and .49 for the model fine-tuned on set B).

A.3 Descriptive Statistics

Table A.2: Summary Statistics: Tweets

Statistic	N	Min	Median	Mean	Max	St. Dev.
Tweet is on gender issues	1,583,917	0	0	0.06	1	0.24
Tweet is on Catalan independence	1,583,917	0	0	0.09	1	0.29
Writer is female politician	1,583,917	0	0	0.38	1	0.49
Writer is left-wing politician	1,583,917	0	1	0.56	1	0.50
Number of retweets	1,583,917	0	3	56.37	42,244	385.18
Number of likes	1,583,917	0	5	96.50	70,085	709.03
Number of replies	1,583,917	0	0	9.75	25,633	82.85
Feedback measure (based on retweets)	1,583,917	−6.81	−0.11	0.00	14.99	1.00
Share of female retweeters	472,959	0.00	0.38	0.39	1.00	0.27
Thread length	1,583,917	1	1	1.00	5	0.02
Tokens	1,583,917	3	21	22.34	95	11.84
Hashtags	1,583,917	0	0	0.45	30	0.92
Mentions	1,583,917	0	1	1.09	50	1.84
Emojis	1,583,917	0	0	0.28	140	1.07
Contains picture	1,583,917	0	0	0.31	1	0.46
Contains link	1,583,917	0	0	0.40	1	0.49
Sentiment score	1,582,931	0.0000	0.19	0.27	1.00	0.26

Note: Tokens are words or other symbols (mentions, emojis, etc.). Mentions are references to other Twitter users. The share of female retweeters is only calculated for tweets starting from 2018 with at least one identified retweeter. The sentiment score could not be computed for approximately 1000 tweets.

Table A.3: Summary Statistics: Politicians

Statistic	N	Min	Median	Mean	Max	St. Dev.
Share of tweets written on gender issues	1,265	0.00	0.04	0.07	0.60	0.08
Share of tweets written on Catalan independence	1,265	0.00	0.03	0.07	0.85	0.11
Female	1,265	0	0	0.44	1	0.50
Left-wing	1,265	0	1	0.57	1	0.50
Followers	1,223	49.00	3,127.00	23,094.08	2,390,647.00	118,166.20
Following	1,223	7.00	1,121.00	1,749.19	98,465.00	3,710.73
Tweets written since joining Twitter	1,223	97.00	7,977.00	12,875.22	134,620.00	15,036.53
Tweets written in sample period	1,265	33	652	1,252.11	29,172	2,059.92
Average number of retweets	1,265	0.13	6.87	35.54	2,651.50	136.06
Standard deviation of retweets	1,265	0.37	11.45	76.72	3,660.48	234.62
Average number of likes	1,265	0.32	9.41	62.00	5,040.03	274.30
Standard deviation of likes	1,265	0.82	15.45	131.20	6,687.89	456.71
Average number of replies	1,265	0.00	0.63	5.68	575.24	27.60
Standard deviation of replies	1,265	0.00	1.77	17.16	2,654.09	88.27
Average number of tokens	1,265	7.36	22.63	23.04	45.14	5.47
Average share of female retweeters	1,257	0.00	0.38	0.39	1.00	0.13

A.4 Evidence for Retweets as Positive Feedback

Figure A.2 plots the network of retweets between Members of Parliament of Spain's four major parties (n=527). Members of Parliaments from one of the smaller parties were excluded to facilitate visualization. Each politician represents one vertex. An edge exists if one politician retweeted another politician in our sampling period or vice versa. The figure shows that most retweets happen within parties. We interpret this as evidence that retweets are used as positive feedback.

Figure A.2: Retweeting Network between Politicians

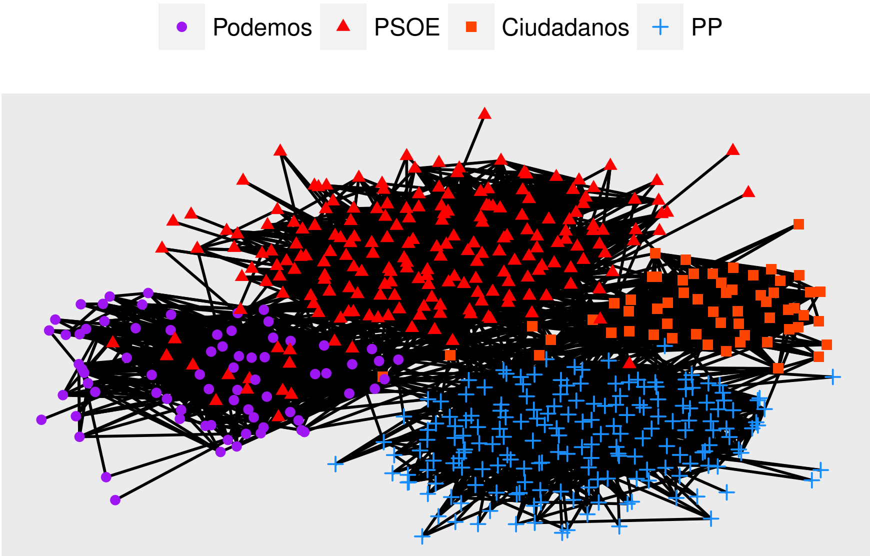

A.5 Gender Issue feedback Advantage – Robustness

A.5.1 Differences Between Left- and Right-Wing Parties

Model 4 in Table 1.3 in the main text tests if the gender issue feedback advantage is driven by one side of the political spectrum. We could conjecture that left-leaning politicians might receive more positive feedback for addressing gender issues or that a stronger feedback advantage for female politicians might be more pronounced among right-leaning politicians.[2] However, when we interact our variables of interest ($GI_{i,m}$, M_i) with a dummy equaling 1 if politician i belongs to a left-leaning party L_i, we do not find that our effects depend on the politician's ideological leaning. We coded parties as follows:

Left-leaning parties: ASG, AVANCEM, Bildu, BNG, CHA, Coalición Caballas, COMPROMIS, CpM, CUP, EM, ERC, Eusko Alkartasuna, GENTxFORMENTERA+PARTIT SOCIALISTA DE LES ILLES BALEARS, Geroa Bai, ICV, INDEPENDENT, MDyC, MÉS PER MALLORCA-PSM-ENTESA- INICIATIVAVERDS, MÉS PER MENORCA, NCa, Podemos, PRC, PSOE, UPL

Right-leaning parties: CCa-PNC, Ciudadanos, EAJ-PNV, EL PI-PROPOSTA PER LES ILLES, Foro Asturias, JxCat, PAR, PDECAT, PP, PPL,

[2]Note that in the case of Spain, a left-right classification also firmly aligns with a divide between progressive and socially conservative parties (Rama et al., 2021).

UPN, VOX

A.5.2 Feedback Measure Based on Retweets, No Within-Politician Normalization

The within-politician normalization stage of the construction of our feedback measure presumes that the psychologically relevant feedback for a politician is the amount of retweets *relative to a politician-specific baseline*. Even though work on adaptive aspirations suggests that this assumption is realistic, it is possible that politicians who tend to receive more retweets pay more attention to retweets than those who tend to receive few retweets. This possibility is assumed away by the within-politician normalization step in the construction of the feedback measure. Relatedly, within-politician normalization implies that the impact of one more retweet on feedback will differ between politicians, such that one more retweet has a larger impact on the feedback measure for politicians who generally receive few retweets. Because female politicians receive fewer retweets than male politicians, this makes the comparison between politicians of these two groups tricky, possibly leading to an inflated estimate of the difference in feedback received by female and male politicians.

To address these possibilities, we replicate our baseline specification (Model 3 in Table 1.3) removing the within-politician normalization step in the construction of the feedback measure. We find that results are similar to our main results (see Model 2 in Table A.4). The difference between female and male politicians is statistically significant ($p < 0.01$).

A.5.3 Feedback Measure Based on the Number of ‘Likes’

Model 3 in Table A.4 replicates Model 3 in Table 1.3 in the main text using likes instead of retweets to construct our feedback measure. In line with the main results, female politicians have a larger gender issue feedback advantage than male politicians. The difference is statistically significant ($p < 0.01$).

A.5.4 Feedback Measure Based on Number of ‘Replies’

Replies are a third possible source of feedback on Twitter. But by contrast to retweets, replies can be positive or negative feedback. Since they often originate from political opponents (Conover et al., 2011), it does not come as a surprise that most of them have a negative tone.

We perform three analyses to examine if replies can be used as an alternative measure of positive feedback. First, we hand-coded a sample of replies (n=100) and found that the most replies (76%) are critical. Second, we applied a sentiment analysis algorithm to all replies from 2019 (n=1,789,000). The results also suggest that replies tend to have a critical tone (average sentiment = 0.22 on a scale from 0 to 1).[3] Third, we compute our measure of feedback based on the number of replies to a tweet rather than the number of retweets and likes and examine the correlation with our main measure. We find that our reply-based feedback measure is positively correlated with the retweet-based feedback measure ($\rho = .49$), but this correlation is much smaller than the correlation between the like-based and the retweet-based feedback measures ($\rho = .85$).

Taken together, these three analyses cast doubt that tweets with a relatively large number of replies were received more positively by the public than other tweets. A large number of replies can also indicate that a tweet was received critically.

Despite this caveat, we replicated Model 3 of Table 1.3, with the reply-based measure of feedback as the dependent variable instead of the retweet-based measure. We find that tweets on gender issues receive more replies but less so for male politicians. The difference between female and male politicians is statistically significant ($p < 0.01$, see Model 4 in Table A.4). However, if we include the retweet-based feedback measure in the regression as an additional control, the coefficient for 'GI' becomes negative and the difference between female and male politicians is not significant anymore ($p > 0.1$, see Model 5 in Table A.4). This result means that controlling for the popularity of a tweet, gender issue tweets receive relatively few replies. We find a similar pattern if we use the like-based feedback measure as a control. Thus, we do not find clear evidence that tweets receive more replies because they are on gender issues.

[3] The sentiment of a reply does not always coincide with the valence of the feedback (positive or negative). A refined understanding of the context is required to judge the intention of the user who publishes a reply. Seemingly negative replies can in fact agree with the original tweet in criticizing a third part or seemingly positive tweets can in fact be meant ironic.

Table A.4: Robustness Checks: Gender issue feedback advantage (replication of Table 1.3).

Dependent Variables:	Retweets	Non-normalized Retweets	Likes	Replies	
Model:	(1)	(2)	(3)	(4)	(5)
Variables					
GI	0.2834***	0.3387***	0.2644***	0.0778***	-0.0603***
	(0.0200)	(0.0322)	(0.0195)	(0.0186)	(0.0126)
GI × Male politician	-0.1250***	-0.1385***	-0.1011***	-0.0652***	-0.0043
	(0.0232)	(0.0376)	(0.0227)	(0.0221)	(0.0157)
Part of thread	0.2920*	0.4153**	0.2535	0.2447***	0.1024
	(0.1514)	(0.1667)	(0.1620)	(0.0940)	(0.1158)
Tweets on day by politician	-0.0072***	-0.0100***	-0.0083***	-0.0025***	0.0010***
	(0.0015)	(0.0024)	(0.0019)	(0.0007)	(0.0003)
Retweets					0.4874***
					(0.0104)
Fixed-effects					
Politician	Yes	Yes	Yes	Yes	Yes
Day	Yes	Yes	Yes	Yes	Yes
Hour of day	Yes	Yes	Yes	Yes	Yes
Fit statistics					
Squared Correlation	0.01806	0.51688	0.02141	0.00961	0.24287
Observations	1,583,917	1,583,917	1,583,917	1,583,308	1,583,308

Note: Model 1 repeats our Model 3 in Table 1.3, our baseline specification. Model 2 uses non-normalized retweets. Standard error are reported in parentheses are clustered by politician *p<0.1; **p<0.05; ***p<0.01

A.6 Responsiveness to Feedback – Robustness

This section extends the discussion about the robustness checks of the responsiveness models discussed in the body of the paper. The model estimation results are reported in Table A.5.

A.6.1 Details on the Additional Specifications Reported in Table 1.4 in the body of the paper

We first account for individual trajectories in politicians' attention to gender issues over time. This is relevant since female politicians increase their attention to gender issues more than male politicians during our study period. As explained in Section 1.4.2, the feedback measure (on which issue valuations are based) already includes a politician-specific trend. This makes valuations more comparable over periods even when politicians are on different time trends. We do the same for issue attention by including a linear time trend for each politician (Table 1.4, Model 4). The coefficient for the trend is highly significant and increases the fit of the model. Yet, the estimated coefficients for the issue valuations hardly change. This is noteworthy as the trend is arguably endogenous to feedback: politicians who consistently receive more positive feedback for tweeting on gender issues will be on a steeper trend.

Next, we want to dispel concerns that serial correlation might bias our estimates. We include the lagged dependent variable (the share of tweets written on gender issues in the last month) as a control (Table 1.4, Model 5). The additional variable has a large positive coefficient and the model fit increases, but the estimated coefficients for the issue valuations remain similar to Model 3 in Table 1.4.

Finally, we address the possibility that issue attention is influenced by peer effects. Even though the month fixed effect already captures common patterns in issue attention that affect all politicians equally, it could be that politicians are more strongly affected by the behavior of politicians from the same gender. In Model 6 in Table 1.4, we include the average attention to gender issues by politicians of the same gender (male or female) as a control. The estimated coefficient is imprecisely estimated. This suggests that this sort of peer effects does not play an important role.

A.6.2 Feedback Measure Based on 'Likes'

We chose retweets over likes to construct our feedback measure because it allowed us to learn about the feedback givers' gender. Still, our theory of reinforcement learning should also apply to likes as a form of feedback. Likes have the advantage that they unambiguously stand for positive feedback. Hence, if our responsiveness results replicate using likes instead of retweets, it provides further evidence that politicians are responsive to positive feedback.

The results are reported in Table A.5, Model 1. They are similar to what we obtained with the feedback measure based on retweets. Responsiveness coefficients all have the same sign as in the main results and are significant. Point estimates are somewhat attenuated, but close. Again, the difference in responsiveness between female and male politicians is not statistically significant (p-value$>$0.1). We conclude from this analysis that politicians are also responsive to likes as a form of feedback.

A.6.3 Feedback Measure Based on 'Replies'

As discussed in Appendix A.5.4, due to the sometimes negative nature of replies and uncertain expectations regarding how politicians respond to negative feedback, we did not necessarily expect that we could replicate our results using the number of replies instead of retweets to construct the feedback measure. Nevertheless, Model 2 of Table A.5 shows that attention to gender issues correlates positively with the reply-based feedback measure. Effect sizes are somewhat smaller but comparable to those obtained in the baseline analyses (Table 1.4, Model 3).

This pattern probably results from a positive correlation between number of replies and number of retweets. Therefore, we replicate the analysis by including both the previously introduced retweet-based feedback measure and the reply-based feedback measure of feedback. As can be seen in Model 3 of Table A.5, the coefficient on the retweet-based feedback measure is almost the same as in the baseline analyses (Table 1.4, Model 3) whereas the effect of replies disappears. This reveals that the reply-based feedback does not affect issue attention in a consistent way. Note that we are *not* claiming that replies do not have an effect. Instead, it is likely that the absence of statistical effect reflects unobserved heterogeneity. For example, depending on the personality of the politician and the tone of the reply, replies might either increase or decrease issue attention. We leave further investigation of replies to future research.

A.6.4 Feedback Measure Based on Retweets, No Within-Politician Normalization

The results are reported in Table A.5, Model 4. They are similar to the baseline results reported in the body of the paper. Both female and male politicians are generally responsive to feedback. Effects are statistically significant but there is no significant difference in responsiveness between the groups in any of the models (p-value>0.1).

A.6.5 Alternative Time Period Used to Compute Issue Attention

To show that our main results do not depend on the particular choice of time period for computing issue attention (months), we replicate the specification using weeks instead of months. Model 5 in Table A.5 shows that our results hold. Again, both female and male politicians are generally responsive to feedback but no difference between the two social categories can be detected. A higher valuation of gender issues increases the attention to the issue whereas a higher valuation of other issues can lead to a crowding out.

A.6.6 Alternative Regression Weights

To avoid concerns that our main results could be driven by the specific weighting scheme we used in the model estimations, we replicate our analyses by weighting each politician month cell equally, independently of the actual number of tweets written in the politician-month cell. Our main result holds, yet, there are some differences.

For female politicians, the estimated responsiveness parameters remain stable. Model 6 in Table A.5 shows that a higher valuation of gender issues is associated with a higher share of tweets written on the issue among female politicians. For male politicians, the effects is attenuated and only barely significant ($p < 0.1$). We believe that this makes sense, considering that many male politicians write few tweets on gender issues and we need to focus on the set of politician-month cells containing more tweets to find significant effects. Giving equal weight to cells with too few underlying tweets creates too much noise. Nevertheless, the difference in effect sizes is not statistically significant in any specification.

A.6.7 Alternative Reference Category

In our analyses, we assumed that a politician who chooses the issue of a tweet selects between 'gender' and 'other' based on her valuation of the two issues. Yet, it is not psychologically realistic that politicians have a clear mental representation of the valuation the 'other' issue. To address this potential limitation, we coded tweets for another substantive category, that of 'Catalan independence' and used it as an alternative reference category.[4]

Model 7 in Table A.5 shows that using the alternative reference category does not affect the main interpretation of our results. If politicians receive more positive feedback for addressing the topic of gender issue, they tend to write more about it. On the other hand, positive feedback for tweets on Catalan independence have a slightly negative (though not significant) effect. Comparing it to the effect of the 'other' category V'_{other} reveals that the effect is relatively small. This can be explained by the fact that two specific issues compete less about attention compared to one specific issue competing with all other issues at once.

A.6.8 Placebo Test

To implement the placebo test, we randomly swap politicians' issue valuations with the feedback-based valuation of another politician of the same gender (male or female). The results are reported in Model 8 in Table A.5. The coefficients of issue valuations are close to zeros and do not reach statistical significance.

[4] We used an approach similar to that adopted to code the 'gender issue'. We hired research assistants to code about 12,000 tweets. We used about 10,000 tweets for training a BERT classifier, about 2,000 tweets for model validation and applied the model on the remaining tweets.

Table A.5: Robustness Checks Responsiveness

Model:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Variables								
V_F^{GI} retweets	0.1250***		0.1323***	0.1154***	0.1507***	0.1749***	0.1574***	-0.0057
	(0.0440)		(0.0439)	(0.0345)	(0.0329)	(0.0468)	(0.0570)	(0.0170)
V_M^{GI} retweets	0.1134***		0.1281***	0.1077***	0.1043***	0.0739*	0.1238***	-0.0056
	(0.0374)		(0.0365)	(0.0255)	(0.0317)	(0.0436)	(0.0461)	(0.0208)
V_F^{other} retweets	-0.0571*		-0.0758**	-0.0831***	-0.0979***	-0.0356		0.0431
	(0.0333)		(0.0298)	(0.0250)	(0.0252)	(0.0449)		(0.0269)
V_M^{other} retweets	-0.0292		-0.0167	-0.0214	-0.0488*	-0.0416		0.0279
	(0.0342)		(0.0357)	(0.0271)	(0.0282)	(0.0356)		(0.0249)
V_F^{GI} replies		0.1145***	0.0129					
		(0.0423)	(0.0178)					
V_M^{GI} replies		0.1045***	-0.0063					
		(0.0319)	(0.0197)					
V_F^{other} replies		-0.0498	-0.0234					
		(0.0343)	(0.0168)					
V_M^{other} replies		-0.0848**	-0.0098					
		(0.0377)	(0.0236)					
V_F^{CAT} retweets							-0.0914	
							(0.0588)	
V_M^{CAT} retweets							-0.0556	
							(0.0493)	
$\widehat{\gamma}$ retweets	0.06		0.07	0.07	0.09	0.08	0.05	0.07
$\widehat{\gamma}$ replies		0.08	0.35					
Fixed-effects								
Politician	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month	Yes	Yes	Yes	Yes		Yes	Yes	Yes
Week					Yes			
Fit statistics								
Squared Correlation	0.57291	0.57233	0.57360	0.57343	0.36229	0.59146	0.64632	0.57199
Observations	18,482	18,469	18,469	18,482	74,588	18,482	15,107	18,481

Note: Estimation of the model in equation 1.8. All regressions use cell-size regression weights, ie. number of tweets published by politician i in period p (N_{ip}). Standard errors are clustered at the level of politicians. *p<0.1; **p<0.05; ***p<0.01

A.7 Mechanisms for the difference in gender issue feedback advantage between female and male politicians – Additional results

Table A.6: Mechanisms for the difference in gender issue feedback advantage between female and male politicians - Regression results

	Dependent variable:						
	Retweets	Predicted Feedback	Retweets			Retweets from female users	Retweets from male users
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
GI	0.283***	0.176***	0.094***	0.146***	0.258***	0.295***	0.161***
	(0.020)	(0.010)	(0.016)	(0.016)	(0.019)	(0.015)	(0.017)
GI * Male Politician	−0.125***	−0.017	−0.107***	−0.104***			
	(0.023)	(0.012)	(0.019)	(0.019)			
Predicted Feedback			1.076***				
			(0.014)				
Sentiment score				−0.073***			
				(0.003)			
Token count				0.312***			
				(0.006)			
Hashtag count				0.080***			
				(0.008)			
Mention count				−0.125***			
				(0.007)			
Emoji count				−0.002			
				(0.003)			
Photo/Video included				0.255***			
				(0.009)			
Link included				0.150***			
				(0.012)			
Tweets on day by Politician	−0.007***	−0.004***	−0.003***	−0.005***	−0.013***	−0.010***	−0.011***
	(0.002)	(0.001)	(0.001)	(0.001)	(0.002)	(0.002)	(0.002)
Part of thread	0.292*	0.304***	−0.035	0.004*	0.379**	0.356***	0.290**
	(0.151)	(0.084)	(0.090)	(0.003)	(0.159)	(0.137)	(0.139)
Politician FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Hour of day FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,583,917	1,583,917	1,583,917	1,582,931	643,511	643,475	643,511

Note: Model 5-7 only consider tweets starting from 2018 with retweeter information. Standard errors are clustered by politician. *p<0.1; **p<0.05; ***p<0.01

A.7.1 Tweet Style

Figure A.3 indicates that, compared to male politicians, female politicians do not systematically use more features that attract positive feedback in their gender issue tweets. More specifically, the left panel shows the effect of different features on standardized feedback. The right panel shows the average usage of those features in gender issue tweets (compared to other tweets) separately for female and male politicians. However, female politicians do not use features that systematically attract more feedback when tweeting on gender issues, relative to male politicians.

Hence, style differences are unlikely to contribute much to the difference in gender issue feedback advantage.

Figure A.3: Effect of Feature usage on Feedback for Male and Female Politicians

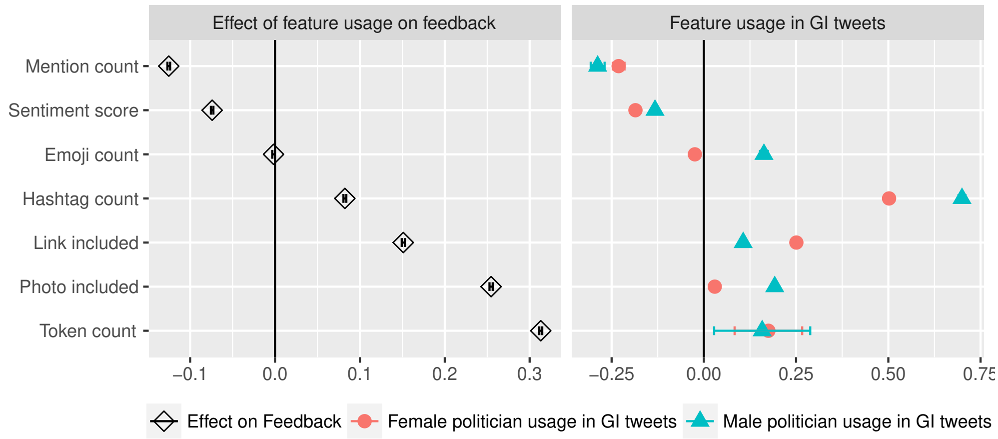

Note: Left panel plots effects of feature usage on feedback. Right panel shows feature usage in gender issue tweets relative to other tweets, separately for female and male politicians. Bars represent 95% confidence interval (confidence intervals sometimes invisible because they are close to zero).

A.7.2 Coding the Gender of Followers and Retweeters

We infer the gender of Twitter users based on the Twitter username. For this, we use Genderize.io, a commercial online service that predicts if a name is male or female.

We do this for all followers of politician and hence even if not all user names can be identified as typical female or male, the large number of followers allows us to obtain a clear picture of the share of female or male followers of each politician.

Regarding the gender of retweeters, Twitter only allows to retrospectively download information about up to 100 retweeters per tweet. Furthermore, some of their Twitter user names were not indicative if the retweeter was male or female. Still, for the average tweet in our sample, we obtained a classification for 61% of the retweeters. We estimate the absolute number of female or male retweeters by multiplying the absolute number of retweeters with the estimated share of female and male retweeters of each tweet.

Finally, we apply the same steps of feedback normalization (see Section 1.4.2) to the retweets of female and male retweets. We used this for Figure 1.3 and Models 6 and 7 in Table A.6.

2

NEITHER LEFT-BEHIND NOR SUPERSTAR: ORDINARY WINNERS OF DIGITALIZATION AT THE BALLOT BOX

Joint with Aina Gallego (Universitat de Barcelona) and Thomas Kurer (Universität Zürich)

Published as: Gallego, A., Kurer, T., and Schöll, N.B. (forthcoming). *The Journal of Politics.* DOI: https://doi.org/10.1086/714920

2.1 Introduction

The latest wave of technological change is profoundly reshaping labor markets. The spread of computers, smart software, robots and, increasingly, artificial intelligence has sparked debates about the future of work and potential repercussions in the political arena. While pessimistic voices emphasize the potential of new technologies to replace human labor and cause political upheaval, tech optimists point to a long history of misguided fears of technological unemployment.[1]

A rich literature in labor economics studies the large but unequally distributed benefits of recent technological innovation. Routine-biased technological change has mostly substituted tasks that can be accomplished by following explicit rules and thus reduces the number of routine jobs in the lower middle of the income

[1] A note on terminology: We use the term digitalization to analytically distinguish from the more generic term of technological change.

distribution. At the same time, digital technologies complement many workers concerned with more complex tasks, increase their productivity, and create high-quality jobs (Autor et al., 2003; Acemoglu and Restrepo, 2019). The resulting process of ”upskilling” in an increasingly digital world of work is a central feature of the emergence of the knowledge economy (Iversen and Soskice, 2019; Hope and Martelli, 2019; Boix, 2019).

Does this crucial economic transformation affect the political preferences of workers? Despite the evident economic benefits of digitalization, most media accounts as well as the nascent scholarly literature dealing with the political consequences of technological change have primarily been concerned with its downsides and risks, and have focused on groups left behind by this process (Frey et al., 2018; Im et al., 2019; Anelli et al., 2019; Kurer, 2020). Another highly visible group that has received considerable attention are exceptionally successful and politically influential technology entrepreneurs (Broockman et al., 2019). Even though both ”left-behinds” and ”superstars” are important constituencies, the majority of workers does not belong to either group.

In this article, we seek to provide a more encompassing understanding of the political consequences of digitalization by studying how increases in ICT capital intensity in an industry affects the political preferences of workers. Our empirical analysis uses longitudinal data from the United Kingdom that encompasses all individuals who remain active in a changing labor market. The core contribution of this paper is to document that digitalization generates a large group of ”ordinary winners”, i.e. skilled workers who have the cognitive abilities to productively use new technologies at the workplace, and to show that the political preferences of such workers who benefit economically from this development change in a stabilizing pro-system direction.

In addition, our innovative empirical approach improves on two weaknesses of existing work about the political consequences of technological change. A first concern is measurement. The aforementioned studies rely either on indirect indicators of exposure to digitalization based on the prevalence of routine tasks in an occupation or on a more direct measure of exposure to robotization. Indicators of routine task intensity (RTI) capture the task content of an occupation at a certain point in time rather than over-time variation in technology exposure. Hence, RTI has difficulty isolating a ”technology effect” from other relevant occupational characteristics. The prevalence of industrial robots, on the other hand, certainly represents a key source of pressure for particular industries, e.g. automotive production. But its consequences are of more limited relevance in the many non-manufacturing domains of the economy. We measure digitalization differently, namely as industry-specific capital stocks of information and communication technology (ICT). Importantly, ICT capital is a time-varying measure of investment in digital technology that applies to all industries. As such, it is well-suited for an analysis of the economic and political implications of digitalization among the entire labor force.

A second limitation of existing work concerns identification. Pioneering publications have relied on cross-sectional or regional data. We merge our indicator of digitalization to rich individual-level panel data from the United Kingdom and fit a series of fixed effects models to provide plausibly causal estimates of the effects of digitalization on political preferences. Panel data substantially reduce concern about omitted variables by focusing on within-individual change, which rules out that the results are driven by selection of individuals to industries or individual- and industry-level time-invariant variables. In addition, we support the causal interpretation of our findings through an instrumental variable approach and a series of robustness checks.

The empirical analysis demonstrates that a large share of the population indeed benefits economically from investment in new technology and that this economic process has political consequences. In contrast to accounts that highlight the disruptive potential of technological change among the ”left behind”, we show that exposure to digitalization increases wages for a majority of workers, a process that does not come at the cost of substantially higher unemployment. These economic benefits in turn entrench support for the political status quo: Digitalization leads to increased (a) support for the Conservative party, (b) support for the incumbent, and (c) voter turnout among ordinary winners of digitalization.

Our finding that digitalization is economically beneficial for a majority of workers and that these workers become more likely to support center-right mainstream and incumbent parties does not preclude that certain subgroups of society suffer in absolute or relative terms and might increasingly support anti-system forces. Indeed, we do find some evidence that unskilled workers, who are most susceptible to the downsides of automation, are increasingly drawn to right-wing populists when their industry digitalizes. Still, our paper shows that technological change does not only shape politics by creating a reservoir of dissatisfied losers who find the political remedies offered by populist or anti-establishment parties appealing, but it also increases support for the establishment and the democratic status quo among the large group of beneficiaries. Rather than creating dissatisfaction across the board, digitalization generates political divergence between a majority of beneficiaries and a minority of non-beneficiaries and thus presumably contributes to increasing political polarization.

To the best of our knowledge, this is the first paper that produces well identified individual-level effects of workplace digitalization on political outcomes using panel data. We contribute to the political economy literature on current political realignments and populist upheaval (Boix, 2019; Iversen and

Soskice, 2019; Rodden, 2019). These important accounts point to the ”knowledge economy” or the ”fourth industrial revolution” as the main economic force underlying changing voting patterns, party realignments, and political geography, but they do not attempt to directly measure technological change and have not examined if the introduction of digital technology modifies workers' political preferences.

We also contribute to the growing literature about how economic shocks and changes in labor market outcomes alter political preferences and vote choices (see Margalit, 2019). These studies typically focus on intense negative changes in economic standing, such as unemployment experiences or large income drops. The question of whether positive changes in the workplace situation affect political behavior has received less attention. The few well-identified studies focus mostly on large, exogenous shocks such as winning the lottery (e.g. Doherty et al., 2006). We extend this literature by focusing on a particular source of changes in the workplace, digitalization, which produces smaller but more continuous economic effects on workers' economic fortunes than shocks studied previously.

2.2 Digitalization: Economic Outcomes and Political Responses

The introduction of new technology at the workplace is a source of continuous change in workers' situation in advanced capitalist democracies. In a nutshell, our argument has three steps: Digitalization has important distributive consequences and impacts wages and unemployment risk. Therefore, digitalization affects voters' attitudes and economic preferences. This in turn links digitalization to voting conservative, voting for the incumbent, or voting for mainstream parties more generally rather than supporting populists or abstaining. Crucially, all of this is moderated by education because the more highly educated benefit more from digitalization while the less educated suffer wage reductions and face more difficult employment prospects in the digital age. Digitalization hence generates political divergence between a majority of beneficiaries and a minority of non-beneficiaries and contributes to increasing political polarization.

2.2.1 The (many) winners and (fewer) losers of digitalization

Recent theoretical work contends that the effects of technological change on wages and employment depend on the outcome of two countervailing forces (Acemoglu and Restrepo, 2019): a displacement effect as machines start to perform tasks previously done by humans and a productivity effect, as they complement workers and free up time spent on dull tasks. The net effect of these two forces on wages and employment is a priori uncertain but empirical estimates suggest that the productivity effect has dominated in past centuries (Mokyr et al., 2015). Technology, along with well-designed complementary institutions, is the most important cause of the unrivaled growth in output and living standards since the Industrial Revolution. Positive net effects also hold during the last wave of technological innovation, which is characterized by the extension of information and communication technologies (ICT). Our first expectation is that a majority of workers economically benefit from the introduction of new digital technologies.

A related, less optimistic expectation is that positive net effects go hand in hand with significant heterogeneity. While digitalization has increased the demand for highly educated workers, it has substituted for less skilled work and those in routine occupations (e.g. Autor et al., 2003; Goldin and Katz, 2009). At the aggregate level, these countervailing effects have produced a pattern of job polarization (Goos et al., 2014). How the well-documented reduction in jobs in mid-paying occupations translates into individual economic fortunes is less clear and represents one of the questions we set to explore. A decline in semi-skilled jobs does not necessarily imply that individual semi-skilled workers suffer downgrading over time. The observed aggregate reductions in mid-paying jobs might be absorbed by retirement without replacement or by exit to other, potentially higher-paying, jobs (Dauth et al., 2017; Cortes, 2016; Kurer and Gallego, 2019). In short, we expect that the introduction of new digital technologies in the workplace has positive economic consequences for a majority of workers. However, these benefits are unevenly distributed and mostly accrue to workers who possess the cognitive abilities to use new technologies productively.

2.2.2 Political implications of digitalization

To derive expectations about political ramifications, we draw on theoretical accounts that view individual's economic self-interest as an important determinant of vote choice. We consider economic channels as a key mechanism linking workplace digitalization to changing political behavior, but do not rule out the existence of non-economic psychological channels. In contrast to most existing work, we do not narrow our focus on workers left behind by technological change. Because technological change might have positive net effects, we are just as interested in the theoretical implications for ordinary winners of digitalization.

Drawing on the small existing literature on political ramifications of digitalization as well as on the broader literature on the impact of economic changes, we discuss four possible effects. The first possibility is that workers at risk of displacement due to automation demand more protection and support for redistribution (Thewissen and Rueda, 2017), which should push them to vote for parties that defend economically left-wing policies. The mechanism is consistent with standard models of voting based on preferences for economic platforms, which depict political competition as a conflict about redistributive issues, where individual material circumstances and economic risk are a main driver of policy preferences and, ultimately, party support (e.g. Iversen and Soskice, 2006; Margalit, 2013; Rehm et al., 2012). In the case of the UK, this argument implies that workers who are harmed economically by digitalization may become more supportive of the Labour Party while workers who benefit become more likely to support to the Conservative Party.[2]

A second possibility is that workers who are economically affected by digitalization respond by voting for or against the incumbent. Frey et al. (2018) find that US counties with a higher exposure to industrial robots experienced larger shifts in vote shares in favor of the Republican Party between 2012 and 2016. They interpret this finding as anti-incumbent voting, an interpretation that is congruent with research about the political consequences of other structural transformations such as off-shoring and trade with China (Margalit, 2011; Jensen et al., 2017; Autor et al., 2016). The basic mechanism in this case is economic voting: negative changes in economic prospects should generate dissatisfaction with the status quo and motivate workers to support parties in the opposition. Conversely, improvements in workers' economic situation due to digitalization should increase satisfaction and increase the likelihood of supporting the incumbent.

A third possibility, and the one that has received most attention so far, is that workers who are threatened in their jobs or lose out economically from being in digitalizing work environments become more likely to vote for anti-system, radical right-wing parties (Im et al., 2019; Kurer and Palier, 2019; Anelli et al., 2019; Kurer, 2020). The key mechanism in this case is related to changing social hierarchies and the lacking trust of the disadvantaged in the political system to improve conditions and provide the left-behind with the recognition they seek. This option might have limited applicability in contexts with majoritarian electoral systems where fringe parties are not electorally viable in many constituencies. Still, we examine this third possibility by studying if workers who lose out economically from digitalization become more likely to support the UKIP (in the years this party is included in the study), while workers who benefit

[2] Although Labour's absolute position on redistributive issues has varied over time, expert survey data on the two major parties' economic left-right position leaves no doubt about the two parties' relative position, even during the Blair era (see Figure B.8 in the SI).

economically do not.

A fourth conceivable way in which technological change affects electoral outcomes is via turnout, i.e. the possibility that digitalization affects the probability to turn out in elections. One possible channel is related to changes in the resources available to participate in politics. In particular, a drop in resources can lead to ”political withdrawal” as citizens concentrate on solving more pressing problems (Rosenstone, 1982). Alternatively, psychological changes, i.e. the realization that tasks previously performed by humans can be carried out by machines, might undermine feelings of self-efficacy and self-esteem, which are important precursors of political engagement (Marx and Nguyen, 2016). The reverse applies to winners of digitalization.

All four possibilities are reasonable ways in which digitalization can affect voting behavior. Previous research in political science about the impact of changes in workers' economic situation provides little guidance about which option is most plausible. In fact, in a recent review of the literature, Margalit (2019) compiles abundant evidence that negative economic shocks, such as becoming unemployed or experiencing income drops, can produce different political effects, including anti-incumbent voting, support for radical parties, support for the left, or a reduction in voter turnout, and concludes that “research to date offers very limited insight on the conditions that lead to one such response over another” (2019, p. 279). For this reason, we examine all possibilities in our empirical analysis and attempt to examine distinct mechanisms, including attitudes about economic issues and overall satisfaction.

Note that the four possibilities apply even in the absence of public debate about the issue of digitalization and even if workers do not consciously relate changes in their workplace due to digitalization (which may affect them economically or psychologically) to their party choice.[3] For instance, voters may just rely on loose cues about general satisfaction to evaluate the performance of the incumbent. Our theoretical expectations could vary if parties more actively politicized the issue of digitalization. However, as in other Western European democracies (König and Wenzelburger, 2018), digitalization remains a marginal issue in UK party manifestos in spite of the pressure for policy change. An analysis of the most recent manifestos shows particularly little attention to digitalization and new technology in the Labour manifesto. The Conservatives talk somewhat more about this topic and, interestingly, do so in an almost exclusively positive tone highlighting business opportunities, prosperity and security (details provided in the SI). If anything, we would hence expect that their

[3]One might reach different conclusions when studying more specific and fine-grained policy preferences instead of general preferences in favor of a center-left vs. a center-right party. For example, Barber et al. (2013) have demonstrated substantial informational barriers when voters are asked to distinguish between the redistributive and insurance elements of public policy.

way to address the issue is particularly appealing to winners of digitalization.

2.3 Data and descriptive overview

Our empirical analyses focus on the case of the UK, an established democracy at the frontier of technological innovation for which rich longitudinal micro-level data are available.

2.3.1 Industry level measure of digitalization

To measure digitalization, we follow Michaels et al. (2014), who use yearly changes in ICT capital stocks within industries (see also Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020). This is our main explanatory variable. We use the September 2017 release of the EU-KLEMS dataset (Jaeger, 2016), which contains yearly measures of output, input and productivity for 40 industries in a wide range of countries, including the UK, and covers the period 1997 to 2015. The data is compiled using information from the national statistical offices and then harmonized to ensure comparability. Most importantly for our purposes, the database provides a breakdown of capital into ICT and non-ICT assets (O'Mahony and Timmer, 2009). This allows for the creation of time-varying, industry-specific indicators of digitalization based on ICT stocks. We extend the existing time-series until 2017 on the basis of cross-classified Eurostat data on fixed assets by industry and asset (stocks), indexed by 2015 EU-KLEMS values.

Our measure of digitalization is constructed as follows:

$$\mathrm{D}_{j,t} = \frac{(\text{ICT capital stock in thousand GBP}_{j,t})}{(\text{Employees}_{j,t})}$$

Where ICT capital $\text{stock}_{j,t}$ is the sum of the fixed capital stocks in computing equipment, communications equipment, computer software and databases in industry *j* in year *t*, at constant 2010 prices, and is normalized by the number of employees in that industry.[4]

[4]Productivity-enhancing and potentially labor-replacing investments can in principle affect our measure in two ways. First, they increase the numerator (the ICT capital stock) and second, they can reduce the denominator if labor-saving technologies are implemented and reduce the number of employees in the industry. This is a manifestation of the two-fold consequences of digitalization: It can be beneficial for workers by increasing productivity or threatening if it reduces labor demand. Our measure hence captures ICT intensity relative to labor in an industry, rather than ICT intensity in an absolute sense.

Figure 2.1: Digitalization: ICT capital stock per employee, by industry

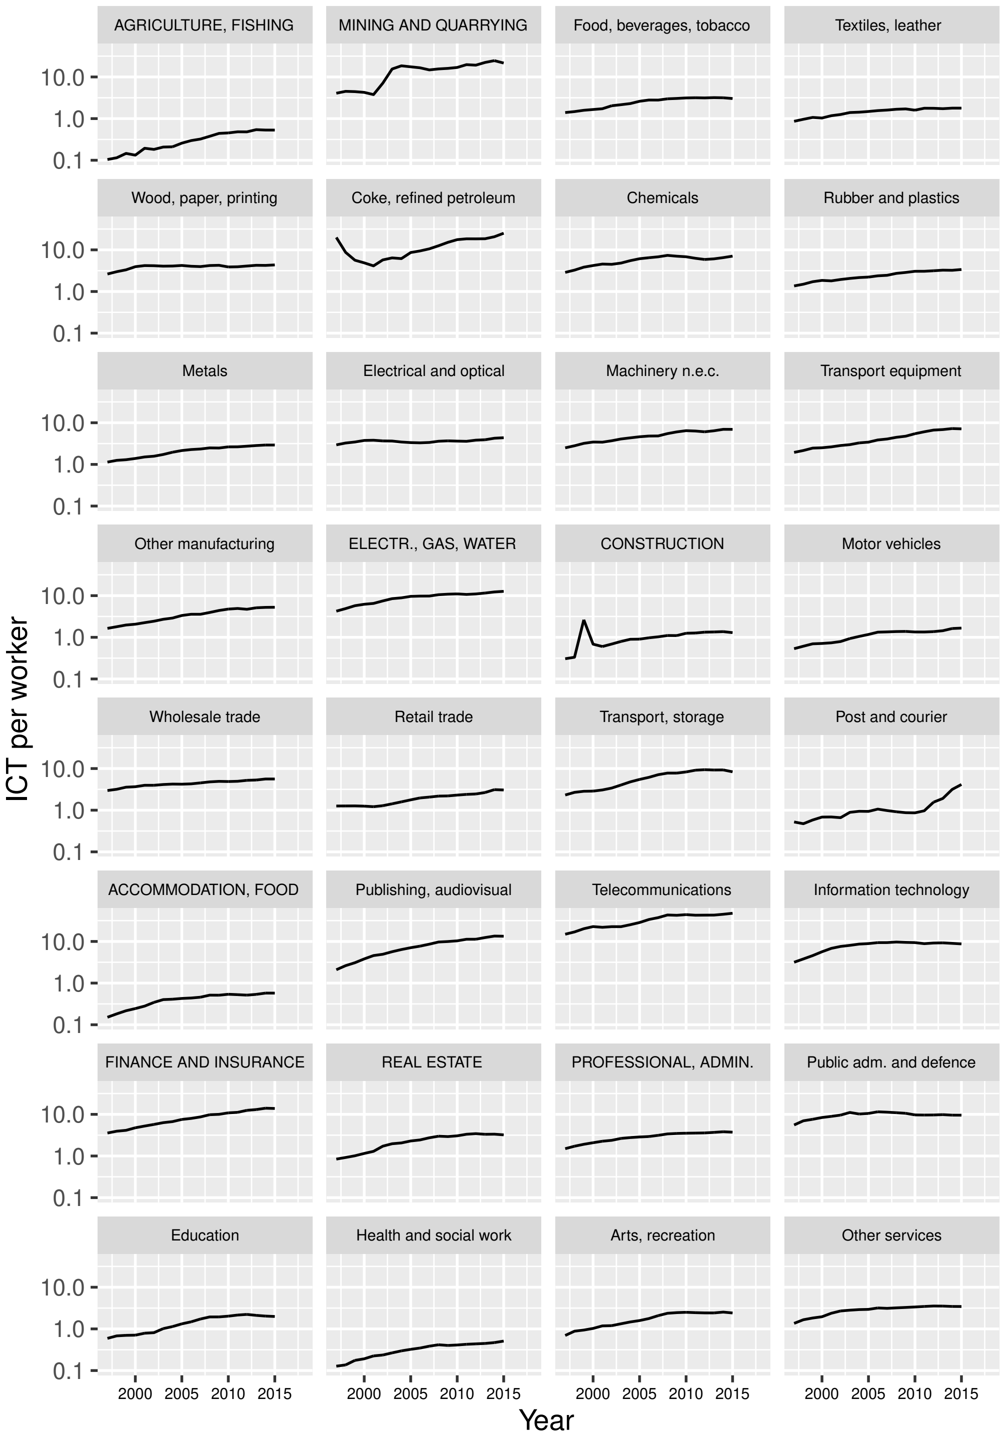

Note: Digitalization measured as yearly ICT capital stock per worker for the industries provided by EU KLEMS. Industries at the 1-digit level are written in capital letters, while industries at the 2-digit level are in lower case letters. The y-axis has a logarithmic scale to facilitate visualization.

Figure 2.1 plots the evolution of our indicator of digitalization over time for the industries provided by EU KLEMS.[5] Some industries are disaggregated only at the 1-digit level (e.g. Agriculture, forestry and fishing), while for other industries EU KLEMS also breaks down the data at the more fine-grained 2-digit level (e.g. manufacturing is disaggregated into 11 categories such as ”food products, beverages and tobacco”).

As expected, we see a general increase in the importance of digital technologies over time. The levels of ICT intensity also vary across industries in a sensible way (e.g. they are highest for telecommunications, or finance and insurance, as we would expect), adding to our confidence that the measure is valid. If anything, the trend shown understates the true degree of digitalization as ICT prices fell over time.

An important difference between our measure and the more widely used measure of robotization (Graetz and Michaels, 2018; Acemoglu and Restrepo, 2020; Anelli et al., 2019; Frey et al., 2018) is that ICT investment has affected all sectors in recent decades, allowing us to study effects of digitalization across the entire labor force. ICT capital reshapes all sectors of the economy and only 40% of total investment takes place in manufacturing industries. By contrast, deployment of robots is more concentrated: In the UK in 2017, according to the International Federation of Robotics, more than 90% of the operational robots were used in manufacturing, by far the largest chunk of it in the automotive industry. Hence, while robotization certainly represents a key source of pressure on workers in certain manufacturing industries, our time-varying measure of technological change appears well-suited to study political repercussions in the broader population. ICT capital affects the entire active labor force and thus nicely complements other studies that focus on particularly disruptive but more concentrated technological innovation in specific sectors of the economy.

2.3.2 Individual-level survey data

We combine our measure of digitalization at the industry level with longitudinal data from the British Household Panel Study (BHPS) and the Understanding Society (UKHLS) survey. The BHPS is a longitudinal study that has interviewed about 10’000 individuals nested in 5’000 households drawn from a stratified random sample of the British population yearly from 1991 to 2008. In 2009 the

[5]EU KLEMS data is disaggregated by 35 industries based on the industry standard classification system used in the European Union (NACE rev1). For 3 industries, ICT data is missing or has only zero values which reduces our sample to 32. NACE codes are consistent with UK SIC codes provided in the BHPS, which allows for a comprehensive merge of the two datasets. The scale of the y axis is logged to facilitate visualization, but the analyses use the original variable, operationalized as discussed above.

BHPS was transformed and expanded into the Understanding Society (UKHLS) survey (see Buck and McFall, 2011). Every year participants are asked detailed questions about their economic situation, current and past employment, as well as a few political questions.

For each year (date of interview), we assign every worker the value of our measure of digitalization (ICT per worker) in his or her current industry. Because the latest release of EU KLEMS only covers the period since 1997, we exclude respondents surveyed between 1991 and 1996 from our study. We also exclude respondents aged 65 and older (who should be less affected by changes in the labor market) and respondents less than 18 year old. From the remaining sample, 71.3% can be linked to one of 32 industries (NACE rev. 2). We exclude extraterritorial organizations and households as employers as there is only very sparse information on ICT capital stocks. Our total final sample contains 287’352 observations for 61’071 unique individuals. Excluded from our sample are people not assigned to an industry (including students or the currently unemployed if no industry is reported), people who never enter the labor force, and people who have exited the labor force. Table B.1.1 provides detailed summary statistics of all variables used.

The dependent variables in our analyses are a set of indicators of the personal economic situation and political attitudes asked consistently over time by BHPS/UKHLS. We compute *hourly net wages* in constant 2010 prices using the variable usual net pay per month, which is derived by BHPS/UKHLS staff using answers to detailed income questions and imputed if this information is missing. This is normalized by hours worked. We exclude observations with less than half time employment (20 hours per week) from this analysis because we found that they contain considerable measurement error.

The *employment status* refers to the week when the respondent was interviewed. Due to the lack of information about unemployment spells between surveys, we can thus only look at the moment of the interview, which most likely provides a lower bound estimate. Since we are interested in the effect of digitalization on the probability to *become* unemployed, we focus on the effect of current digitalization on a worker’s probability to being unemployed at the time of the *next* interview.

Our measure of *voter turnout* is self-reported participation in the last general election, which is asked in all waves until 2008 and then in 2010, 2015, and 2017. We construct a *party support* variable using a series of questions asked every year on whether respondents consider themselves supporters of a party or (if they are not) if they feel closer to one political party than to the others.

To measure *support for the incumbent*, we code respondents as supporters of the incumbent party if they supported the Labour Party before the government change on May 7 2010 and the Conservative Party after it changed. The Liberal

Democrats are coded as incumbents during their spell in the coalition government between May 2010 and May 2015.

Our key moderator variable, *education*, is coded in six categories: university degree (27% on average over the entire period); other higher degree (such as teaching or nursing, 12%), A-Level and other higher secondary qualifications (24%); General Certificate of Secondary Education, O-level and other lower secondary qualifications (22%); other qualifications (8%); and no formal qualifications (7%).

We concentrate on education rather than on task content, i.e. the distinction between routine vs non-routine occupations (Autor et al., 2003), for theoretical and empirical reasons. Education is a generally stable individual characteristic, as relatively few people acquire higher educational credentials after finishing schooling in young adulthood. Intra-individual stability makes education more suited for our longitudinal analysis than routine task intensity (RTI), which is measured on the level of occupations and changes as workers switch between different jobs. RTI is hence a fluid and potentially endogenous characteristic giving rise to varied trajectories. More importantly, education should be correlated with individuals' unobserved cognitive skills and ability to learn and hence with their potential to adapt to and reap the benefits of the introduction of new digital technologies in the workplace. By contrast, it is unclear if the current RTI of a worker's job is informative about his or her ability to adapt to digitalization. In our empirical setting, which interacts an industry-level measure of digitalization with an individual trait capturing the capability to deal with this development, education is more informative about the ability to learn, retrain, and ultimately benefit from digitalization than routine task content of the current job. We support this claim with empirical evidence in section B.2 where we show that education is a stronger moderator than RTI in predicting whether workers are positively or negatively affected by digitalization in their industries.

2.4 Estimation and identification

2.4.1 Fixed-effects model

We use individual industry-spell fixed-effects models to estimate the effects of digitalization in a worker's industry on labor market and political outcomes. Our modelling strategy controls for all time-invariant individual and industry-level characteristics, and only uses over time variation in the level of digitalization within industries for workers who remain in the same industry for two or more periods to identify the effect of digitalization.

To test the expectation that the effects of digitalization on labor market and political outcomes are heterogeneous depending on workers' education level, we estimate separate slopes for the effect of digitalization in a worker's industry for workers with different education levels. Our baseline specification is:

$$Y_{ijt} = \sum_{s*=1}^{6} I_{[S_{it}=s*]}\delta_{s*} + \theta_0 \times D_{jt} + \sum_{s*=1}^{6} I_{[S_{it}=s*]}\theta_{s*} \times D_{jt} + \eta_{ij} + \mu_t + \boldsymbol{\gamma}'\mathbf{C}_{it} + \epsilon_{ijt} \tag{2.1}$$

Where Y_{ijt} is the outcome of interest (economic or political) for individual i in industry j at time t. It is a function of six dummy variables $I_{[S_{it}=s*]}$, which take the value 1 if an individual has the corresponding education level and 0 otherwise. The coefficient vector $\boldsymbol{\delta}$ identifies separate intercepts for each education level.[6] We further add the time-varying measure of digitalization (ICT capital stock per worker) at the industry level D_{jt} described above and interact it with the education level dummy variables $I_{[S_{it}=s*]}$ to estimate a different slope for the effect of digitalization on economic and political outcomes for each education group. This is important as we argued that a worker's education level is a key moderator to understand the implications of being exposed to digitalization.

In our baseline specification, we include the term η_{ij}, a vector of individual by industry fixed effects (or industry-spell fixed effects) which captures all time-invariant variables that might affect labor market and political outcomes, self-selection of workers into specific workplaces, such as their gender, personality or family origin, as well as time-invariant industry-level characteristics. The industry-spell fixed effects include separate intercepts for the same individual in periods when he or she has worked in a different industry, which allows us to rule out that switchers to different industries are driving the results.[7] However, we also conduct extensive robustness checks to examine if our conclusions hold using alternative fixed effects specifications.

Furthermore, we include a year fixed effect μ_t. The fixed effect absorbs the impact of any contextual factors that are common to all individuals such as the growth of the economy or the performance of a given party. Hence, our analyses rely only on within-individual variation, controlling for circumstances that are common for all individuals. While the fixed effect capture most unobserved

[6]For most individuals, the education level is constant in all waves of the study. In our fixed effect model, the coefficient vector $\boldsymbol{\delta}$ will only be identified by the few who upgrade their education level as education is otherwise absorbed by the individual fixed effect. Therefore, we do not focus on the direct effect of education when interpreting the results.

[7]This is important because differences in digitalization across industries are much larger than differences within industries from one year to another. Any changes occurring when workers move to a different industry (which may coincide with many other relevant changes besides digitalization) would dominate the more subtle effects of digitalization at a given workplace we are interested in.

heterogeneity, we still add a vector $\mathbf{C}_{it}$ of time-varying individual-level controls. Here, we include age as a non-linear control because there is a sharp increase in the average values of most variables (such as hourly wages or voter turnout) during the 20s and 30s while their values level off later in life.
To allow for the correlation of error terms of the same individual over time and when they work in different industries, we cluster the error term ϵ_{ijt} at the individual level. We report an alternative specification with standard errors clustered at the level of the variation of the treatment, that is on the industry-year level, in the SI.

2.4.2 Threats to identification

A key concern with our empirical approach is the possible endogeneity of our measure of digitalization. In particular, ICT capital stocks per worker in the UK could be influenced by governmental policies that also affect workers' economic and political outcomes, e.g. policies adopted to shelter some industries from competition or subsidies to accelerate or slow down the adoption of digital technologies in some industries in response to their political power. In return, workers employed in that industry could have a more favorable view of the party in power.

To address this concern, we follow recent work on the Chinese import shock (Autor et al., 2013) and instrument our measure of ICT capital stocks per worker in the UK (D_{jt}) with an analogous measure from the USA (D_{jt}^{USA}):

$$D_{j,t}^{USA} = \frac{(\text{ICT capital stock in the USA in thousand USD }_{j,t})}{(\text{Employees in the UK}_{j,t})}$$

In the second stage, $\tilde{D}_{jt}^{USA}$ represents digitalization in the UK instrumented with values from the USA:

$$Y_{ijt} = \sum_{s*=1}^{6} I_{[S_{it}=s*]}\delta_{s*} + \theta_0 \times \tilde{D}_{jt}^{USA} + \sum_{s*=1}^{6} I_{[S_{it}=s*]}\theta_{s*} \times \tilde{D}_{jt}^{USA} + \gamma \mathbf{C}_{it} + \eta_{ij} + \mu_t + \epsilon_{ijt} \tag{2.2}$$

The first stage of the IV analysis is strong (all F-statistics are larger than 75). This is to be expected given that the USA is clearly at the technological frontier and competition and profit maximization motivate industries in other countries to adopt these productivity-enhancing technologies once they exist. Digital technologies adopted in an industry in the US are likely to be adopted in the UK as well, perhaps with a time lag.

The exclusion restriction of our IV strategy is that changes in ICT capital stocks in the USA do not produce changes in the economic outcomes or political views of workers from the same industry living in the UK *if ICT stocks in the UK are held constant*. Channels other than technology diffusion are likely to impact workers in the UK too indirectly and too slowly to drive the effects we capture. Furthermore, given the unequal size of the countries, politics and economics in the UK are unlikely to affect the adoption of technology in the USA.

We address further concerns including the specificity of ICT investment as opposed to general investment, within-subject switching between industries, displacement effects of technology, regional effects, the impact of trade, and panel attrition, among others, in the robustness section.

2.5 Results

This section presents the marginal effect of a one-unit increase in digitalization (a 1000 GBP increase in the ICT capital stock per worker, which is equal to 1.4 standard deviations of within industry variation in ICT), for workers of different education levels. The complete regression tables are presented in the SI.

2.5.1 Digitalization and Labor Market Outcomes

The first part of our analysis tests our expectations about the distributive consequences of digitalization and helps validate our novel longitudinal approach. Figure 2.2 presents the marginal effects of digitalization on net hourly wages and the probability of unemployment at the time of the next interview for workers with varying levels of education.

We find a strong positive effect of increases in digitalization in an industry on the hourly net wages of workers with higher education levels, especially university degrees. At the same time, individuals with low levels of education or no qualifications experience a reduction in their hourly wages in periods when their industry digitalizes rapidly.[8] The coefficients can be interpreted as follows: a one unit increase in digitalization (1000 GBP ICT capital stock per worker) increases the average hourly net wage of a university graduate by 0.4 GBP which is equivalent to a yearly net wage increase of 768 GBP. By contrast, a one unit increase in digitalization decreases the average hourly wage of workers with no qualifications by 0.16 GBP or 312 GBP per year.

Second, we study the effect of digitalization on employment status. In this case, we use lead models because we are interested in the probability of becoming unemployed in the future. We find some evidence that digitalization

[8]We tested if the differences in the effect of digitalization across education groups are statistically significant. All of them are, except for the difference between no qualification and other qualification.

Figure 2.2: Effect of ICT capital stock increases on labor market outcomes

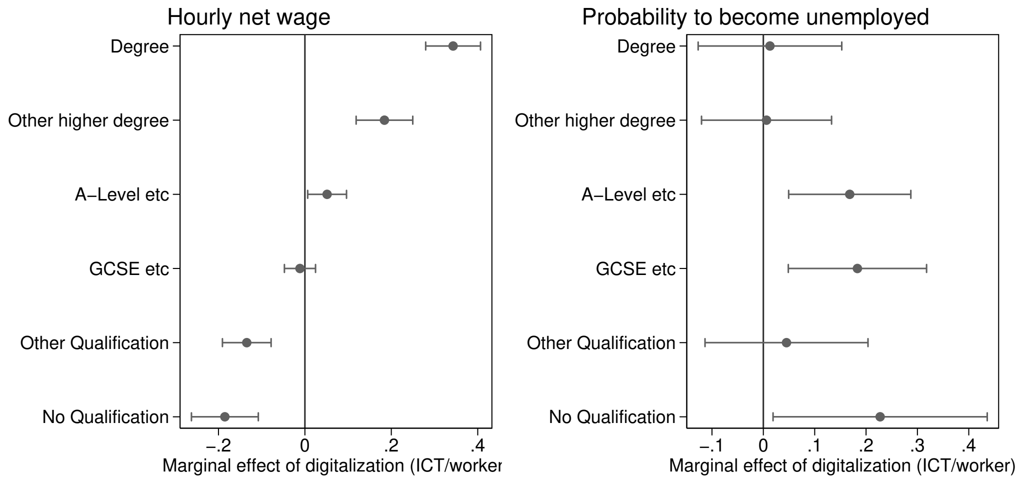

Note: Results show the marginal effect of one unit increase in digitalization (1000 GBP in ICT capital/worker) on hourly net wages (left) and the probability to become unemployed in percentage points (right).

increases the likelihood that less educated workers report being unemployed when they are reinterviewed after digitalization occurred. This finding is in line with the task-based literature emphasizing that primarily routine jobs in the middle and low end of the wage and education distribution are susceptible to automation (Autor et al., 2003). However, the effects are substantively small. For example, a one-unit increase in our measure of digitalization, i.e. a 1000 GBP increase in the ICT capital stock per worker (0.4 std), is associated with an increase in the probability to report being unemployed at the next interview of 0.24 percentage points for the no qualification group. This constitutes a 7% increase in the odds to become unemployed from 1:30 to 1:28.5. As noted above, a caveat is that we do not observe unemployment spells between interviews. The reported increase thus likely represents a lower bound estimate.

Our findings are in line with previous studies and suggests that our novel empirical approach is valid. For example, Autor et al. (2015) conclude that digitalization has rather limited net employment effects despite its profound impact on the overall employment structure. For the UK, Kurer and Gallego (2019) show that most routine workers stay in their jobs and the decline in the share of routine jobs happens through retirement and lower entry rates rather than layoffs.

So far, the analysis yields two important take-away points. The impact of faster than average digitalization on hourly wages is positive for a majority of workers, but digitalization has unequal effects on highly and less educated workers. Those with a higher degree represent 39% of our sample in 2015 and are unambiguous economic winners, as digitalization increases their wages without any adverse employment effects. Adding workers holding A-Level certificates (upper secondary education), whose wage gains come at the cost of slightly increased unemployment risk, this share increases to 61% of the population. Workers with secondary education (GCSE and similar) make for about a fifth of the population and experience neither positive nor negative income effects from digitalization. Unambiguous economic losers of digitalization are concentrated in groups with low formal educational credentials, which account for about 20% of the population. In sum, digitalization first and foremost benefits those who have the skills to thrive in a rapidly world of work and reinforces patterns of wage polarization.

2.5.2 Digitalization and Political Outcomes

Our primary interest is in whether and how these distributive effects lead to changes in individual political behavior. Figure 2.3 presents the main results regarding voter turnout, support for the Conservative Party, for the Labour Party, and for the incumbent.

We find evidence of increasingly unequal political participation due to technological change. Highly educated workers in industries digitalizing more quickly become more likely to vote. A one unit increase in digitalization raises turnout among voters with university degrees by 0.64 percentage points. On the other hand, we find no effects or negative effects among less educated workers. Recent work has shown that the gaps in the turnout rates of citizens with high and low socio-economic status has increased over time in the UK (Heath, 2018). Our results suggest that digitalization contributes to increasing inequalities in voter turnout by (weakly) augmenting existing gaps.

Next, we examine the relationship between digitalization and support for parties. The results provide clear evidence for increased support for the Conservatives among winners of technological change. For example, a 1000 GBP increase in the capital stock per worker is associated with an increase in support for the Conservatives of approximately 0.6 percentage points among the highly educated. For less educated workers, digitalization is associated with a reduction in support for the Conservatives.[9]

The results are consistent with our expectation that workers who benefit from

[9]The differences in the effects of digitalization for workers with university degrees and workers of the three lower education groups are statistically significant at conventional levels. The same is true for the difference between the top three education groups and the no qualification group.

Figure 2.3: Effect of digitalization on political outcomes, industry-spells fixed effect specification

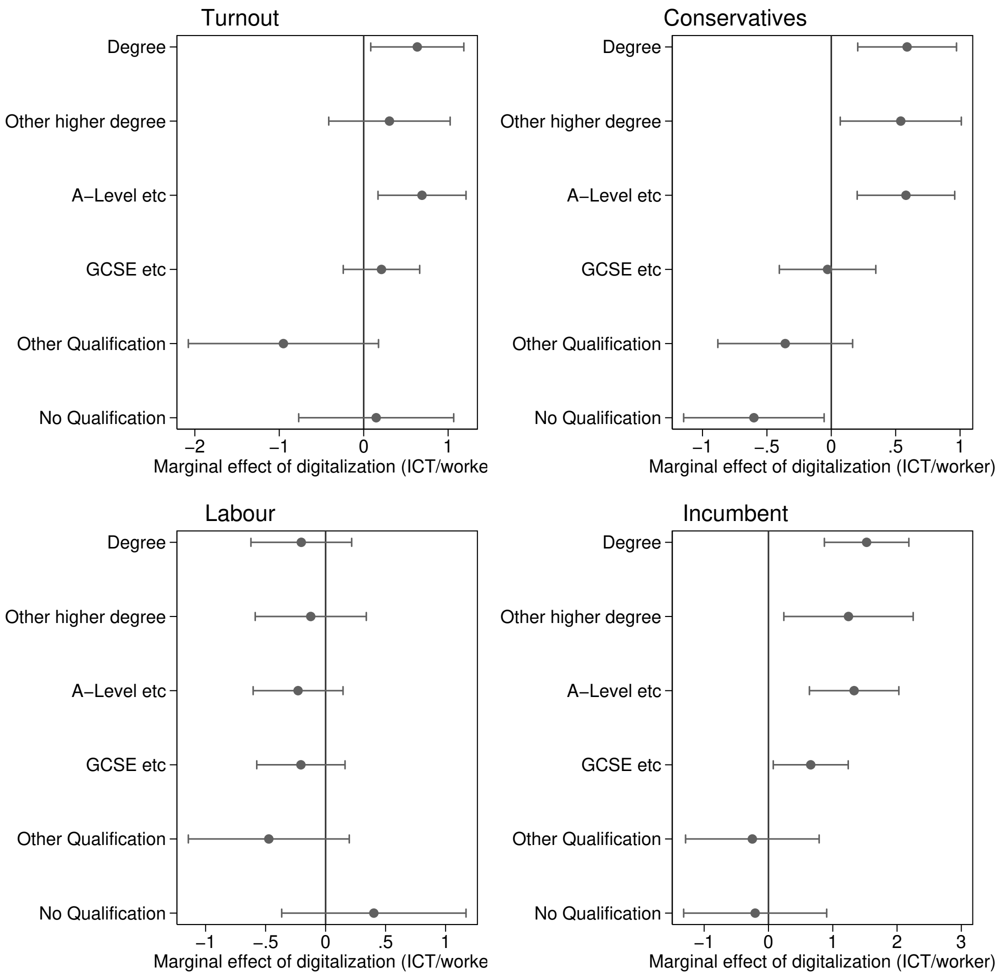

Note: Results show marginal effect of one unit increase in digitalization (1000GBP in ICT capital/worker) on the probability to report having voted or supporting a given political party. All results are in percentage points.

digitalization become more likely to support an economically right-wing party which could be due to changes in economic preferences about redistribution. In line with other studies on economic shocks and voting behavior (see Margalit, 2019), the effect is limited in magnitude. Still, the reported effects are short-term and can accumulate over time, leading to more significant shifts in party support. Moreover, even modest changes in political behavior can be politically consequential as elections are often won by small margins.

With respect to support for the Labour Party, we do not find clear results. While the pattern is to some extent a weak mirror image of support for the Conservative party, the effects are small and imprecisely estimated. This is true even among less qualified workers, which contrasts with previous research suggesting that losers of digitalization ask for more redistribution (Thewissen and Rueda, 2017). However, it should be noted that our industry-spell fixed-effect approach may underestimate the effects on the behavior of losers of digitalization since our analyses only capture political reactions of workers who remain in the labor market (see section B.4.2 for an approach that includes displaced workers).

Finally, we also theorized effects on support for the incumbent that are analytically distinct from voting decisions based on support or opposition to redistribution. The main hypothesis in this case is that through a simple reward-punishment mechanism, winners of digitalization become more likely to support the incumbent while losers withdraw support. The lower right panel of Figure 2.3 reports marginal effects of digitalization on support for the incumbent party. The results provide clear-cut evidence in line with the egotropic economic voting hypothesis: Being in a digitalizing environment increases the likelihood to support the incumbent, but only for highly educated workers (who benefit more from digitalization).

2.5.3 Incumbency effect: Analysis by period

So far, our analysis finds that digitalization increases support for the Conservative party and for the incumbent among highly educated workers. In an attempt to distinguish between these two possibilities, we re-ran our analysis separately before and after the government change in May 2010.[10]

Table 2.1 shows that our results about political effects are mainly driven by the years after 2010. Column 1 shows that digitalization did not result in significantly increased support for the Labour party during their period in government (until 2010). Columns 6 and 7, on the other hand, speak in favor of an incumbency effect because the coefficients for incumbent voting are twice as large than for vote for Conservatives. Also, the Conservative Party did not benefit from digitalization when they were in opposition (pre-2010, column 4).

The findings are consistent with the interpretation that digitalization affects support for parties through two distinct mechanisms (spatial voting and economic voting), which can cancel each other out or reinforce each other depending on

[10]Note that results are not driven by differential economic effects of digitalization before and after the Great Recession. Additional analyses presented in section B.3 in the Supplementary Information (SI) show that the estimates of the effects of digitalization on hourly wages and unemployment are comparable across periods.

Table 2.1: Sub-period Analysis: Until May 2010 and after May 2010

	Vote for Labour			Vote for Conservatives			Incumenbent
	(1) Pre May 2010	(2) Post May 2010	(3) Overall	(4) Pre May 2010	(5) Post May 2010	(6) Overall	(7) Overall
Degree × ICT	0.432	-0.694	-0.203	0.172	0.598	0.589**	1.527***
	(0.245)	(0.370)	(0.214)	(0.197)	(0.400)	(0.196)	(0.336)
Other higher degree × ICT	0.146	-0.313	-0.124	0.289	0.757	0.540*	1.245*
	(0.327)	(0.448)	(0.237)	(0.318)	(0.447)	(0.240)	(0.514)
A-Level etc × ICT	0.0302	-0.441	-0.229	0.425	0.717	0.580**	1.333***
	(0.233)	(0.386)	(0.191)	(0.222)	(0.377)	(0.193)	(0.355)
GCSE etc × ICT	0.0392	-0.406	-0.206	-0.181	0.563	-0.0288	0.657*
	(0.246)	(0.413)	(0.188)	(0.258)	(0.413)	(0.191)	(0.298)
Other Qualification × ICT	-0.240	-1.308*	-0.473	-0.402	0.650	-0.358	-0.251
	(0.443)	(0.645)	(0.345)	(0.331)	(0.609)	(0.268)	(0.534)
No Qualification × ICT	0.275	-0.528	0.402	-0.467	-0.601	-0.601*	-0.207
	(0.434)	(0.861)	(0.391)	(0.305)	(0.743)	(0.278)	(0.567)
Age	-0.393	0.143	0.128	0.0995	0.881	0.383	-0.730
	(0.339)	(0.521)	(0.268)	(0.275)	(0.462)	(0.226)	(0.409)
Age × Age	0.00420	-0.00959**	-0.00453*	-0.00198	-0.00561	-0.00330*	-0.000287
	(0.00270)	(0.00340)	(0.00182)	(0.00235)	(0.00300)	(0.00163)	(0.00317)
Constant	61.88***	64.76**	59.78***	13.02	0.508	11.99	81.14***
	(11.77)	(19.86)	(9.050)	(9.410)	(16.78)	(7.639)	(13.30)
Individual*Industry FE	X	X	X	X	X	X	X
Education Group FE	X	X	X	X	X	X	X
Year FE	X	X	X	X	X	X	X
Region FE	X	X	X	X	X	X	X
Observations	106387	114663	221050	106387	114663	221050	221050

Note: All results are in percentage points. Standard errors in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. Liberal Democrats are coded as incumbent party during the 2010-2015 coalition government. We present, for each education group, the marginal effect of digitalization (direct effect + interaction effect). This allows readers to immediately infer what is the effect of digitalization among workers with a given education level: e.g. if a university degree holder working in a digitalizing industry starts earning X more than if this industry were not digitalizing. The standard approach proposed by Brambor et al. (2006) involves including the main effect and interaction effects separately, which yields identical results. However, the coefficients would then be relative to the base category, i.e. we would compare affected workers with different education levels. Marginal effects on the other hand compare affected and non-affected workers with the same education level, and are better suited in a longitudinal framework because they emphasize within-person changes.

the ideological profile of the party in power. Although both parties' relative position on the economic left-right axis has varied over time, the Tories have had a clearly more pronounced pro-market stance during the entire time span of our analysis (see Figure B.8 in the Supporting Information). Accordingly, when the Tories are in power, both mechanisms push in the same direction for winners of digitalization, resulting in more clear-cut effects. In contrast, when the Labour party is in power, winners of digitalization face a trade-off: on one hand, the improvements in their economic situation push them to vote for the incumbent.

On the other side, this incumbent has policy positions on the economic left-right dimension that are not in line with their economic interest. Such tension may be smaller when Labour governments are in favor of promoting the advanced sectors of economy than under a more sharply left-wing party.

2.5.4 Do the left-behind turn to the populist right?

An important question attached to our primary focus on winners of digitalization is if the minority of workers who lose out in the same process politically respond by increasing support for populist or anti-system parties. Admittedly, our case and data is not ideal to fully examine this question: In a majoritarian electoral system, protest and populist parties are rarely electorally viable, making their political presence marginal. In the case of the UK, the UK Independence Party (UKIP) has been a fringe party over most of the period studied and support for UKIP has only been coded since 2013 in BHPS/UKHLS. Hence, the data available to examine this question is limited to the latest period.

Nevertheless, our results, which should be interpreted with caution, support the possibility that the ”left-behind” might turn to the populist right when their workplace digitalizes. Figure 2.4 shows marginal effects of digitalization on UKIP support. We find increased support among the small group of unambiguous losers of digitalization (the ”no qualification” group is about 4% of our sample since 2013). This is consistent with previous findings that digitalization makes losers more likely to support anti-establishment parties (Im et al., 2019; Anelli et al., 2019; Kurer, 2020). The magnitude of the effect is impressive but it is very imprecisely estimated.[11] While the negative effect of digitalization on low-skilled workers' wages might rather suggest support for a pro-welfare party than for the populist right, the below section on attitudinal mechanisms offers some evidence that welfare chauvinism and competition for social expenditure might be part of the explanation.

2.6 Instrumental variables analysis

Since one might worry about endogeneity of our measure of digitalization, e.g. due to governmental policy support for specific sectors, we instrument ICT capital stocks in the UK with analogous data from the United States. Tables 2.2 and 2.3 present the results of the instrumental variables analysis next to the baseline results.

[11] A possible concern is that a large share of low-skilled workers has migration background, which in turn mutes right-wing populist support but Table B.11 in the SI shows that the results are substantively unchanged when excluding people born outside of the UK.

Figure 2.4: Effect of digitalization on UKIP support, industry-spells fixed effect specification

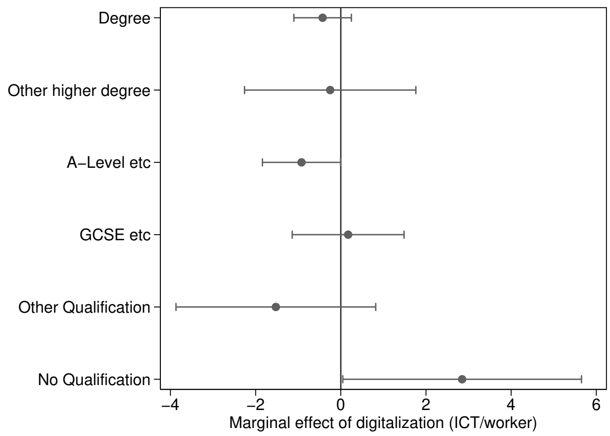

Note: Results show marginal effect of one unit increase in digitalization (1000GBP in ICT capital/worker) on probability to report to have voted or support UKIP. All results are in percentage points.

All economic and political results remain qualitatively unchanged, although the instrumental variable approach tends to produce larger point estimates. Obtaining larger IV estimates is not unusual and could be due to different reasons. A small part of the difference between our main specification and the IV can be attributed to differences in the sample used. EUKLEMS does not provide data for two industries in the USA (telecommunications and wholesale and repair of motor vehicles) resulting in a slightly smaller and more homogeneous sample. When we rerun the main analyses excluding these industries, the coefficients become somewhat closer to the IV results. Measurement error may also contribute to explain the larger IV coefficients if ICT capital stocks are better measured in a larger economy like the USA.

More substantively, the difference between the coefficients suggests that our measure of digitalization in the UK is indeed partly endogenous. One possible reason is that policy in the UK may work to limit the polarizing effects of digitalization on economic and political outcomes. Another reason could be that industrial policy in the UK might lead to an inefficient allocation of ICT investment across industries. Yet another explanation is related to trade unions pressure on firms to mitigate the strongest symptoms of digitalization on workers' material and psychological well-being. All three processes would result in attenuation bias in our main specification.

Table 2.2: Instrumental Variable Results: Economic Outcomes

	Hourly net wage		Probability to become unemployed	
	(1) Main specification	(2) Instrumental variable	(3) Main specification	(4) Instrumental variable
Degree × ICT	0.343***	0.435***	0.0129	0.241
	(0.0324)	(0.0808)	(0.0713)	(0.197)
Other higher degree × ICT	0.184***	0.301***	0.00620	0.354
	(0.0336)	(0.0745)	(0.0644)	(0.211)
A-Level etc × ICT	0.0514*	0.104	0.168**	0.421*
	(0.0229)	(0.0860)	(0.0608)	(0.203)
GCSE etc × ICT	-0.0114	-0.0477	0.183**	0.631
	(0.0185)	(0.0598)	(0.0686)	(0.413)
Other Qualification × ICT	-0.135***	-0.228**	0.0451	0.572*
	(0.0288)	(0.0876)	(0.0807)	(0.274)
No Qualification × ICT	-0.185***	-0.305***	0.227*	0.620
	(0.0398)	(0.0894)	(0.106)	(0.446)
Degree	-1.995***	-2.513***	0.883	1.496
	(0.209)	(0.308)	(0.793)	(1.257)
Other higher degree	-2.028***	-2.622***	1.446	1.549
	(0.218)	(0.294)	(0.778)	(1.174)
A-Level etc	-1.628***	-1.970***	0.607	1.169
	(0.156)	(0.250)	(0.691)	(1.094)
GCSE etc	-1.141***	-1.254***	0.773	0.741
	(0.147)	(0.218)	(0.657)	(1.183)
Other Qualification	-0.441**	-0.420	1.124	0.900
	(0.137)	(0.222)	(0.652)	(0.964)
Age	0.345***	0.346***	-0.435***	-0.442***
	(0.0271)	(0.0277)	(0.0994)	(0.101)
Age × Age	-0.00312***	-0.00311***	0.00158**	0.00166**
	(0.000212)	(0.000220)	(0.000604)	(0.000624)
Constant	-2.821***		13.76***	
	(0.797)		(3.681)	
Individual*Industry FE	X	X	X	X
Year FE	X	X	X	X
Region FE	X	X	X	X
Observations	179477	151642	216130	187153
First stage F-stat		104.6		90.11

Note: Probability to become unemployed is the probability of being unemployed at the time of the next interview (reported in percentage points). Standard errors in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 2.3: Instrumental Variable Results: Political Outcomes

	Turnout		Conservatives		Labour		Incumbent	
	(1) Main	(2) IV	(3) Main	(4) IV	(5) Main	(6) IV	(7) Main	(8) IV
Degree × ICT	0.635*	1.396*	0.589**	2.198**	-0.203	0.324	1.527***	2.877*
	(0.282)	(0.622)	(0.196)	(0.672)	(0.214)	(0.529)	(0.336)	(1.444)
Other higher degree × ICT	0.305	2.299*	0.540*	1.759*	-0.124	0.272	1.245*	2.365*
	(0.366)	(1.051)	(0.240)	(0.696)	(0.237)	(0.666)	(0.514)	(1.182)
A-Level etc × ICT	0.691**	1.998*	0.580**	1.513*	-0.229	-0.550	1.333***	2.683**
	(0.264)	(0.992)	(0.193)	(0.592)	(0.191)	(0.532)	(0.355)	(0.943)
GCSE etc × ICT	0.211	1.186	-0.0288	0.917	-0.206	0.464	0.657*	2.034*
	(0.231)	(0.983)	(0.191)	(0.657)	(0.188)	(0.598)	(0.298)	(0.952)
Other Qualification × ICT	-0.951	1.863	-0.358	1.468	-0.473	0.451	-0.251	2.645
	(0.575)	(1.860)	(0.268)	(0.996)	(0.345)	(0.975)	(0.534)	(1.776)
No Qualification × ICT	0.148	2.235	-0.601*	0.443	0.402	0.216	-0.207	0.556
	(0.470)	(3.140)	(0.278)	(1.073)	(0.391)	(1.761)	(0.567)	(2.138)
Degree	-0.617	2.391	-7.420***	-8.232**	2.319	0.601	-12.11***	-12.67*
	(3.336)	(6.396)	(1.937)	(3.101)	(2.371)	(4.350)	(3.591)	(6.092)
Other higher degree	-2.424	-2.807	-5.326**	-5.324	0.522	-0.803	-9.677*	-10.15
	(4.038)	(6.884)	(2.053)	(3.238)	(2.439)	(4.405)	(3.982)	(5.948)
A-Level etc	-5.519	-3.938	-6.227***	-5.190	0.879	1.462	-9.460**	-10.63*
	(2.846)	(5.948)	(1.786)	(2.762)	(2.164)	(4.043)	(3.136)	(5.152)
GCSE etc	-4.484	-2.404	-3.577*	-3.018	1.581	-0.428	-9.527**	-10.57*
	(2.881)	(5.919)	(1.744)	(2.822)	(2.028)	(3.919)	(3.147)	(5.088)
Other Qualification	0.548	-1.107	-0.00495	-1.629	-0.495	-3.125	-1.458	-6.587
	(2.274)	(6.176)	(1.703)	(2.942)	(1.824)	(3.600)	(2.602)	(4.909)
Age	-1.143**	-1.112**	0.383	0.354	0.128	0.189	-0.730	-0.739
	(0.390)	(0.404)	(0.226)	(0.232)	(0.268)	(0.274)	(0.409)	(0.417)
Age × Age	-0.00913***	-0.00951**	-0.00330*	-0.00276	-0.00453*	-0.00531**	-0.000287	0.000314
	(0.00264)	(0.00290)	(0.00163)	(0.00170)	(0.00182)	(0.00191)	(0.00317)	(0.00325)
Constant	133.1***		11.99		59.78***		81.14***	
	(12.47)		(7.639)		(9.050)		(13.30)	
Individual*Industry FE	X	X	X	X	X	X	X	X
Year FE	X	X	X	X	X	X	X	X
Region FE	X	X	X	X	X	X	X	X
Observations	103739	81054	221050	187899	221050	187899	221050	187899
First stage F-stat		109.9		86.76		86.76		86.76

Note: All outcomes are in percentage points. Standard error in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

2.7 Robustness Checks

We run a series of robustness checks in order to rule out alternative interpretations and further concerns about endogeneity. Perhaps the most important concern with respect to the main findings relates to the possibility that an increase in ICT capital investment simply reflects the fact that an industry is doing well and thus able to offer higher wages and better working conditions.

This could invalidate the interpretation of our results since they would not capture the specific consequences of digitalization but rather the effect of working in a thriving industry. To assess this possibility, we conduct an additional analysis using *non-ICT investments* as the main explanatory variable. Non-ICT investments are simply the sum of all assets minus our three ICT categories (capital stocks in computing equipment, communications equipment, and computer software and databases) divided by employees. (We discuss different disaggregations of the residual asset categories in the Supplementary Information.) Changes in an industry's non-ICT capital stock per worker do not predict any of the outcomes we are interested in, suggesting that our results specifically capture the consequences of digitalization rather than a thriving industry.

Further analysis deal with potential outliers (e.g. rapidly digitalizing industries or regions); additional controls for trade exposure to isolate the impact of technology; different fixed-effects structures and clustering at the industry instead of the individual level. In addition, we replicated all analyses using lead models to better capture negative effects on workers who lose their job and hence drop out of the labor force. Finally, we have a closer look at attrition. Overall, the result of the robustness checks are reassuring. We can recover our substantive results in all of these additional models. We present a more detailed description of both the empirical concerns and our proposed remedy including full regression tables in the Supplementary Information (section B.4).

2.8 Mechanisms

The causal chain underlying our argument assumes three steps, namely that (i) digitization creates winners and losers through its differential impact on wages and employment along an education gradient. These distributive consequences (ii) affect individuals political preferences and attitudes, which leads beneficiaries of digitalization to (iii) voting for conservative parties, voting for the incumbent, and higher turnout rates. We have provided robust evidence for (i) and (iii) in the above analysis.

As a final step, we assess some attitudinal mechanisms possibly linking digitalization's implications to electoral behavior. To be clear, our panel data is not ideally suited to trace attitudinal mechanisms. The number of questions on preferences and subjective perceptions of respondents is small and they are infrequently included, as most attitudes are only asked in a few waves. The few questions asked repeatedly are imperfect indicators of the theoretical concepts of interest, introducing measurement error, which attenuates results and is particularly relevant in a longitudinal analysis. This final auxiliary analysis helps us assess the plausibility of attitudinal channels, but it is not powerful enough to clearly refute any of them.

We argued that workplace digitalization can increase support for right-wing parties through a change in preferences for economic policies if winners of digitalization become less likely to support a redistributive welfare state. Additionally, we argued that digitalization can increase support for the incumbent party if winners become more satisfied in general and more supportive of whoever is in government. For both processes, we anticipate the opposite reaction for losers. We operationalize preferences about economic policies through a battery about preferences for state intervention which asks about governments' capacity to solve economic problems and their obligation to provide jobs, and satisfaction with a question asking respondents about general life satisfaction. The exact wording and results figures are provided in the Supplementary Information.

Digitalization is associated with at best small changes in life satisfaction, but we do observe a clear pattern of divergence between winners and losers. Workers with no formal qualification become significantly less satisfied compared to all workers who hold at least a GCSE when their sector digitalizes (p ¡ 0.01). This divergence mirrors the pattern with respect to incumbency support.

We find support for the claim that digitalization reduces support for state intervention in the economy among university degree holders. This result is consistent with the possibility that very skilled workers, the main economic beneficiaries of digitalization, adjust their economic preferences in a more pro-market direction, which makes them increasingly attracted to the Tories' program. However, we also find an unexpected result: the group with the lowest qualifications, i.e. unambiguous losers of digitalization, also seem to become less supportive of state intervention. A plausible explanation in light of this specific group's support for UKIP (see Figure 2.4) might be related to the particular social policy position of many right-wing populist parties who strongly differentiate between deserving segments of society (veterans, elderly, "ordinary people") and the rest (Fenger, 2018). Indeed, UKIP has been shown to support insurance-based welfare interventions, especially pensions, but in general opposes a more equity-based, universalist expansion of the welfare state (Ennser-Jedenastik, 2018). It is possible that concerns about deservingness and competition for increasingly scarce welfare benefits is reflected among the lowest skilled group's critical stance on general state intervention that benefits the broader population.

Lastly, we also tested a competing mechanism, namely that digitalization may affect political preferences through changes in attitudes about non-economic issues. It has long been argued that economic modernization and rising living standards increase the importance of non-material goods and help spread social progressiveness on issues such as gender, the environment, or gay rights (Inglehart, 1977). This argument is in conflict with our finding of increased support for the Conservative party and lead us to test the competing hypothesis that increases in digitalization make workers more liberal on social issues. Note that the prediction, if this mechanism holds, would be a shift of winners towards socially progressive parties, such as Labour or the LibDems rather than the Conservative Party. The best suited indicator of socially progressive attitudes available for a sufficiently large number of years in our data is an item battery on support for gender equality. Interestingly, but in line with our main results, we do not find any evidence that changes in digitalization affect progressiveness about gender issues among skilled beneficiaries.

This final result clashes with a common depiction of digitalization winners in the media: the socially progressive celebrity tech entrepreneurs or creators of innovative start-up companies in dynamic urban areas. It is worth reiterating at this point that our analysis is not concerned with such exceptional beneficiaries. We do not study superstars and we do not primarily cover individuals who self-select into thriving technology industries. Our analysis is concerned with the large but less visible group of regular beneficiaries of new technologies who continue to work in their factories, laboratories and offices, become more productive when new digital tools are introduced at their workplace, and benefit from limited but steady improvements of their material conditions.

Our analysis of wage effects has provided strong support for an economic channel linking digitalization and political behavior. Moreover, in light of our auxiliary results on attitudinal variables, an economic voting mechanism seems plausible. Reflecting the polarization of wages, we find a gradient in life satisfaction between winners and losers of digitalization. Furthermore, winners' relatively stable economic situation makes them less supportive of state intervention, especially compared to semiskilled workers with more ambiguous economic prospects. This aspect may help explain their tendency to lean towards center-right rather than center-left incumbents. Finally, we do not find any evidence of particularly progressive values on the cultural dimension. Taken together, ordinary winners of digitzalization are unspectacular supporters of the status quo. For them, mainstream pro-market parties, especially those in government, are a reasonable choice on election day.

2.9 Discussion

The digital revolution is accompanied by two fears: that many workers will be displaced from their jobs and that this will lead to political unrest. Public debate and the scarce academic literature on this topic has primarily been concerned with its downsides and focused on the losers of technological progress. While this focus is comprehensible in light of recent political disruptions, we contend that this one-sided attention is at odds with standard economic theories emphasizing productivity gains as well as with historical experience, which has proved many gloomy projections wrong.

We document both economic and political effects of digitalization. Contrary to pessimistic accounts, a majority of workers benefit economically from rapid digitalization in their industries. Yet, these benefits are not equally distributed and they disproportionately accrue to the highly educated. Our most novel finding is that these diverging economic trajectories are mirrored in diverging political trajectories. First of all, regarding party choice, the beneficiaries of digitalization become more likely to support the Conservative Party, in particular when they are the incumbent party. Second, with respect to turnout, we observe that digitalization reinforces inequalities along education lines: The highly educated turn out more to vote if their sector digitalizes whereas we do not find such mobilizing effects among the less educated. The large but often neglected pool of voters who benefit from technological innovation thus seems willing to support mainstream parties and uphold the existing social contract.

There are several reasons why our results are more optimistic than previous work. First of all, we look at the effects of a general-purpose technology (ICT) on the workforce. This approach is likely to produce different results than if we had focused on more specific technologies, such as industrial robots, that may have particularly strong displacement effects. Indeed, Acemoglu and Restrepo (2020) show that industrial robots have strong negative effects on employment and wages, whereas the effects of increases in other ICT capital, such as computers per worker or investment in software and computers, are often *positive*. Clearly, some technologies have stronger labor-displacement effects, and possibly political effects, than others. We see our contribution as an important complement to studies with a focus on technologies with a more concentrated and more unequivocally negative impact on employment. Our approach allows us to include all sectors rather than mostly manufacturing, a sector which has seen particularly sharp reductions in employment in advanced economies, but is overall rather small (according to the Office for National Statistics, the UK share of people in manufacturing is below 10%). Our coverage of all sectors with a general measure of digitalization possibly facilitates identifying gains of technological change and results in a more optimistic picture.

Another reason why our conclusions may be relatively optimistic is related to our empirical approach. We study the political implications of digitalization on the active labor force, not on the population as a whole, and we focus on individual effects, which can differ from contextual effects. Using a longitudinal approach, we find little indication of political unrest among regular workers. We do not include in our sample retired or disabled people, students or people doing housework, even though workplace digitalization may affect them through various channels including the changes in communities and spillovers within the household. Some segments of this population might react more negatively, e.g. workers who lose their job and cannot find a new one or young citizens with troubles entering the labor market in the first place, although the size of these groups is too small to produce large differences. For these reasons, we do not make inferences based on our findings to population-wide political effects.

To conclude, our findings reveal a complex picture of the political consequences of digitalization. The innovative empirical analysis provides abundant and robust evidence that digitalization is economically beneficial for a majority of the labor force and is politically consequential in two contrasting ways: First, the large group of winners become more likely to support incumbent mainstream parties and thus can act as a stabilizing force in democratic systems. Second, while we only find weak evidence of an anti-establishment backlash among unskilled workers as a reaction to digitalization, we demonstrate that the economic polarization associated with digitalization is accompanied by differential political effects on winners and losers of this process. The resulting divergence in political behavior between the two groups might translate quite directly into increasing political polarization.

For good reasons, much of the reporting on recent political disruptions like Brexit has been on the grievances among the disadvantaged and the likely reasons for their support of leaving the European Union. The Brexit vote should certainly be attributed to a wide range of causes, but it is plausible that the economic and political polarization between beneficiaries of digitalization and other citizens we document in this paper generated political alienation among a subset of the electorate that is exposed to the downsides of economic modernization. While the group of clear-cut losers of digitalization in absolute terms is small, a larger segment of the population in the lower middle class is confronted with relative decline as they observe how others thrive in a digital world while they themselves stagnate.

At the same time, our results remind us that the emergence of anti-establishment forces in most advanced capitalist democracies up to now remains a minority phenomenon. Certainly, how large exactly that minority grows is a question of crucial importance and in some cases, most notably Brexit, anti-establishment forces even managed to mobilize a tight majority of the population. Nevertheless, even in exceptionally disruptive events like Brexit, there was a less attention-grabbing but equally sized group of Remainers who seem content with current circumstances and support the political status quo. All in all, we thus contend that the implications of digitalization at the workplace are more multi-faceted than the narrative of the ”revenge of the left-behind” suggests.

Appendix B

B.1 Description of the data

B.1.1 Summary Statistics

Table B.1: Summary Statistics

	count	mean	sd	min	max
Year	288009	2009.45	5.40	1997	2018
Turnout	108558	0.71	0.46	0	1
Conservatives	233521	0.22	0.41	0	1
Labour	233521	0.33	0.47	0	1
Liberal Democratic Party	233521	0.10	0.29	0	1
UKIP	65920	0.45	0.21	0	1
Incumbent	233521	0.31	0.46	0	1
Industry ID from EUKELMS.	288009			1	38
ICT	257241	3.71	4.58	0.10	47.46
Non-ICT machinery capital stock	257241	27.87	44.20	2.20	540.77
Non-ICT capital stock	257241	133.43	392.35	6.46	4955.94
ICT stock USA / workers in UK	250883	50.28	147.52	0.33	1771.66
Imports in goods from China	40365	9.22	20.77	0.01	189.74
Government region ID	287157			1	13
Female	288009	0.50	0.50	0	1
Born outside the UK	288009	0.03	0.17	0	1
Age	288009	40.55	12.07	18	64
Age squared	288009	1789.68	984.44	324	4096
Education level	288009	4.08	1.52	1	6
Hourly net wage	201830	9.48	5.39	0.00	100.80
Becomes unemployed	224907	0.02	0.15	0	1
Above median RTI	267833	0.47	0.50	0	1
Supports government intervention PCA	69004	-0.06	1.03	-3.22	2.90
Social progressiveness PCA	146729	0.24	1.32	-3.34	3.04
Life satisfaction	262063	5.22	1.28	1	7
Total observations	288009				

Note: ICT defined as ”real fixed ICT capital stock (in 1000 GBP or USD, respectively, in constant 2010 prices) normalized by number of employees”. The Supplementary Information to this article contains a detailed description of the evolution of all dependent variables over time for each educational group.

B.1.2 Dependent Variables by Education

This section presents the longitudinal evolution of our dependent variables between 1997 and 2017, dividing the sample by education level. Figure B.1 plots the average net hourly wage. As in the main analysis, we use constant 2010 prices. The wages of all educational groups have increased over time. In the period until the financial crisis, the growth was largely similar for all income groups, but there is a divergence after the crisis between respondents with university degrees and the rest.

Figure B.1: Average hourly net wage by education

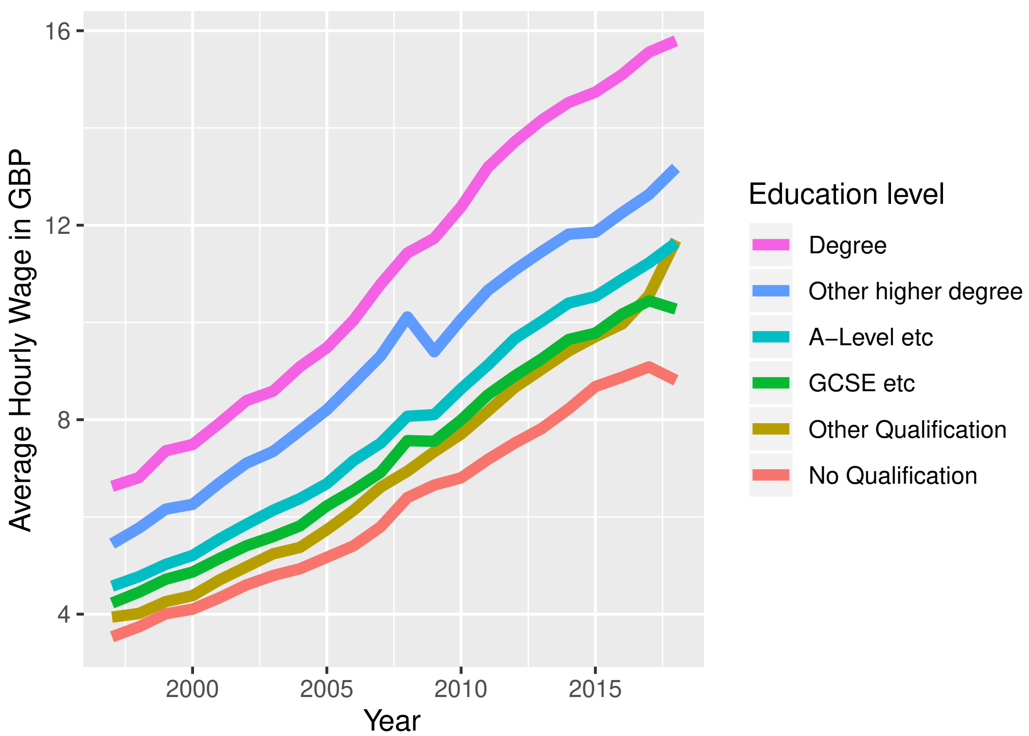

Note: Hourly net wage calculated as monthly net wage in constant 2010 prices normalized by average hour worked. In 2009, BHPS is changed into US which results in the inclusion of new households into the sample.

Figure B.2 presents the percentage of respondents who were unemployed in the week when the interview was conducted. Here again we observe some divergence, as increases in unemployment after the crisis were particularly visible among citizens with less education. Note that unemployment shares in our actual sample are smaller because those who stay unemployed for two periods are not captured by our operationalization.

Figure B.2: Share unemployed by education

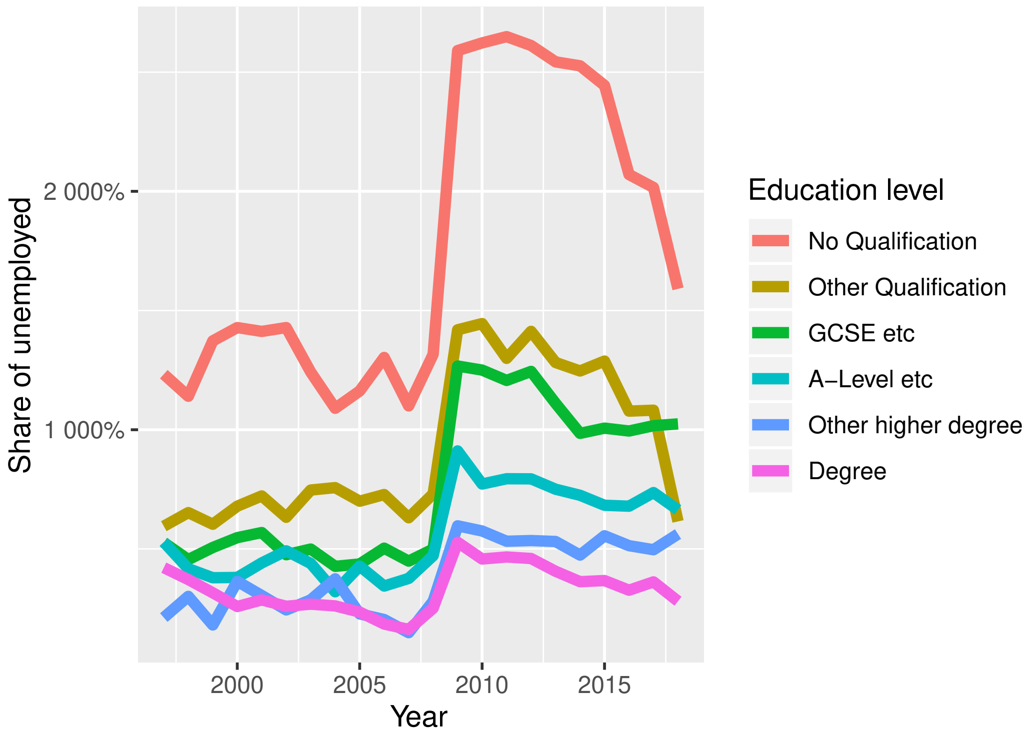

Note: Share unemployed at the time of the interview.

Figure B.3 describes the probability to become unemployed (i.e. to be unemployed at the time of next interview). Again, we see that less educated respondents are more likely to become unemployed and there is an increase after the financial crisis of 2008.

Figure B.4 plots reported turnout for different education levels. Note that this was only asked infrequently after 2008. There was a steady decline in turnout until the mid 2000s and then a partial recovery. Turnout is consistently higher for the highly educated.

Figure B.5 plots the average support for the political parties included in the analyses: the Conservative Party, the Labour Party, as well as the Liberal-Democratic Party, and UKIP (since 2013). We observe a markedly different evolution of support for parties for different education groups, with support for the Conservatives having grown most among workers with university degrees, at the expense of the Liberal-Democratic Party. Some of the time trends will be captured by the year fixed effects.

Figure B.3: Probability to become unemployed in the next period by education

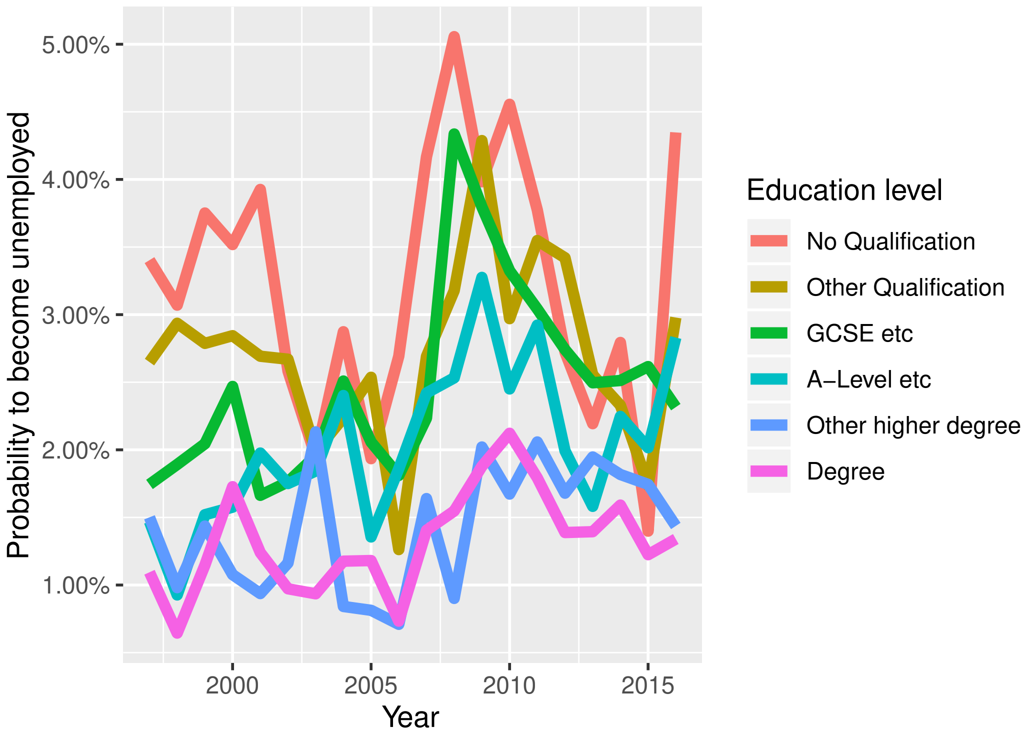

Note: Average probability to become unemployed in the next interview for different education groups. Currently unemployed and respondents without any industry assignment are excluded to ensure equivalence with the main analysis. In 2009, BHPS is changed into US which results in the inclusion of new households into the sample.

Crosswalking and Merging Data Sets

The BHPS, UKHLS and the EU KLEMS datasets are provided using different classifications, which we address by constructing cross-walks. We are able to match the 2007 version of the Standard Industrial Classification (SIC07), used between 2009 and 2015 comprehensively to the classification scheme used by EU KLEMS (NACE Rev. 2). We also manually construct cross-walks from SIC 1992, used in 1994, 1997 and from 2001 to 2008, and are able to match the vast majority of respondents. Between 1991 and 2001 the BHPS used the SIC 1980, which differs markedly from the following versions. We use another crosswalk to translate SIC-80 codes into SIC-92 codes, which then allows to merge the remaining years of EU-KLEMS data. This procedure generates an individual-level data set with information on ICT capital per industry ranging from 1997 to 2017.

Figure B.4: Reported voter turnout by education

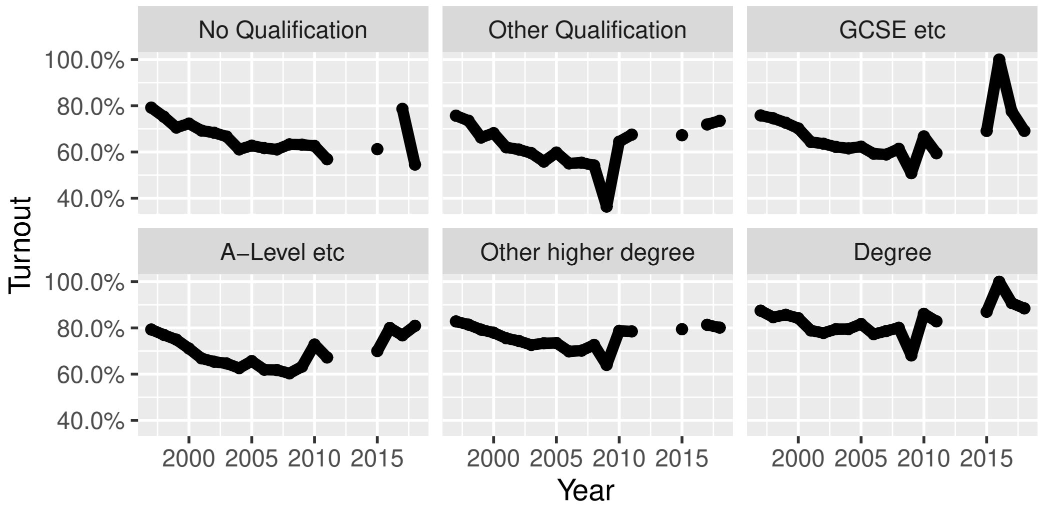

Note: Participation in elections was asked in all waves of BHPS which ended in 2008. In the Understanding Society Survey, participation in elections was only asked in 2010, 2015 and to the few participants of the latest wave who were interviewed after the snap-elections of 2017 which makes the group averages less representative of the election turnout of the whole education group. This does not affect our main results as we focus at within-individual variation.

Figure B.5: Support for political parties by education

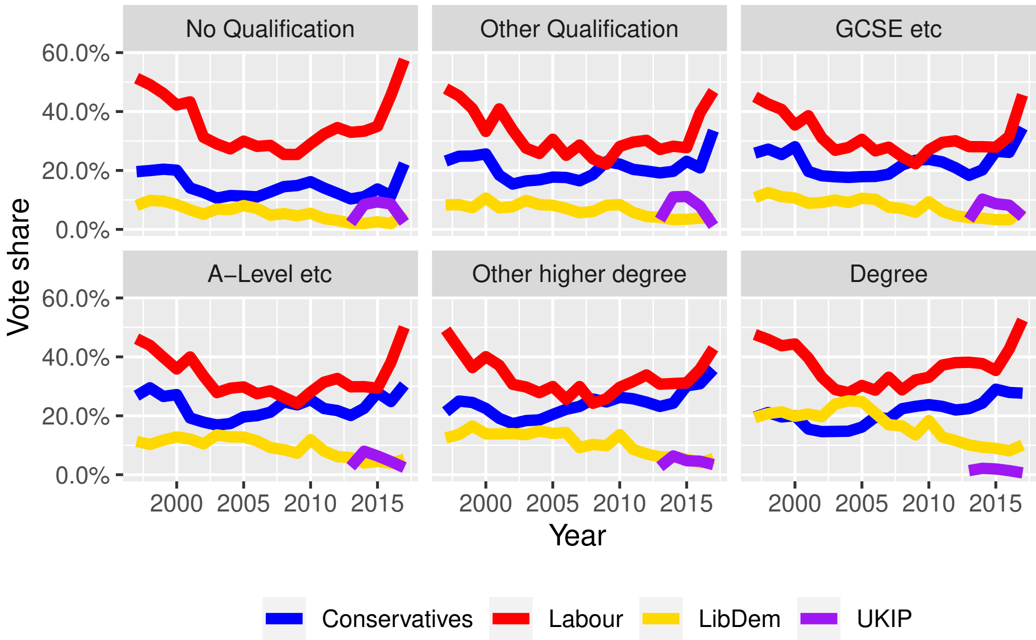

Note: Vote shares calculated based on sample responses answering they voted for the respective party divided by the number of responses for any party including other parties not reported here.

B.2 Comparison of RTI and education as key dimension

In this section, we show that while education is a strong moderator predicting if workers stand to gain or lose from workplace digitalization, RTI seems to be less relevant.

Specifically, we created occupation-specific RTI scores from ONET data following the standard approach of Autor and Dorn (2013), i.e. subtracting log abstract and log manual content from log routine content of each occupation, and relying on a crosswalk by Hardy and colleagues (2018) to merge data with European occupational codes. We then split the observations in high and low RTI groups if they are above or below the median of RTI in the sample.

Figure B.6 shows that high RTI workers in general benefit less from digitalization in terms of wages, as we would expect, but the differences are not statistically significant. By contrast, the strong education gradient suggests that digitalization affect highly and less educated workers in very heterogeneous ways. We learn from this analysis that when looking at individual trajectories, education seems to be a more important source of heterogeneity in the impact of digitalization than RTI.

Given the strong emphasis in the economics literature on the distinction between routine and non-routine occupations, this finding is somewhat surprising. However, this literature looks mostly at aggregate level economic outcomes and we discuss in the text several reasons why our within-individual effects may diverge. We believe that education may be a better proxy than RTI for the ability of workers to adapt to and benefit from digitalization. RTI may predict which jobs are more likely to be partially or fully conducted by machines, but it does not predict well if the individual worker performing a job will benefit or lose from digitalization. The difference between the aggregate level and micro level results are worth further empirical exploration.

In any case, the empirical findings reported here are a strong motivation for our decision of concentrating on education as the key moderator of the effects of workplace digitalization on economic and political outcomes.

Figure B.6: Main outcomes split by high and low RTI

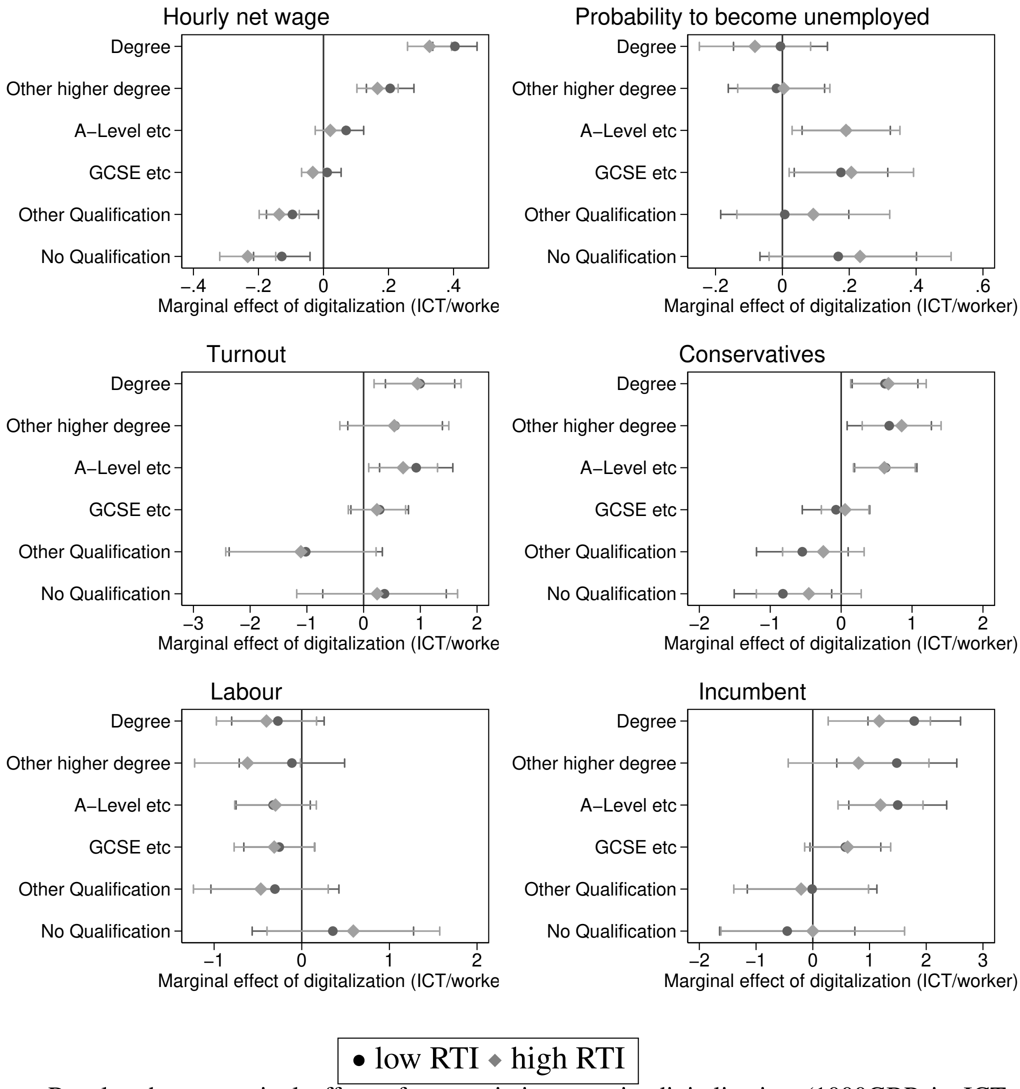

Note: Results show marginal effect of one unit increase in digitalization (1000GBP in ICT capital/worker) on hourly wage, probability to become unemployed and probability to report to have voted or support a given political party. All results except for the hourly wage are in percentage points. High RTI and low RTI is defined relative to the median RTI of the sample.

B.3 Economic Effects Before and After the 2010 Government Change

Table B.2 shows a sub-period analysis for our economic outcomes. It compares the results for hourly net wages and the probability to become unemployed for the time before and after the government change in 2010. The results are comparable to the composite effects. Main difference seems to be that in the 2010 onward period, low educated workers did not seems to lose out in terms of wages in absolute term when they were effected by digitalization. Nevertheless, digitalization decreased their relative wage performance as the effect of digitalization on the wages of the higher educated increases over time.

Table B.2: Economic effects pre and post Government change in May 2010

	Hourly Wage		Unemployment	
	(1)	(2)	(3)	(4)
	Pre May 2010	Post May 2010	Pre May 2010	Post May 2010
Degree × ICT	0.327***	0.302***	-0.0641	0.108
	(0.0350)	(0.0484)	(0.113)	(0.124)
Other higher degree × ICT	0.169***	0.207***	-0.138	0.0625
	(0.0479)	(0.0431)	(0.0843)	(0.140)
A-Level etc × ICT	0.0518	0.103*	0.221*	0.297*
	(0.0274)	(0.0424)	(0.110)	(0.147)
GCSE etc × ICT	-0.0300	0.0894*	0.222*	0.189
	(0.0216)	(0.0392)	(0.102)	(0.173)
Other Qualification × ICT	-0.116**	-0.00793	0.0620	-0.00930
	(0.0371)	(0.0612)	(0.119)	(0.209)
No Qualification × ICT	-0.206***	-0.0381	0.229	0.172
	(0.0490)	(0.0693)	(0.131)	(0.237)
Degree	-1.387***	-1.800***	3.188*	-0.509
	(0.242)	(0.353)	(1.352)	(1.778)
Other higher degree	-1.442***	-1.773***	4.476**	0.934
	(0.273)	(0.323)	(1.427)	(1.581)
A-Level etc	-1.265***	-1.257***	1.406	-0.571
	(0.164)	(0.293)	(1.085)	(1.550)
GCSE etc	-0.765***	-1.089***	1.797	-0.264
	(0.167)	(0.278)	(1.098)	(1.447)
Other Qualification	-0.333*	-0.491	1.541	1.486
	(0.142)	(0.251)	(0.926)	(1.547)
Age	0.339***	0.455***	-0.250	-0.367
	(0.0281)	(0.0487)	(0.130)	(0.206)
Age × Age	-0.00296***	-0.00422***	-0.000344	0.00118
	(0.000261)	(0.000329)	(0.000962)	(0.00133)
Constant	-3.759***	-1.880	6.771	14.92
	(0.871)	(1.757)	(4.142)	(8.167)
Id*Ind FE	X	X	X	X
Year FE	X	X	X	X
Region	X	X	X	X
Observations	85782	93695	100612	115518

Note: All columns use our main specification. Column (1) and (2) report a sub-period analysis for net hourly wages (calculated as monthly net wage in constant 2010 prices normalized by average hour worked. Column (3) and (4) report a sub-period analysis for probability to become unemployed in percentage points (ie. to be unemployed at the next interview conditional on currently working). Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.4 Robustness checks in detail

This section extends the discussion about the robustness checks offered in the text. The full regression tables are presented at the end of this section.

Non-ICT capital investment

First, we need to rule out the possibility that an increase in ICT capital stocks simply reflects the fact that booming industries have a larger capacity to invest and offer their workers higher wages and better conditions. If the general propensity to invest of a sector has an effect on workers' economic outcomes and political preferences, this could invalidate our interpretation of our results. They would not capture the specific consequences of digitalization but rather the effect of working in a thriving industry.

To assess this possibility, we conduct an additional analysis using non-ICT capital stock per worker as the main explanatory variable:

$$\text{Non-ICT capital intensity}_{jt} = \frac{\text{Total capital stock}_{jt} \text{ - ICT capital stock}_{jt}}{(\text{Employees}_{jt})}$$

Changes in an industry's non-ICT capital stock do not predict any of the outcomes we are interested in. As can be seen in column (3) in the tables presented in this section, the coefficients are very small and imprecisely estimated. This was to be expected since we argued that investment in digitalization substitutes or complements labor in a specific way depending on their skill level. The same is not true for other kinds of capital investments (e.g. building a new production plant or buying a new office building).

This result increases our confidence in the interpretation that the main results are driven specifically by ICT capital, since other kinds of capital do not affect workers' political preferences in a similar way.

In addition, we have tested more specific aggregations of residual asset categories among the non-ICT group. Certain asset categories we categorize as non-ICT but might not be seen as "digital" assets but still relate to technological change more broadly, e.g. other machinery equipment besides ICT equipment. As we argue in the manuscript, our goal is to specifically study the impact of digitalization, not the impact of the broader and more elusive concept of technological change. That said, since the data allows for more fine-grained analysis, we have explored further operationalizations to examine implications for the presented main results. We replicated our analysis with a dependent variable consisting only of the two categories related to non-digital machinery ("transportation", "other machinery equipment and weapons"). We find that investment in machinery has somewhat comparable economic effects in that it has positive wage implications on high-skilled workers. However, crucially, the effect sizes are much smaller than the effects of ICT investment. In terms of standard deviations, a one standard deviation in ICT capital stocks produces an increase of 0.25 GBP per hour worked among workers with university degrees, but non-ICT machinery only translates into an increase of 0.05 GBP per hour. Consequently, and unsurprisingly, these much smaller effects do not translate into changes in workers' political behavior. In line with the original non-ICT analysis, we do not find any evidence that investment in machinery affects political outcome variables.

B.4.1 Excluding industry and regional outliers

One might object that our results could be driven by a few rapidly digitalizing industries. To rule out this possibility, we excluded the three industries with the largest increase in digitalization in recent years (Telecommunications, Mining and Quarrying and Coke, Refined petroleum) in the models in column (4). The exclusion of these outliers does not change results. If anything, it even increases the precision of our estimates.

Relatedly, our results could also be driven by some particularly rapidly digitalizing regions such as the metropolitan area of London. To account for this, we include separate set of time fixed effects for each region. Column (5) in the tables presented in the SI confirms that the results are not driven by these regions, as point estimates remain largely unchanged for all outcomes while standard errors decrease for some outcomes.

B.4.2 Lead models and simple fixed effects

Another key concern is that our models are too restrictive towards losers and thus may underestimate the effects of digitalization because they miss the negative effects on workers who are displaced by digitalization and do not work in the same industry in the next period when they are re-interviewed. This could happen for two different reasons. If displaced workers drop out of the labor force they would not be assigned to an industry in the next interview and would therefore drop out of our analysis. If they switch to a different industry, the industry-spell fixed effects would absorb part of the effect of job displacement on economic and political outcomes. In any case, our models may fail to capture the effects of digitalization on some displaced workers workers.

We deal with this concern by relaxing the sample restriction in two ways and thus potentially capturing more losers: First, we replicate all analyses using lead models in which we examine how our measure of digitalization affects labor market and political outcomes measured at the time of the next interview. In this way, we keep in our sample all workers who may have been displaced by digitalization (and either exit the labor force or work in a different industry). This results in a slightly smaller sample (because we lose the last year), but the coefficients reported in column (6) confirm that the results remain unchanged when using leads. The only exception is voter turnout, as several of the coefficients of interest become statistically non-significant.

Second, we replicate all analyses using a unique individual fixed effect by respondent instead of industry-spell fixed effects. Using this approach, workers who change industries (perhaps in response to job displacement due to technology) contribute to the average estimates of the effect of digitalization on labor market and political outcomes, although workers who drop out of the labor force entirely are still excluded from the sample. The results are reported in column (7) in the full tables below. Although the polarizing effect of digitalization on wages is still clearly visible, this specification results in smaller estimates of the effects of digitalization on hourly pay for both highly and less educated workers. This was to be expected as using unique individual fixed effects adds measurement error to our explanatory variable which causes attenuation bias in the estimated coefficients.[1] An alternative explanation is that economic benefits of digitalization are reaped mostly by educated workers who stay in their industries while the costs may be borne also by less educated workers who choose to stay in the same industries. Using this specification, we do not find effects of digitalization on voter turnout, but we still observe that digitalization is associated with increased support for the Conservatives and the incumbent party among workers with more education.

B.4.3 Including controls for trade

A possible threat to identification is that our indicator of technology may be correlated with changes in international trade in an industry. In that case, our estimates would partially capture effects of international trade on economic outcomes and political behavior. However, previous work on the geography of trade shocks and technological change in the US shows that the two types of shocks have largely distinct distributions in space (Autor et al., 2015), suggesting that there is limited overlap. In any case, we replicate all the analysis controlling for international trade in the industries for which we can collect data. Specifically, we use yearly UN Comtrade data on exports from China to the UK

[1] The variation in digitalization created by industry switches is much larger than the year to year variation for stayers which is problematic for two reasons. First, frequent back and forth switches between two industries within individuals is possibly due to measurement error in the interviews. Second, we theorize that a digitalizing workplace is what affects political attitudes, not the jumps when switching between highly and low digitalized industries.

as an indicators of international trade.[2] This measure is only available for manufacturing industries, resulting in a much smaller sample size. The results presented in column (8) of the complete tables show that the results remain unchanged when controlling for changes in trade within the industries for which data are available.

B.4.4 Cross-sectional OLS

For the sake of completeness, we also add a cross-sectional OLS regression including only industry and year fixed effects to see how between-worker differences in ICT intensity relate to our outcomes (column 9). Results have to be interpreted with a large grain of salt as we now cannot control for unobserved worker-level characteristics anymore. Instead, except for the inclusion of a gender dummy, we tried to stay as close as possible to our main specification to ensure the comparability of results while avoiding post-treatment bias. The results for political outcomes are surprisingly similar to the fixed-effects specification. Especially, they confirm the finding that digitalization increase support for the Conservatives for the incumbent among highly educated workers.

Regarding economic outcomes, the results change slightly. The highly educated are still the main beneficiaries when it comes to wages. However, looking at unemployment, less educated people already working in digitalized industries appear to benefit from digitalization as they have lower probabilities to become unemployed. This is somewhat counter-intuitive and seemingly opposite to our findings from the baseline specification. Yet, the two diverging results make sense considering the different nature of the two analyses. The cross-sectional analysis shows that working in an already digitalized industry reduces the risk of unemployment whereas the fixed-effects specification shows that for a given worker in a given industry, increasing digitalization might threaten the jobs of less educated workers if tasks are automated. We interpret this more nuanced reading as a validation that it is important to only consider within-individual variation if we want to study how a *given worker* is affected when his or her work environment digitalizes.

[2]The data is provided for different types of goods which we first crosswalk to SIC and from there to NACE rev. 2 codes which is used in EUKLEMS.

Table B.3: Net hourly wages in GBP

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	Main	IV	Placebo	Region*Year FE	Excl. outliers	Lead	ID FE	Trade	Cross Sect
Degree × ICT	0.343***	0.435***	-0.000809	0.331***	0.432***	0.307***	0.153***	0.478***	0.133***
	(0.0324)	(0.0809)	(0.000705)	(0.0307)	(0.0410)	(0.0349)	(0.0161)	(0.0802)	(0.00866)
Other higher degree × ICT	0.184***	0.301***	-0.000485	0.182***	0.225***	0.174***	0.109***	0.328***	0.104***
	(0.0336)	(0.0745)	(0.000539)	(0.0331)	(0.0435)	(0.0337)	(0.0165)	(0.0627)	(0.00921)
A-Level etc × ICT	0.0514*	0.104	-0.000726	0.0496*	0.0824*	0.0651*	0.0720***	0.124*	0.131***
	(0.0229)	(0.0860)	(0.000449)	(0.0227)	(0.0362)	(0.0255)	(0.0143)	(0.0542)	(0.00787)
GCSE etc × ICT	-0.0114	-0.0477	-0.000728	-0.0141	-0.00711	-0.00707	0.0462***	0.0119	0.114***
	(0.0185)	(0.0598)	(0.000432)	(0.0185)	(0.0282)	(0.0208)	(0.0130)	(0.0422)	(0.00808)
Other Qualification × ICT	-0.135***	-0.228**	-0.00109*	-0.141***	-0.145***	-0.122***	0.0300	-0.0968	0.0972***
	(0.0288)	(0.0876)	(0.000498)	(0.0286)	(0.0347)	(0.0305)	(0.0177)	(0.0588)	(0.0105)
No Qualification × ICT	-0.185***	-0.305***	-0.00128**	-0.188***	-0.212***	-0.109**	-0.00863	-0.224*	0.0351**
	(0.0398)	(0.0894)	(0.000444)	(0.0391)	(0.0500)	(0.0415)	(0.0209)	(0.0965)	(0.0112)
Degree	-1.995***	-2.513***	-0.679***	-1.953***	-2.243***	-1.675***	-1.125***	-2.793***	4.712***
	(0.209)	(0.308)	(0.178)	(0.209)	(0.225)	(0.215)	(0.169)	(0.628)	(0.0457)
Other higher degree	-2.028***	-2.622***	-1.242***	-2.019***	-2.148***	-1.876***	-1.603***	-2.733***	2.714***
	(0.218)	(0.294)	(0.179)	(0.219)	(0.232)	(0.227)	(0.177)	(0.702)	(0.0442)
A-Level etc	-1.628***	-1.970***	-1.276***	-1.604***	-1.707***	-1.496***	-1.409***	-1.899***	1.584***
	(0.156)	(0.250)	(0.130)	(0.158)	(0.171)	(0.162)	(0.135)	(0.363)	(0.0373)
GCSE etc	-1.141***	-1.254***	-0.903***	-1.128***	-1.179***	-0.978***	-1.000***	-1.492***	0.976***
	(0.147)	(0.218)	(0.127)	(0.150)	(0.158)	(0.148)	(0.130)	(0.358)	(0.0351)
Other Qualification	-0.441**	-0.420	-0.448***	-0.408**	-0.458**	-0.395**	-0.521***	-0.490	0.436***
	(0.137)	(0.222)	(0.112)	(0.137)	(0.144)	(0.135)	(0.118)	(0.332)	(0.0419)
Age	0.345***	0.346***	0.391***	0.374***	0.343***	0.334***	0.369***	0.235***	0.446***
	(0.0271)	(0.0277)	(0.0280)	(0.0275)	(0.0271)	(0.0316)	(0.0260)	(0.0588)	(0.00527)
Age × Age	-0.00312***	-0.00311***	-0.00331***	-0.00315***	-0.00307***	-0.00345***	-0.00332***	-0.00184***	-0.00454***
	(0.000212)	(0.000220)	(0.000217)	(0.000212)	(0.000213)	(0.000241)	(0.000190)	(0.000431)	(0.0000678)
Imports								-0.00292	
								(0.00331)	
Dummy=1 if person identifies as female									-1.189***
									(0.0210)
Constant	-2.821***	-2.585**	-3.960***	-3.439***	-2.667***	-1.817*	-3.356***	-0.642	-7.709***
	(0.797)	(0.832)	(0.850)	(0.845)	(0.799)	(0.873)	(0.777)	(1.756)	(0.145)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	179477	174723	179477	179477	176659	153751	178458	32817	179477

Note: Hourly net wage calculated as monthly net wage in constant 2010 prices normalized by average hours worked. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.4: Probability to become unemployed

	(1) Main	(2) IV	(3) Placebo	(4) Region*Year FE	(5) Excl. outliers	(6) Lead	(7) ID FE	(8) Trade	(9) Cross Sect
Degree × ICT	0.0129	0.241	0.000297	0.0157	-0.0849	0.124	0.0715	-0.0231	0.0182
	(0.0713)	(0.197)	(0.000635)	(0.0711)	(0.0812)	(0.0816)	(0.0454)	(0.120)	(0.0246)
Other higher degree × ICT	0.00620	0.354	0.000337	0.0161	-0.0601	0.0601	0.0713	-0.198	0.0201
	(0.0644)	(0.211)	(0.000602)	(0.0646)	(0.101)	(0.0764)	(0.0692)	(0.147)	(0.0289)
A-Level etc × ICT	0.168**	0.421*	0.000642	0.180**	0.152*	0.122	0.159**	0.0697	0.0190
	(0.0608)	(0.203)	(0.000689)	(0.0612)	(0.0711)	(0.0926)	(0.0533)	(0.139)	(0.0280)
GCSE etc × ICT	0.183**	0.631	0.000129	0.186**	0.176	0.257**	0.101	0.112	-0.00192
	(0.0686)	(0.413)	(0.000664)	(0.0686)	(0.0917)	(0.0843)	(0.0518)	(0.134)	(0.0296)
Other Qualification × ICT	0.0451	0.572*	0.000118	0.0496	0.00768	0.0158	-0.0195	-0.368	-0.0196
	(0.0807)	(0.274)	(0.00125)	(0.0807)	(0.0924)	(0.109)	(0.0864)	(0.195)	(0.0437)
No Qualification × ICT	0.227*	0.620	0.000146	0.225*	0.241	0.259	-0.0423	0.0821	-0.0628
	(0.106)	(0.446)	(0.00111)	(0.106)	(0.138)	(0.149)	(0.0903)	(0.168)	(0.0462)
Degree	0.883	1.496	0.208	0.872	1.203	0.258	-2.162*	3.255	-2.314***
	(0.793)	(1.258)	(0.739)	(0.794)	(0.817)	(0.946)	(0.840)	(2.059)	(0.214)
Other higher degree	1.446	1.549	0.812	1.450	1.655*	1.393	-1.199	3.775	-1.927***
	(0.778)	(1.174)	(0.726)	(0.776)	(0.817)	(0.934)	(0.856)	(2.972)	(0.222)
A-Level etc	0.607	1.169	0.465	0.593	0.685	0.750	-0.855	0.563	-1.765***
	(0.691)	(1.094)	(0.634)	(0.691)	(0.720)	(0.846)	(0.743)	(1.525)	(0.217)
GCSE etc	0.773	0.741	0.676	0.757	0.835	0.499	-0.478	0.657	-1.093***
	(0.657)	(1.183)	(0.596)	(0.655)	(0.692)	(0.810)	(0.711)	(1.809)	(0.217)
Other Qualification	1.124	0.900	0.625	1.089	1.238	1.561*	-0.0571	2.275	-0.702**
	(0.652)	(0.964)	(0.584)	(0.653)	(0.670)	(0.764)	(0.709)	(1.985)	(0.259)
Age	-0.435***	-0.442***	-0.435***	-0.441***	-0.445***	-0.195	-0.580***	-0.383	-0.468***
	(0.0994)	(0.101)	(0.102)	(0.102)	(0.100)	(0.111)	(0.106)	(0.253)	(0.0238)
Age × Age	0.00158**	0.00166**	0.00152*	0.00154*	0.00154*	0.00259***	0.00407***	0.00217	0.00489***
	(0.000604)	(0.000624)	(0.000602)	(0.000604)	(0.000606)	(0.000719)	(0.000600)	(0.00172)	(0.000275)
Imports								0.00529	
								(0.0140)	
Dummy=1 if person identifies as female									-0.530***
									(0.0749)
Constant	13.76***	13.16***	14.68***	14.41***	14.13***	3.520	17.79***	-0.382	13.11***
	(3.681)	(3.827)	(3.792)	(3.811)	(3.712)	(3.300)	(3.691)	(8.508)	(0.641)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	216130	210708	216130	216130	213075	183311	214741	34841	216130

Note: Probability to become unemployed in percentage points among those currently working. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.5: Voted in last general elections

	(1) Main	(2) IV	(3) Placebo	(4) Region*Year FE	(5) Excl. outliers	(6) Lead	(7) ID FE	(8) Trade	(9) Cross Sect
Degree × ICT	0.635*	1.396*	0.00864	0.562*	1.005**	0.353	0.364*	0.349	-0.00195
	(0.282)	(0.623)	(0.00540)	(0.281)	(0.376)	(0.280)	(0.153)	(0.725)	(0.110)
Other higher degree × ICT	0.305	2.299*	0.00726	0.293	0.806	0.656	0.230	0.145	-0.145
	(0.366)	(1.052)	(0.00458)	(0.364)	(0.553)	(0.397)	(0.187)	(0.629)	(0.132)
A-Level etc × ICT	0.691**	1.998*	0.00704	0.718**	0.976**	1.073***	0.460**	-0.0351	0.175
	(0.264)	(0.992)	(0.00644)	(0.264)	(0.365)	(0.290)	(0.153)	(0.558)	(0.116)
GCSE etc × ICT	0.211	1.186	-0.00261	0.180	-0.235	0.295	0.335*	0.0326	0.164
	(0.231)	(0.983)	(0.00524)	(0.229)	(0.396)	(0.256)	(0.155)	(0.508)	(0.119)
Other Qualification × ICT	-0.951	1.863	-0.00671	-1.000	-0.839	-0.180	0.207	-2.007	-0.430*
	(0.575)	(1.860)	(0.00766)	(0.575)	(0.558)	(0.417)	(0.251)	(1.240)	(0.182)
No Qualification × ICT	0.148	2.235	0.00204	0.205	0.277	0.536	0.637*	-1.262	0.342
	(0.470)	(3.141)	(0.00570)	(0.468)	(0.672)	(0.489)	(0.265)	(0.810)	(0.189)
Degree	-0.617	2.391	-0.874	-0.274	-1.302	-2.443	-1.022	-18.41	22.60***
	(3.336)	(6.397)	(3.257)	(3.334)	(3.436)	(3.478)	(2.900)	(9.758)	(0.723)
Other higher degree	-2.424	-2.807	-3.678	-2.658	-3.241	-5.876	-3.094	-22.62*	15.12***
	(4.038)	(6.886)	(3.821)	(4.016)	(4.215)	(4.131)	(3.439)	(11.41)	(0.782)
A-Level etc	-5.519	-3.938	-5.114	-5.339	-5.875*	-6.698*	-3.732	-14.13*	10.94***
	(2.846)	(5.949)	(2.719)	(2.836)	(2.958)	(2.932)	(2.540)	(6.130)	(0.704)
GCSE etc	-4.484	-2.404	-4.343	-4.265	-3.445	-4.106	-5.122*	-9.008	6.196***
	(2.881)	(5.920)	(2.750)	(2.871)	(3.041)	(3.081)	(2.488)	(6.064)	(0.695)
Other Qualification	0.548	-1.107	-1.092	0.888	0.771	1.274	-0.848	5.268	2.753**
	(2.274)	(6.177)	(2.029)	(2.243)	(2.314)	(2.373)	(1.811)	(5.314)	(0.844)
Age	-1.143**	-1.112**	-0.505	-0.553	-1.185**	0.455	-1.002**	-2.001*	1.979***
	(0.390)	(0.404)	(0.396)	(0.396)	(0.393)	(0.398)	(0.354)	(0.945)	(0.0801)
Age × Age	-0.00913***	-0.00951**	-0.00981***	-0.00947***	-0.00870**	-0.00968***	-0.00919***	-0.000692	-0.0114***
	(0.00264)	(0.00290)	(0.00263)	(0.00263)	(0.00266)	(0.00272)	(0.00229)	(0.00631)	(0.000959)
Imports								-0.0674	
								(0.0619)	
Dummy=1 if person identifies as female									0.152
									(0.302)
Constant	133.1***	129.5***	112.8***	113.5***	134.2***	80.16***	125.5***	162.7***	14.20***
	(12.47)	(14.40)	(12.88)	(12.91)	(12.60)	(13.44)	(12.47)	(31.78)	(2.202)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	103739	100881	103739	103739	102060	91381	102642	19183	103739

Note: Probability to report to have voted in last general election in percentage point. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.6: Support for the Conservative Party

	(1) Main	(2) IV	(3) Placebo	(4) Region*Year FE	(5) Excl. outliers	(6) Lead	(7) ID FE	(8) Trade	(9) Cross Sect
Degree × ICT	0.589**	2.198**	0.00334	0.533**	0.793**	0.548**	0.366***	1.240	0.282***
	(0.196)	(0.673)	(0.00276)	(0.195)	(0.275)	(0.202)	(0.100)	(0.634)	(0.0727)
Other higher degree × ICT	0.540*	1.759*	0.00711*	0.494*	0.975**	0.831**	0.153	-0.00127	0.0784
	(0.240)	(0.696)	(0.00334)	(0.238)	(0.309)	(0.257)	(0.124)	(0.553)	(0.0837)
A-Level etc × ICT	0.580**	1.513*	0.00653*	0.538**	1.078***	0.579**	0.295**	0.384	0.137
	(0.193)	(0.592)	(0.00300)	(0.190)	(0.277)	(0.196)	(0.101)	(0.343)	(0.0754)
GCSE etc × ICT	-0.0288	0.917	0.000506	-0.0719	0.428	0.158	0.166	-0.948*	0.203*
	(0.191)	(0.657)	(0.00296)	(0.188)	(0.253)	(0.177)	(0.109)	(0.389)	(0.0791)
Other Qualification × ICT	-0.358	1.468	-0.00432	-0.457	-0.240	-0.344	0.0478	-0.670	-0.107
	(0.268)	(0.996)	(0.00566)	(0.276)	(0.328)	(0.265)	(0.142)	(0.533)	(0.108)
No Qualification × ICT	-0.601*	0.443	-0.00163	-0.575*	-0.638	-0.278	-0.225	-1.062	-0.247*
	(0.278)	(1.073)	(0.00422)	(0.278)	(0.347)	(0.271)	(0.160)	(0.819)	(0.110)
Degree	-7.420***	-8.232**	-5.513**	-7.281***	-7.551***	-6.832***	-5.144**	-20.95***	8.362***
	(1.937)	(3.102)	(1.836)	(1.939)	(1.989)	(1.907)	(1.607)	(5.703)	(0.440)
Other higher degree	-5.326**	-5.324	-4.087*	-5.380**	-6.157**	-8.448***	-3.591*	-8.627	11.21***
	(2.053)	(3.238)	(1.881)	(2.044)	(2.121)	(2.090)	(1.680)	(6.362)	(0.485)
A-Level etc	-6.227***	-5.190	-4.822**	-6.208***	-7.259***	-7.810***	-4.711**	-7.653	9.361***
	(1.786)	(2.763)	(1.675)	(1.796)	(1.839)	(1.698)	(1.488)	(4.018)	(0.431)
GCSE etc	-3.577*	-3.018	-3.093	-3.582*	-4.545*	-5.510***	-3.834**	-0.402	7.040***
	(1.744)	(2.822)	(1.648)	(1.753)	(1.791)	(1.660)	(1.427)	(4.325)	(0.428)
Other Qualification	-0.00495	-1.629	0.462	0.270	-0.154	-1.442	-0.641	1.448	4.148***
	(1.703)	(2.942)	(1.447)	(1.693)	(1.749)	(1.593)	(1.297)	(5.503)	(0.522)
Age	0.383	0.354	0.631**	0.584*	0.386	-0.0277	0.238	0.814	0.144**
	(0.226)	(0.232)	(0.230)	(0.230)	(0.227)	(0.241)	(0.208)	(0.572)	(0.0480)
Age × Age	-0.00330*	-0.00276	-0.00356*	-0.00313	-0.00301	-0.00443*	-0.00149	-0.00531	0.00278***
	(0.00163)	(0.00170)	(0.00163)	(0.00163)	(0.00164)	(0.00178)	(0.00139)	(0.00419)	(0.000596)
Imports								0.0164	
								(0.0277)	
Dummy=1 if person identifies as female									-0.116
									(0.192)
Constant	11.99	10.43	4.270	6.185	11.75	25.44**	11.35	-1.467	11.10***
	(7.639)	(8.107)	(7.866)	(7.887)	(7.696)	(8.726)	(6.824)	(17.59)	(1.474)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	221050	215784	221050	221050	218065	189046	219758	34586	221050

Note: Probability to report to support the Conservative Party in percentage point. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.7: Support for the Labour Party

	(1) Main	(2) IV	(3) Placebo	(4) Region*Year FE	(5) Excl. outliers	(6) Lead	(7) ID FE	(8) Trade	(9) Cross Sect
Degree × ICT	-0.203	0.324	-0.00242	-0.210	-0.0700	-0.120	-0.185	-0.921	-0.441***
	(0.214)	(0.529)	(0.00350)	(0.213)	(0.281)	(0.223)	(0.103)	(0.849)	(0.0802)
Other higher degree × ICT	-0.124	0.272	-0.00102	-0.125	-0.183	-0.231	0.0584	-0.682	-0.218*
	(0.237)	(0.666)	(0.00415)	(0.240)	(0.321)	(0.321)	(0.117)	(0.434)	(0.0910)
A-Level etc × ICT	-0.229	-0.550	-0.00402	-0.213	-0.459	-0.260	-0.207	-0.279	-0.399***
	(0.191)	(0.532)	(0.00414)	(0.190)	(0.275)	(0.211)	(0.109)	(0.583)	(0.0826)
GCSE etc × ICT	-0.206	0.464	-0.00433	-0.193	-0.500	-0.268	-0.208	0.609	-0.576***
	(0.188)	(0.599)	(0.00422)	(0.188)	(0.267)	(0.186)	(0.114)	(0.569)	(0.0885)
Other Qualification × ICT	-0.473	0.451	-0.00512	-0.455	-0.767	-0.282	0.0199	0.722	-0.276*
	(0.345)	(0.976)	(0.00726)	(0.343)	(0.406)	(0.341)	(0.168)	(0.707)	(0.121)
No Qualification × ICT	0.402	0.216	-0.00628	0.357	0.297	-0.0357	0.196	0.564	-0.0417
	(0.391)	(1.761)	(0.00393)	(0.389)	(0.513)	(0.469)	(0.217)	(0.550)	(0.148)
Degree	2.319	0.601	0.480	2.000	1.350	0.598	2.550	5.599	0.188
	(2.371)	(4.350)	(2.182)	(2.370)	(2.469)	(2.576)	(2.094)	(6.375)	(0.577)
Other higher degree	0.522	-0.803	-1.018	0.424	0.271	-2.260	-0.756	1.857	-4.898***
	(2.439)	(4.405)	(2.214)	(2.444)	(2.553)	(2.721)	(2.160)	(5.875)	(0.612)
A-Level etc	0.879	1.462	-0.819	0.547	0.875	-0.927	0.560	-1.041	-3.063***
	(2.164)	(4.044)	(1.990)	(2.171)	(2.265)	(2.409)	(1.979)	(5.189)	(0.562)
GCSE etc	1.581	-0.428	0.0660	1.358	1.747	-1.159	0.780	-2.184	-3.539***
	(2.028)	(3.919)	(1.895)	(2.033)	(2.128)	(2.253)	(1.849)	(4.682)	(0.559)
Other Qualification	-0.495	-3.125	-2.565	-0.649	-0.0227	-0.994	-1.093	-5.388	-4.708***
	(1.824)	(3.600)	(1.548)	(1.824)	(1.898)	(2.012)	(1.506)	(4.562)	(0.657)
Age	0.128	0.189	0.0542	0.0739	0.0993	0.477	0.142	-0.681	0.538***
	(0.268)	(0.274)	(0.272)	(0.272)	(0.269)	(0.291)	(0.249)	(0.685)	(0.0545)
Age × Age	-0.00453*	-0.00531**	-0.00431*	-0.00447*	-0.00458*	-0.000146	-0.00475**	0.00234	-0.00612***
	(0.00182)	(0.00191)	(0.00180)	(0.00181)	(0.00183)	(0.00198)	(0.00158)	(0.00436)	(0.000664)
Imports								-0.0236	
								(0.0353)	
Dummy=1 if person identifies as female									-1.569***
									(0.215)
Constant	59.78***	59.35***	64.57***	62.59***	61.05***	35.24***	58.44***	71.52***	39.24***
	(9.050)	(10.02)	(9.628)	(9.694)	(9.120)	(9.902)	(8.238)	(21.46)	(1.598)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	221050	215784	221050	221050	218065	189046	219758	34586	221050

Note: Probability to report to support the Labour Party in percentage point. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.8: Support for the Incumbent

	(1) Main	(2) IV	(3) Placebo	(4) Region*Year FE	(5) Excl. outliers	(6) Lead	(7) ID FE	(8) Trade	(9) Cross Sect
Degree × ICT	1.527***	2.877*	0.0130	1.415***	2.493***	1.220***	0.955***	0.446	0.871***
	(0.336)	(1.444)	(0.00733)	(0.324)	(0.499)	(0.363)	(0.172)	(0.809)	(0.0834)
Other higher degree × ICT	1.245*	2.365*	0.0109	1.230**	2.342***	1.293*	0.835***	0.404	0.752***
	(0.514)	(1.183)	(0.00610)	(0.468)	(0.602)	(0.553)	(0.218)	(1.107)	(0.0942)
A-Level etc × ICT	1.333***	2.683**	0.00374	1.259***	2.155***	1.065**	0.844***	-0.627	0.739***
	(0.355)	(0.943)	(0.00568)	(0.331)	(0.440)	(0.395)	(0.189)	(0.798)	(0.0855)
GCSE etc × ICT	0.657*	2.034*	0.000193	0.605*	1.214*	0.857**	0.505**	-0.186	0.555***
	(0.298)	(0.952)	(0.00594)	(0.285)	(0.475)	(0.330)	(0.181)	(0.614)	(0.0908)
Other Qualification × ICT	-0.251	2.645	-0.0150	-0.537	-0.149	-0.352	0.701**	-1.385	0.572***
	(0.534)	(1.776)	(0.00839)	(0.544)	(0.651)	(0.538)	(0.254)	(1.228)	(0.121)
No Qualification × ICT	-0.207	0.556	-0.0225*	-0.306	-0.0166	-0.179	0.544	-0.210	0.155
	(0.567)	(2.139)	(0.0104)	(0.571)	(0.753)	(0.609)	(0.294)	(0.812)	(0.145)
Degree	-12.11***	-12.67*	-11.48***	-11.11**	-13.70***	-12.69***	-10.51***	-20.47*	5.279***
	(3.591)	(6.093)	(3.255)	(3.493)	(3.803)	(3.826)	(3.036)	(8.923)	(0.569)
Other higher degree	-9.677*	-10.15	-9.504**	-8.994*	-11.61**	-14.59***	-8.767**	-22.55*	3.555***
	(3.982)	(5.948)	(3.469)	(3.810)	(4.134)	(4.239)	(3.295)	(10.44)	(0.609)
A-Level etc	-9.460**	-10.63*	-8.815**	-9.015**	-10.80**	-11.58***	-8.770**	-8.456	1.999***
	(3.136)	(5.152)	(2.850)	(3.047)	(3.287)	(3.287)	(2.776)	(7.123)	(0.556)
GCSE etc	-9.527**	-10.57*	-10.13***	-9.276**	-10.23**	-14.04***	-9.810***	-16.11*	0.328
	(3.147)	(5.089)	(2.842)	(3.039)	(3.339)	(3.315)	(2.742)	(6.586)	(0.553)
Other Qualification	-1.458	-6.587	-3.347	-1.282	-0.883	-2.890	-3.239	2.393	-1.624*
	(2.602)	(4.909)	(2.277)	(2.581)	(2.707)	(2.869)	(2.183)	(6.474)	(0.653)
Age	-0.730	-0.739	-0.347	-0.464	-0.835*	-1.305**	-1.051**	-1.263	0.392***
	(0.409)	(0.417)	(0.408)	(0.408)	(0.411)	(0.451)	(0.389)	(1.022)	(0.0548)
Age × Age	-0.000287	0.000314	-0.000818	0.000190	0.000587	0.00340	0.000328	-0.000378	-0.00193**
	(0.00317)	(0.00325)	(0.00308)	(0.00309)	(0.00318)	(0.00356)	(0.00279)	(0.00818)	(0.000673)
Imports								-0.139	
								(0.0773)	
Dummy=1 if person identifies as female									0.420
									(0.216)
Constant	81.14***	80.49***	86.36***	87.13***	83.87***	88.14***	91.61***	87.08**	33.27***
	(13.30)	(14.29)	(14.31)	(14.37)	(13.43)	(14.28)	(12.85)	(28.61)	(1.662)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	221050	215784	221050	221050	218065	189046	219758	34586	221050

Note: Probability to report to support the incumbent in percentage point. Until May 2010, Labour is coded as the incumbent whereas the Conservatives after 2010. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.4.5 Panel Attrition

Attrition is a key concern in panel data analysis. In our case, one may worry that digitalization causes differential attrition rates between winners and losers. For instance, workers displaced by digitalization can be more likely to move and become more difficult to be located for reinterview. In addition, as discussed above, displacement may force workers to change industries. Higher attrition rates and more industry switches would both make it difficult for us to capture the adverse effects of digitalization, painting an exceedingly optimistic picture.

To examine if digitalization in an industry predicts sample attrition and industry switches, table B.9 first presents the results of regressing the likelihood of dropping out of the sample or changing industries on ICT capital per worker. Next, we examine if these effects are heterogeneous for workers with different education levels by regressing both outcomes on the education dummies and the interaction of ICT capital per worker and education.

The results are reassuring as we do not find clear evidence that ICT capital per worker is associated with increased attrition. While the average effect of our key measure of digitalization is in fact negative, suggesting that workers in rapidly digitalizing industries are less likely to drop out of the panel, this difference is very small. Second, digitalization is not clearly associated with a stronger likelihood to change to a different industry in the next period for none of the education groups. In sum, differences between groups are small. It thus seems unlikely that differential attrition is driving our main results.

Table B.9: Predictors of attrition

	Leave sample		Change industry	
	(1)	(2)	(3)	(4)
ICT	-0.000605**		0.000164	
	(0.000208)		(0.000207)	
Degree × ICT		-0.00130		0.000943
		(0.00193)		(0.00181)
Other higher degree × ICT		0.00241		-0.000600
		(0.00219)		(0.00205)
A-Level etc × ICT		0.00182		0.000965
		(0.00174)		(0.00159)
GCSE etc × ICT		0.00344		0.00158
		(0.00184)		(0.00147)
Other Qualification × ICT		0.00134		-0.000813
		(0.00314)		(0.00345)
No Qualification × ICT		0.00650		0.00794
		(0.00388)		(0.00411)
Degree		0.0994***		0.0524*
		(0.0224)		(0.0248)
Other higher degree		0.0986***		0.0308
		(0.0237)		(0.0245)
A-Level etc		0.0643**		0.000180
		(0.0197)		(0.0216)
GCSE etc		0.0437*		0.00100
		(0.0200)		(0.0211)
Other Qualification		0.0393*		0.0172
		(0.0179)		(0.0197)
Age		0.0395***		-0.0225***
		(0.00372)		(0.00290)
Age × Age		-0.000163***		0.000143***
		(0.0000182)		(0.0000182)
Constant	0.0833***	-1.077***	0.275***	0.639***
	(0.00632)	(0.113)	(0.00813)	(0.0954)
Id*Ind FE	X	X	X	X
Year FE	X	X	X	X
Region	X	X	X	X
Observations	234662	234662	200579	200579

Note: Column (1) reports the direct effect of ICT intensity on probably to leave the sample. Column (2) reports the effect of ICT intensity on the probability to leave the sample by education group. Column (3) reports the direct effect of ICT on the probably to change industries. Column (4) reports the effect of ICT on the probably to change industries by education group. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.4.6 Alternative Clustering

Table B.10 shows that our results are robust when we cluster standard errors at the industry-year level rather than the individual level. This table shows that when clustering at the industry-year level, standard errors tend to be somewhat smaller than in the results presented in the main text.

Table B.10: All Outcomes with Standard Errors Clustered at the Industry-Year Level

	(1) Hourly wage	(2) Unemployed	(3) Turnout	(4) Conservative	(5) Labour	(6) Incumbent
Degree × ICT	0.343***	0.0129	0.635*	0.589***	-0.203	1.527***
	(0.0359)	(0.0875)	(0.247)	(0.173)	(0.166)	(0.365)
Other higher degree × ICT	0.184***	0.00620	0.305	0.540**	-0.124	1.245***
	(0.0299)	(0.0777)	(0.333)	(0.199)	(0.203)	(0.330)
A-Level etc × ICT	0.0514**	0.168*	0.691*	0.580***	-0.229	1.333***
	(0.0172)	(0.0690)	(0.272)	(0.169)	(0.163)	(0.341)
GCSE etc × ICT	-0.0114	0.183*	0.211	-0.0288	-0.206	0.657*
	(0.0162)	(0.0827)	(0.295)	(0.188)	(0.173)	(0.278)
Other Qualification × ICT	-0.135***	0.0451	-0.951	-0.358	-0.473	-0.251
	(0.0262)	(0.0989)	(0.603)	(0.247)	(0.342)	(0.415)
No Qualification × ICT	-0.185***	0.227*	0.148	-0.601**	0.402	-0.207
	(0.0372)	(0.0998)	(0.497)	(0.227)	(0.316)	(0.474)
Degree	-1.995***	0.883	-0.617	-7.420***	2.319	-12.11***
	(0.247)	(0.846)	(3.446)	(1.861)	(2.173)	(2.964)
Other higher degree	-2.028***	1.446	-2.424	-5.326**	0.522	-9.677**
	(0.216)	(0.822)	(3.864)	(2.022)	(2.283)	(3.334)
A-Level etc	-1.628***	0.607	-5.519	-6.227***	0.879	-9.460***
	(0.145)	(0.720)	(3.017)	(1.691)	(2.033)	(2.546)
GCSE etc	-1.141***	0.773	-4.484	-3.577*	1.581	-9.527***
	(0.116)	(0.670)	(2.959)	(1.736)	(2.071)	(2.767)
Other Qualification	-0.441***	1.124	0.548	-0.00495	-0.495	-1.458
	(0.110)	(0.716)	(2.521)	(1.469)	(1.710)	(2.186)
Age	0.345***	-0.435***	-1.143**	0.383	0.128	-0.730*
	(0.0262)	(0.102)	(0.408)	(0.219)	(0.282)	(0.362)
Age × Age	-0.00312***	0.00158*	-0.00913**	-0.00330*	-0.00453**	-0.000287
	(0.000195)	(0.000699)	(0.00281)	(0.00155)	(0.00171)	(0.00223)
Constant	-2.821***	13.76***	133.1***	11.99	59.78***	81.14***
	(0.787)	(3.594)	(13.59)	(6.961)	(8.775)	(11.53)
Individual*Industry FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
Region FE	X	X	X	X	X	X
Observations	179477	216130	103739	221050	221050	221050

Note: All columns use the main specification. Column (1) reports the results for hourly wage, column (2) for the probability to become unemployed, column (3) for voter turnout, column (4) for vote for the Conservatives, column (5) for vote for Labour and column (6) for vote for the incumbent. Except for the the wage variable, all results in percentage points. Standard error reported in parenthesis are clustered at the industry-year level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.4.7 Excluding Migrants

Last but not least, we dealt with the concern that migrants affected our results in a systematic way as they might have a different reaction to digitalization when it comes to political preferences. For example, workers with a migration background might be less inclined to turn to the UK Independence Party if they feel left behind by workplace digitalization.

For this reason, we replicate the analyses excluding workers who were born outside of the UK. This reduces the sample size by about 5%. Table B.11 shows the results for our main outcomes and the support for UKIP. They are almost indistinguishable from the presented results in the main body of the text.

Table B.11: All Outcomes Excluding Foreign-Born Workers

	(1) Hourly wage	(2) Unemployed	(3) Turnout	(4) Conservative	(5) Labour	(6) Incumbent	(7) UKIP
Degree × ICT	0.342***	0.0135	0.624*	0.585**	-0.206	1.535***	-0.428
	(0.0327)	(0.0723)	(0.282)	(0.197)	(0.216)	(0.339)	(0.345)
Other higher degree × ICT	0.185***	-0.000835	0.309	0.524*	-0.116	1.245*	-0.251
	(0.0338)	(0.0661)	(0.367)	(0.241)	(0.238)	(0.517)	(1.026)
A-Level etc × ICT	0.0497*	0.167**	0.687**	0.561**	-0.244	1.305***	-0.924*
	(0.0230)	(0.0624)	(0.264)	(0.193)	(0.192)	(0.356)	(0.469)
GCSE etc × ICT	-0.0123	0.178*	0.215	-0.0175	-0.232	0.679*	0.170
	(0.0186)	(0.0697)	(0.231)	(0.192)	(0.189)	(0.298)	(0.670)
Other Qualification × ICT	-0.138***	0.0357	-0.932	-0.307	-0.531	-0.275	-1.525
	(0.0289)	(0.0820)	(0.578)	(0.264)	(0.350)	(0.542)	(1.199)
No Qualification × ICT	-0.186***	0.237*	0.152	-0.549*	0.340	-0.140	2.845*
	(0.0400)	(0.106)	(0.470)	(0.276)	(0.392)	(0.569)	(1.431)
Degree	-1.998***	0.811	-0.531	-7.260***	2.072	-12.31***	13.22
	(0.212)	(0.820)	(3.345)	(1.959)	(2.397)	(3.644)	(7.029)
Other higher degree	-2.071***	1.479	-2.394	-5.299*	0.416	-9.937*	9.767
	(0.221)	(0.804)	(4.063)	(2.071)	(2.470)	(4.042)	(7.266)
A-Level etc	-1.637***	0.562	-5.441	-6.167***	0.565	-9.644**	9.401
	(0.157)	(0.714)	(2.852)	(1.801)	(2.189)	(3.176)	(6.711)
GCSE etc	-1.139***	0.713	-4.447	-3.761*	1.487	-9.933**	8.603
	(0.148)	(0.680)	(2.891)	(1.767)	(2.047)	(3.195)	(7.128)
Other Qualification	-0.435**	1.049	0.553	-0.363	-0.726	-1.750	17.05*
	(0.139)	(0.671)	(2.286)	(1.725)	(1.849)	(2.644)	(7.698)
Age	0.352***	-0.419***	-1.146**	0.408	0.0934	-0.826*	0.292
	(0.0273)	(0.0996)	(0.391)	(0.229)	(0.271)	(0.414)	(0.580)
Age × Age	-0.00312***	0.00153*	-0.00920***	-0.00323*	-0.00403*	0.000408	0.00596
	(0.000214)	(0.000607)	(0.00264)	(0.00164)	(0.00183)	(0.00320)	(0.00481)
Constant	-3.036***	13.43***	133.3***	11.43	60.36***	84.45***	-20.42
	(0.800)	(3.718)	(12.50)	(7.820)	(9.198)	(13.58)	(21.48)
Individual*Industry FE	X	X	X	X	X	X	X
Year FE	X	X	X	X	X	X	X
Region FE	X	X	X	X	X	X	X
Observations	174697	210773	103358	215730	215730	215730	53893

Note: All columns use the main specification. Column (1) reports the results for hourly wage, column (2) for the probability to become unemployed, column (3) for voter turnout, column (4) for vote for the Conservatives, column (5) for vote for Labour, column (6) for vote for the incumbent and column (7) for vote for UKIP. Except for the the wage variable, all results in percentage points. Standard error reported in parenthesis are clustered at the industry-year level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.5 Other political outcomes

The following tables report the full regression results of additional analyses examining if digitalization affects support for the Liberal Democratic Party and UKIP.

We do not find a change in the support for the Liberal Democratic Party among workers who experience digitalization. The Liberal Democratic Party is a centrist party that includes both classical economic liberals as well as social-democrats. The two main wings have varying strengths across constituencies and over time. One possible interpretation of this finding is that these different factions within the party cancel each other out. It is furthermore noteworthy that it seems that Libdem could not capitalize from an incumbency advantage.

As already graphically presented in the main text, we find some tentative evidence for increased UKIP support among the lowest qualified respondents in our sample, which would be consistent with the possibility that digitalization makes losers more likely to support anti-establishment parties, in this case from the radical right. Among workers with no formal qualification, an increase in ICT intensity produces a substantively large increase in the likelihood to support UKIP. However, the point estimates are never significant. These results have to be interpreted with caution since they are based on a short period of time and small sample. The option to report support for the UKIP is only provided since 2013 and the no qualification group only constitutes 4% of responses in those years.

Table B.12: Support for the Liberal Democratic Party

	(1) Main	(2) IV	(3) Placebo	(4) Region*Year FE	(5) Excl. outliers	(6) Lead	(7) ID FE	(8) Trade	(9) Cross Sect
Degree × ICT	-0.0846	-1.550**	-0.00298	-0.0825	-0.138	-0.0309	-0.0741	-0.676	0.146**
	(0.145)	(0.596)	(0.00212)	(0.145)	(0.226)	(0.174)	(0.0788)	(0.428)	(0.0550)
Other higher degree × ICT	-0.0506	-0.979	-0.00350	-0.0346	0.0478	-0.159	-0.134	-0.146	0.0763
	(0.207)	(0.662)	(0.00241)	(0.205)	(0.274)	(0.227)	(0.0926)	(0.422)	(0.0575)
A-Level etc × ICT	0.184	0.218	-0.00141	0.216	0.188	0.327*	-0.0794	0.222	0.181***
	(0.129)	(0.685)	(0.00195)	(0.130)	(0.186)	(0.136)	(0.0879)	(0.328)	(0.0545)
GCSE etc × ICT	0.0690	-0.836	-0.00363	0.0862	0.112	0.125	0.0420	-0.119	0.273***
	(0.133)	(0.441)	(0.00227)	(0.135)	(0.202)	(0.128)	(0.0840)	(0.523)	(0.0592)
Other Qualification × ICT	0.220	-0.540	0.00253	0.247	0.267	0.333	-0.0890	0.0991	0.181*
	(0.191)	(0.599)	(0.00428)	(0.189)	(0.239)	(0.254)	(0.150)	(0.382)	(0.0725)
No Qualification × ICT	0.259	0.158	0.00190	0.192	0.341	-0.00742	-0.00486	0.217	0.0705
	(0.244)	(0.826)	(0.00245)	(0.239)	(0.327)	(0.324)	(0.118)	(0.291)	(0.0754)
Degree	3.384*	6.710**	3.066*	3.401*	3.839*	1.602	3.363**	5.797	9.510***
	(1.476)	(2.386)	(1.286)	(1.472)	(1.570)	(1.609)	(1.268)	(3.418)	(0.335)
Other higher degree	3.034	4.961*	2.915*	2.989	3.091	1.142	3.007*	2.146	4.842***
	(1.615)	(2.523)	(1.409)	(1.609)	(1.680)	(1.776)	(1.399)	(4.166)	(0.351)
A-Level etc	2.452	1.490	2.767*	2.323	2.797*	0.793	3.700***	1.433	3.561***
	(1.255)	(2.136)	(1.096)	(1.246)	(1.320)	(1.356)	(1.113)	(2.367)	(0.306)
GCSE etc	1.272	3.057	1.445	1.140	1.581	0.725	1.955	0.106	1.546***
	(1.167)	(2.015)	(1.009)	(1.159)	(1.241)	(1.283)	(1.046)	(2.330)	(0.300)
Other Qualification	0.980	2.181	0.582	0.543	0.906	0.00199	1.533	3.396	0.716*
	(1.085)	(1.951)	(0.921)	(1.077)	(1.128)	(1.369)	(0.951)	(2.551)	(0.359)
Age	0.114	0.107	-0.0644	-0.0617	0.130	0.213	-0.0234	0.228	-0.244***
	(0.203)	(0.208)	(0.205)	(0.205)	(0.205)	(0.213)	(0.189)	(0.465)	(0.0349)
Age × Age	0.000904	0.00110	0.00116	0.00116	0.000838	0.000359	0.00159	-0.00381	0.00369***
	(0.00137)	(0.00142)	(0.00136)	(0.00137)	(0.00138)	(0.00151)	(0.00121)	(0.00300)	(0.000428)
Imports								0.00888	
								(0.0234)	
Dummy=1 if person identifies as female									0.928***
									(0.142)
Constant	-1.016	-1.146	3.405	3.125	-1.743	4.134	4.545	-12.04	7.774***
	(6.538)	(6.870)	(6.739)	(6.778)	(6.591)	(6.900)	(5.931)	(15.98)	(1.007)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	221050	215784	221050	221050	218065	189046	219758	34586	221050

Note: Probability to report to support the Liberal Democratic Party in percentage point. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table B.13: Support for UKIP (only asked since 2013)

	(1) Main	(2) IV	(3) Placebo	(4) Region*Year FE	(5) Excl. outliers	(6) Lead	(7) ID FE	(8) Trade	(9) Cross Sect
Degree × ICT	-0.426	-1.374	0.0128	-0.344	-0.198	0.220	-0.264	-0.947	0.0917
	(0.344)	(1.494)	(0.0148)	(0.343)	(0.543)	(0.432)	(0.344)	(0.851)	(0.341)
Other higher degree × ICT	-0.249	-1.043	0.00984	-0.214	-0.617	-1.020	-0.0701	-0.578	0.176
	(1.026)	(2.196)	(0.0151)	(1.016)	(0.725)	(0.589)	(0.468)	(2.090)	(0.345)
A-Level etc × ICT	-0.922*	-1.926	0.00530	-0.847	-0.731	0.248	-0.389	-3.499	0.134
	(0.469)	(1.802)	(0.0137)	(0.468)	(0.581)	(0.367)	(0.348)	(2.238)	(0.345)
GCSE etc × ICT	0.173	-0.219	-0.00877	0.132	-0.224	0.691	-0.353	2.968	0.250
	(0.670)	(1.815)	(0.0175)	(0.668)	(0.735)	(0.402)	(0.399)	(2.040)	(0.350)
Other Qualification × ICT	-1.525	0.426	-0.0324	-1.539	-1.688	-1.835	-0.346	6.022	0.252
	(1.197)	(2.522)	(0.0262)	(1.198)	(1.301)	(1.054)	(0.509)	(3.761)	(0.362)
No Qualification × ICT	2.849*	6.805*	0.0845	2.947*	2.763	1.005	0.0681	20.66*	0.316
	(1.430)	(3.401)	(0.0496)	(1.430)	(1.478)	(0.717)	(0.625)	(10.00)	(0.373)
Degree	13.24	27.79*	8.957	13.02	11.98	6.343	0.885	105.0*	-2.749***
	(7.025)	(10.96)	(6.626)	(6.948)	(7.106)	(4.138)	(4.930)	(52.92)	(0.691)
Other higher degree	9.769	23.44*	6.560	9.910	10.08	11.54**	-1.048	87.51	-0.681
	(7.262)	(10.74)	(6.545)	(7.180)	(7.077)	(4.337)	(5.020)	(51.36)	(0.733)
A-Level etc	9.393	23.38*	4.075	9.245	8.727	8.134*	-1.275	89.48	0.506
	(6.707)	(10.49)	(6.269)	(6.631)	(6.766)	(3.874)	(4.722)	(50.07)	(0.719)
GCSE etc	8.586	20.63*	7.865	9.177	9.380	3.007	1.156	71.45	1.327
	(7.124)	(10.45)	(6.815)	(7.050)	(7.187)	(3.934)	(5.090)	(49.47)	(0.731)
Other Qualification	17.01*	23.77*	14.43*	17.48*	17.18*	8.085	4.201	67.47	1.652
	(7.686)	(10.61)	(7.017)	(7.574)	(7.801)	(4.869)	(5.616)	(47.77)	(0.881)
Age	0.273	0.283	0.239	0.255	0.198	-0.476	-0.149	3.394	-0.0123
	(0.578)	(0.582)	(0.581)	(0.579)	(0.582)	(0.478)	(0.558)	(2.016)	(0.0529)
Age × Age	0.00614	0.00634	0.00628	0.00611	0.00736	0.00766*	0.00845	-0.0187	0.000924
	(0.00480)	(0.00485)	(0.00482)	(0.00481)	(0.00480)	(0.00377)	(0.00454)	(0.0164)	(0.000655)
Imports								-0.0259	
								(0.0664)	
Dummy=1 if person identifies as female									-1.410***
									(0.192)
Constant	-19.97	-30.30	-18.27	-20.05	-18.70	-8.824	-2.980	-163.8	2.922
	(21.42)	(24.64)	(21.26)	(21.39)	(21.69)	(17.73)	(20.57)	(84.88)	(1.707)
Individual*Industry FE	X	X	X	X	X	X		X	
Year FE	X	X	X		X	X	X	X	X
Region FE	X	X	X		X	X	X	X	X
Year*Region FE				X					
Individual FE							X		
Industry FE							X		X
Observations	54137	52995	54137	54137	53495	60141	53992	7103	54137

Note: Probability to report to support the United Kingdom Independence Party in percentage point. Column (1) is our main specification with industry-spell fixed-effects. In column (2), we instrument ICT with data from the USA. Column (3) uses non-ICT capital per worker as main regressor. Column (4) is equivalent to the main specification with adding region by year fixed-effects. In column (5) we exclude the most digitalized industries. Column (6) uses the lead of the dependent variable. Column (7) uses individual fixed-effects and industry fixed effects. Column (8) includes a control for trade. Column (9) is a cross sectional analysis without individual fixed effects. Standard error reported in parenthesis are clustered at the individual level. * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.6 Mechanisms

B.6.1 Operationalization and Data Availability

The three dependent variables of the mechanism section are operationalized as follows:

- Satisfaction with Life: Likert scale of:
  - ”Satisfaction with Life overall” (lfsato, sclfsato), 1=completely dissatisfied, 7=completely satisfied. Linearly imputed within individual if missing between two non-missing values.
- Supports Government Intervention: Principal component analysis (PCA) of:
  - ”Private enterprise solves economic probs” (opsocc), 1=strongly agree, 5=strongly disagree. Linearly imputed within individual if missing between two non-missing values.
  - ”Government has obligation to provide jobs” (opsoce), 1=strongly disagree, 5=strongly agree (recoded). Linearly imputed within individual if missing between two non-missing values.
- Social Progressiveness: Principal component analysis (PCA) of:
  - ”Pre-school child suffers if mother works” (scopfama), 1=strongly agree, 5=strongly disagree. Linearly imputed within individual if missing between two non-missing values.
  - ”Family suffers if mother works full-time” (scopfamb), 1=strongly agree, 5=strongly disagree. Linearly imputed within individual if missing between two non-missing values.
  - ”Husband and wife should contribute to hh income” (scopfamd), 1=strongly disagree, 5=strongly agree (recoded). Linearly imputed within individual if missing between two non-missing values.

The underlying survey items are only included infrequently in BHPS/UKHLS. Table B.14 provides an overview of their availability. Talbe B.15 gives basic descriptive statistics.

Table B.14: Availability of Survey Items over Time (N obs)

Year	Satisfaction	Gov Intervention	Progressiveness
1997	5896	5847	5835
1998	5859	5057	104
1999	6206	4574	5972
2000	7246	5715	1821
2001	7705	6960	7385
2002	7781	5750	1440
2003	8908	5957	7652
2004	8298	7807	355
2005	8495	6738	7680
2006	8163	6477	207
2007	7935	7196	7184
2008	7663	273	233
2009	11425	0	0
2010	24302	0	13480
2011	24040	0	8665
2012	22388	0	12626
2013	21525	0	7993
2014	20407	0	556
2015	18814	0	0
2016	19262	0	0
2017	8213	0	0
2018	909	0	0

Table B.15: Mechanism Items: Descriptives

	count	mean	sd	min	max
Satisfaction	261’440	5.2	1.285	1	7
Government Intervention	68’351	0	1.081	-3.356	3.153
Progressiveness	89’188	0	1.323	-3.491	2.713

B.6.2 Results

Figure 3.7 presents the results of the analyses about mechanisms, which are discussed in the main text.

Figure B.7: Effect of digitalization on satisfaction and attitudes

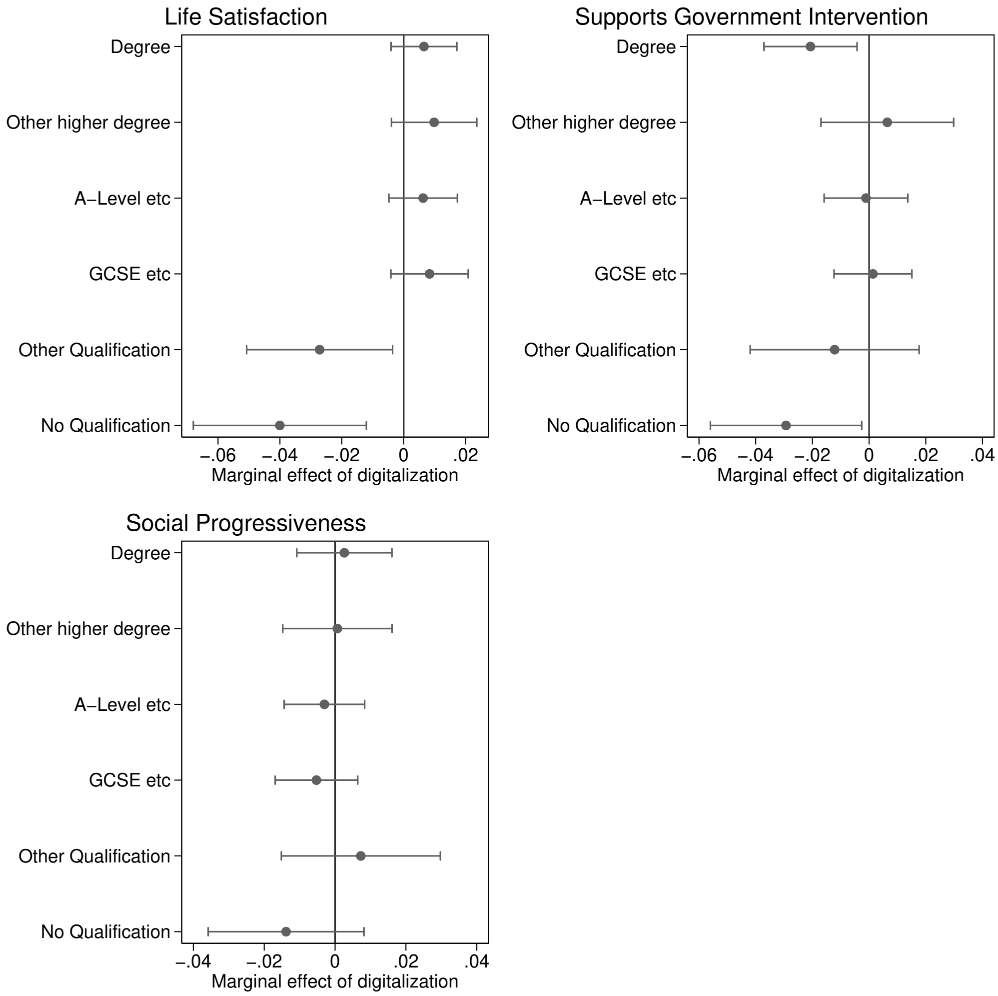

Note: Results show marginal effect of one unit increase in digitalization (1000 GBP in ICT capital/worker) on specified dependent variable, industry-spell fixed-effects specification.

B.6.3 Additional description of the UK political context

Positions of the parties over time

We use Chapell Hill Expert Survey to back the claim in the main text that the Labor Party has been more pro-redistribution throughout the time period studied.

Figure B.8: Position on Redistributive Issues

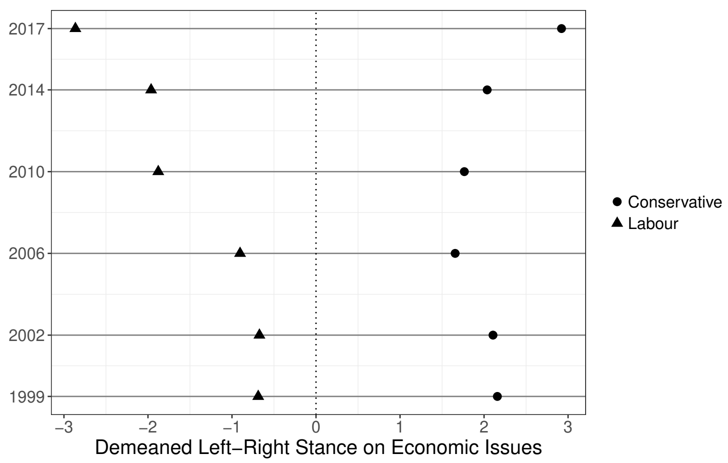

Source: Chapel Hill Expert Survey. Values of economic left-right position (lrecon) demeaned by year across all available party positions. Party positions weighted by vote share.

B.6.4 Party Manifestos

In order to get a more precise idea of potential supply-side effects related to the framing of the digitalization debate, we undertook an original analysis of the two large parties' most recent manifestos. We studied the content of the Conservative and Unionist Party Manifesto 2017 (”FORWARD, TOGETHER. Our Plan for a Stronger Britain and a Prosperous Future”, 88 pages, available online [access date: November 22, 2019]) and the Labour Party Manifesto 2019 (”It's time for real change”, 107 pages, available online [access date: November 22, 2019]). The Conservative 2019 Manifesto was not yet available at the time of writing. If anything, we would expect the less recent manifesto to result in a downward bias of attention to digitalization compared to the Labour Party.

We examine if the two parties differed in the extent to which they discuss digitalization and technology in their manifestos. A simple key word analysis demonstrates that the Conservative Party speaks more about these issues than the

Labour party. In general, attention to the topic is surprisingly limited in both manifestos, which might reflect the difficulty to claim ownership of a newly emerging issue (König and Wenzelburger, 2018). Still, while apparently not being a priority, the relevant concepts at least appear among the Conservative's top-30 terms. This is not the case for the Labour manifesto, which has been released very recently. Figure B.9 gives a broad overview and provides a comparison between the two parties.

Figure B.9: Digitalization: ICT capital stock per employee, by industry

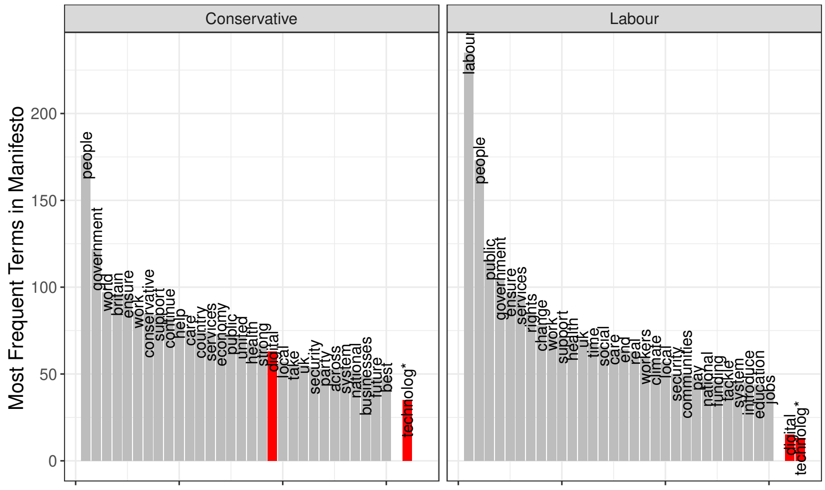

We next looked at the relevant keywords in context to get a better sense of the way the Conservative Party tried to frame the debate. A simple overview in Table B.16 suggests that they address the issue in an almost exclusively positive sense, in which digitalization benefits businesses and the economy in general. Digital technology, according to the Conservative Party, promises prosperity and security. Another frequent feature is the use of new technology to increase government efficiency and public services, e.g. related to NHS. A final important aspect is investment in skills to seize the opportunities provided by new technologies.

To summarize, it can be said (a) that digitalization has not featured very prominently in the two main parties' manifesto in absolute terms, (b) that the Conservative Party was considerably more attentive to the issue in relative terms, and (c) that it discussed almost exclusively the beneficial aspects of new technologies. We conclude that our simple supply-side analysis supports the idea that the Conservative Party is a reasonable political choice for ordinary winners of digitalization throughout the whole period.

Table B.16: Conservative Manifesto: Top Features among Keyword ('Digital') in Context

top features	count
technology	10.0
economy	9.0
services	8.0
digital	8.0
age	8.0
prosperity	7.0
security	6.0
government	6.0
help	6.0
use	6.0
charter	6.0
new	5.0
companies	5.0
businesses	5.0
infrastructure	5.0
right	4.0
skills	4.0
public	4.0
creative	3.0
data	3.0
strategy	3.0
ensure	3.0
provide	3.0
online	3.0
support	3.0
access	3.0
also	3.0
need	3.0
people	2.0
working	2.0

3

HOW TECHNOLOGICAL CHANGE AFFECTS REGIONAL ELECTORATES

Joint with Thomas Kurer (Universität Zürich)

3.1 Introduction

The widespread use of new technology at the workplace has raised fears about wage pressure and employment loss. Influential work in labor economics shows that capital in the form of industrial robots or specialized software directly replaces certain routine tasks previously done by human labor in both white- and blue-collar occupations (Autor et al., 2003; Acemoglu and Restrepo, 2019). This has sparked a vivid debate about the political and societal consequences of such an uncertain future of work as those who lose out in this process are likely to seek for ways to express their discontent. Indeed, a growing literature in political science has gathered mounting evidence suggesting that workers directly threatened by a transforming employment structure disproportionately support anti-establishment parties (Frey et al., 2018; Im et al., 2019; Anelli et al., 2019, 2021; Kurer, 2020; Milner, 2021).

This paper explicitly recognizes that technological innovation affects regional voting outcomes in two ways. On the one hand, there is a *direct effect* on workers who are threatened by technology and arguably become more supportive of radical right and populist forces. On the other hand, technological innovation also affects regional voting through a *compositional effect*. Over time, more and more workers belong to occupations which are associated with more progressive values. The direction of the net effect of technological innovation on regional voting outcomes is thus theoretically ambiguous. We advance the existing literature by an empirical analysis of the relative importance of the direct and compositional effect in West Germany. This case is highly relevant because (a) West Germany is both one of the largest information and communication technology (ICT) markets in the world and home to the overwhelming majority of industrial robots currently installed in Europe, (b) West Germany has still the largest manufacturing share of employment compared to other advanced economies and (c) has recently seen the rapid rise of a radical right party, thus putting an end to a historic taboo.

Fine-grained labor market data with high levels of geographical disaggregation from the German Institute for Employment Research (IAB) allow for a more detailed regional analysis than most existing accounts. We combine these detailed labor market data with two distinct empirical measures of technological change. First, we use data from the International Federation of Robotics (IFR) to measure county-level exposure to robotization and how it has changed over time. This mainly captures automation in the manufacturing sector. Second, we measure county-level exposure to digitalization in the form of ICT by relying on EU-KLEMS data (Jaeger, 2016). This constitutes a distinct form of technological change which (in contrast to robotization) also affects the service sector. Following pioneering work in the field (Acemoglu and Restrepo, 2020), identification stems from a shift-share approach where we use pre-sample-period local employment composition to estimate the exposure to new technologies in a time-varying fashion. We employ a panel model with region and time fixed effects (generalized diff-in-diff) to control for unobserved factors.

Unlike most existing work studying the political implications of the most recent wave of technological change, our approach allows to document technology-induced changes in the labor market that are typically invoked to explain political reactions. This is important as all studies on the topic –more or less explicitly– argue that technological change affects political outcomes through the labor market. In line with previous work in labor economics, our approach reveals that robot adoption and ICT investment shift employment from manufacturing and routine jobs to the service sector. Furthermore, regions with faster growing technological innovation experience stronger labor market polarization: Semi-skilled and routine occupations decline at the expense of non-routine work at both ends of the skill spectrum. Robots primarily displace manual routine jobs whereas ICT investment more powerfully substitutes for cognitive routine jobs. However, importantly, overall employment does not decrease in West German counties with higher exposure to technological change. To the contrary, we find weakly positive net employment effects.[1]

[1]This finding helps correct a common misperception. Investment in new technologies is

Our analysis of political outcomes at the regional level does not support the narrative that new technology at the workplace first and foremost results in right-authoritarian success. Instead, we find that, on average, regions more strongly affected by technological innovation shift their political support towards socially progressive parties. The *regional* vote shares of center-right and right-authoritarian parties *decline* as a result of the labor market transitions caused by robot adoption and ICT investment. We provide evidence that these results are indeed the consequence of changing local labor market composition. In line with the literature on occupational preference formation, we demonstrate that a lower number of regional manufacturing jobs is associated with less support for right-authoritarian parties whereas a larger interpersonal service sector is associated with more support for progressive left parties.

By highlighting that new technologies not only replace human work (the replacement effect) but also create new jobs (the productivity effect), we challenge rather gloomy perspectives on the political repercussions of technological change. In the conclusion of the paper, we provide an extensive discussion of the relative importance of the direct individual-level effect of automation risk that could create support for radical-right parties vis-à-vis the aggregate-level effects of economic modernization that may work in the opposite direction by changing the composition of local labor markets away from manufacturing towards more high-skilled non-routine jobs. Concerning the important case of West Germany, we show that compositional effects of technology adoption on local labor markets can outweigh the political resentment among workers directly affected by the adverse consequences of technological change. Hence, our results suggest that technological innovation need not result in local political disruption. While we acknowledge that automation contributes to the emergence of anti-establishment forces through electoral support from the segment of society directly exposed to the negative consequences of this process, our results show that, overall, technology adopting regions do not necessarily turn into right-authoritarian strongholds.

actually a sign of a relatively healthy, future-oriented local economy. While it could be imagined that the alternative to robot adoption were thriving manufacturing plants relying on human work, recent research suggests that the more realistic counterfactual scenario seems to be substantial job loss and closed factories as companies without robots fall behind in global competition (Koch et al., 2019).

3.2 Labor Market Implications of Technological Change

In their seminal work on routine biased technological change (RBTC), Autor et al. (2003) argue that new technologies primarily substitute for routine tasks that follow clearly defined rules, which makes jobs that heavily rely on such tasks ”codifiable” and hence replaceable by computers or robots. This *substitution effect* mainly hits workers located at the middle of the income and skill distribution and in particular workers in the manufacturing sector.

On the other hand, new technologies raise productivity which leads to an increased demand for workers whose skills are complementary to automation. Newly created jobs tend to pertain either to the growing group of white-collar professionals with college education focusing on cognitive and interpersonal tasks (management, education and cultural and health sector) or to a rather precarious group of low-skilled manual services (retail, restaurants and hospitality). Most of them benefit from automation indirectly through lower prices of goods and new demands for their products and services. This was dubbed the *productivity* or *reinstatement effect* (Acemoglu and Restrepo, 2019).

While there seems to be a general consensus among scholars that these are the main forces at work, it is still hotly debated whether the substitution or productivity effect dominates. With respect to robotization, an influential paper on the US found that the substitution effect dominates as regions adopting more robots experienced weaker employment growth (Acemoglu and Restrepo, 2020). On the other hand, studies focusing on Europe and on Germany in particular found null or slightly positive employment effects (Dauth et al., 2021; Klenert et al., 2020). With regard to ICT, existing work appears slightly less controversial and tends to show that investment in technology has not led to a decline in employment (Biagi and Falk, 2017) but shifted jobs from mid-skill to high-skilled sectors, consistent with ICT-based employment polarization (Michaels et al., 2014).[2]

Our own original analysis points in the same direction: although we do find that mid-skilled routine jobs generally and manufacturing employment in particular are negatively affected by technological innovation, this decline is more than offset by an increase in work in other sectors. While no single analysis will be able conclusively answer the question of whether technology adoption tilts the balance towards more or less employment, for our purpose the distributive implications of robotization and ICT investment and how they

[2]It should be noted, however, that it remains unclear to what extent the findings on traditional ICT investment can be applied to the most recent and, especially, future developments in the domain of software development and artificial intelligence (Graetz, 2020).

transform the composition of local labor markets are particularly relevant. Parts of society, namely manufacturing and routine workers, stand to lose from this process whereas non-routine occupations requiring cognitive and social skills (and oftentimes a university education) are growing in numbers.

3.3 Political Implications of Technological Innovation

The distributive implications of technological innovation as described in the previous section give rise to two distinct and most likely countervailing political implications. On the one hand, studies that focus on the *direct* effect are interested in the individual-level consequences of direct exposure to automation. On the other hand, a different strand of literature has studied the consequences of economic modernization and occupational change at the aggregate level and emphasizes the changing *composition* of postindustrial societies, i.e. general upskilling and the emergence of modern ”knowledge economies”. These two perspectives have most often been studied in isolation and it should not come as a surprise that they come to fundamentally different conclusions about the prospects for advanced capitalist democracy exposed to automation. While the first is often motivated by a concern about the potential substitution of human labor and resulting political disruption, the second provides a much more optimistic outlook emphasizing economic opportunity and mobility through widespread higher education. Interestingly, the net impact of the two effects remains unclear and the relative importance of winners and losers is at the root of much of the debate about the political implications of technological change.

3.3.1 Direct Effect

Existing papers that according to our grouping study the direct effect of automation risk focus on individual-level effects on political preferences and voting behavior. Despite the fact that technological change creates both winners and losers, it seems safe to say that most existing work investigates the political reactions of workers who stand to lose from technological change. Alluding to the historical examples of Luddites destroying machines during the Industrial Revolution, pundits and academics alike have raised concerns that the left-behind would turn against the system. In short, it is argued that losers of technological become more attracted to anti-establishment forces due to their economic decline (Frey et al., 2018; Anelli et al., 2019; Kurer and Gallego, 2019; Im et al., 2019). Specifically looking at the impact of robots, Frey et al. (2018) showed an association between robot adoption and anti-incumbent voting in the U.S. and Anelli et al. (2019, 2021) and Milner (2021) provide evidence for a link between local robot penetration and support for right-authoritarian parties across Western Europe.

The political reactions of winners of technological change have received considerably less attention in individual-level research. Gallego et al. (2021) examine political preferences of ”ordinary winners” of digitalization in the United Kingdom. They show that a majority of the population, but especially high-skilled workers, benefit from ICT capital investment and that these economic benefits translate into more support for moderate incumbent parties, in particular those from the center right. The intragenerational experience of growing economic prosperity as a consequence of technological innovation hence creates a stabilizing pro-system force. The more specific literature on political consequences of robotization has so far exclusively focused on the downsides of this process. We are not aware of scholarly work analyzing how robotization affected the political attitudes of those who benefited from this process.

Summing up, workers imminently threatened by automation tend to become more supportive of radical parties challenging the political status quo. The direct effect of automation seems to primarily benefit authoritarian-right parties. Voters who benefit at least moderately from the ”digital revolution”, in contrast, tend to vote for more centrist ideological positions and support incumbent parties. Technological change hence potentially creates political divergence between winners and losers and can contribute to increasing political polarization.

3.3.2 Compositional Effect

While research on individuals’ susceptibility to automation has concentrated on the downsides of the technological revolution, its upside is at the heart of a different body of work that describes the transition of modern society into ”knowledge economies”. Starting back in the late 1970s, technological progress has facilitated a transition in advanced capitalist democracies from a manufacturing-based to a more services dominated economy with an ever greater reliance on intellectual capabilities rather than on physical inputs or natural resources (Powell and Snellman, 2004). Influential recent accounts highlight the value of the educational expansion (Boix, 2019) and a broad (upper) middle class enjoying economic growth, wealth and opportunity (Iversen and Soskice, 2019).

The emergence of the knowledge economy is intimately linked to the distributional implications of technological change discussed above. Non-routine and service sector jobs, especially higher skilled ones, have expanded at the expense of mid-skilled routine jobs. Importantly, this change in the composition of local labor markets has important political implications since occupations are known as important sites of preference formation (Kitschelt, 1994; Oesch, 2006; Kitschelt and Rehm, 2014). Occupations shape political preferences through both a market logic reflecting vertical divisions in marketable skills and economic self-interest, and an important additional horizontal differentiation in terms of work logic. Key contributions to the literature differentiate between a technical, organizational/bureaucratic and interpersonal work logic depending on the education level required, setting of the work process, the relation to authority, the primary type of client relation and the kind of skills applied (Oesch, 2006). At the risk of simplification, the theory of occupational preference formation thus posits that lower education levels, strict hierarchies and dealing with objects and files (rather than people) are associated with authoritarian views. Occupations that require university educations, which are based on cooperation (rather than hierarchies), which focus on social interactions and culture tend to entertain more cosmopolitan and progressive values (Kitschelt, 1994).[3] Translating this into actual occupational groups and milieus means that mid-skilled, routine occupations in the manufacturing sector are characterized by disproportionate support for authoritarian-right parties (see, e.g. Oesch, 2008). Much in contrast, the growing number of highly educated workers engaging in more analytical and interactive work ("socio-cultural professionals") tend to belong to a milieu which is more left-leaning and cosmopolitan. Gingrich and Häusermann (2015) show how this transformation of the employment structure has resulted in a decline of traditional class voting: contemporary progressive left parties draw substantial electoral support from among an expanding highly educated middle class.[4]

Going back to the expected distributional implications of automation, it becomes clear that the compositional effect shifts political support to progressive left parties. The manufacturing sector and in particular semi-skilled routine work is shrinking through the substitution effect. At the same time, occupations with higher educational requirements and a more client-interactive work logic are growing due to the productivity effect of automation.[5]

[3] Kitschelt and Rehm (2014) provide an extensive discussion on whether these correlations may result from a selection into occupations ("weak theory") or from socialization within occupations ("strong theory").

[4] Section 3.5.2 provides additional micro-level evidence from household panel data bolstering this conjecture.

[5] Note that the regional labor force composition can change through different mechanisms: (a) workers can retrain and change occupations as demand for their original occupation declines, (b) the region can attract migrant workers from elsewhere if they have the education and skills now required and (c) new generations will choose a different educational path and take up occupations that are now in high demand.

3.3.3 Net Effect

The political space in Germany and many other postindustrial democracies is composed of an economic and a cultural dimension (see, e.g., Kitschelt, 1994). The lion's share of voters as well as the relevant political actors tend to cluster along the diagonal, which is characterized by a progressive, economically left-leaning pole and an authoritarian, economically right-leaning pole with progressive left parties and authoritarian-right parties representing ”polar normative ideals” (Bornschier, 2010). Appendix Figure C.1 provides a descriptive overview of the contemporary German partisan landscape. From a theoretical perspective, the direct and the compositional effect of automation work as opposing forces. While the direct effect of automation risk and substitution may fuel individual support for the authoritarian right, the accompanying shift in the composition of the labor force fuels party support for more progressive, cosmopolitan left parties. Hence, a priori, technological innovation could affect regional party support in either way. We treat the question of which factor dominates as an empirical issue and strive to provide an answer, at least for the German case, in below analysis.

3.4 Data

Our empirical analysis focuses on the important case of West Germany.[6] West Germany is a highly relevant case for various reasons: (i) it still has a large manufacturing sector, (ii) it has deployed the largest number of robots anywhere outside Asia, and (iii) it experienced large investments in ICT over the years (see Figure 3.1). We apply a regional approach similar in spirit to previous studies in economics (Acemoglu and Restrepo, 2020; Dauth et al., 2021), choosing West German counties (*Landkreise und kreisefreie Städte*) to be the regional unit of analysis (n = 324, NUTS-3).[7] In the following, we describe how we constructed a yearly panel of economic and political variables on county level for the time period 1994 to 2017.

3.4.1 Robot exposure

To calculate regional robot exposure over time, we use data from the International Federation of Robotics (IFR). A robot is defined as an "automatically controlled, re-programmable, and multipurpose machine". This means that robots are "fully autonomous machines that do not need a human operator and that can be programmed to perform several manual tasks such as welding, painting, assembling, handling materials, or packaging" (IFR, 2016). The yearly data differentiates between 25 industries, mostly in manufacturing.

To measure robot exposure at a time-varying county level, we follow Acemoglu and Restrepo's (2020) approach to exploit information on pre-sample regional employment composition. The idea is to distribute the robots of a given sector to regions based on the the number of employees in the region working in the sector relative to the nation-wide employment in the sector. Since we are interested in the robot intensity of a region, i.e. the number of robots per workers,

[6]We do not consider regions of the former GDR due to profoundly distinct economic and political trajectories. The structure of the manufacturing sector differs fundamentally, with technology penetration being much slower in the East. In addition, while East German manufacturing has imploded right after the fall of the iron curtain in 1989, resulting in a much lower but eventually slightly growing manufacturing share (19% in 2017, up from 17% in 1994), West German manufacturing has seen a more steady decline (25% in 2017, down from 32% in 1994). Finally, with respect to the political arena, existing research shows that support for radical right parties is systematically higher in the East (Lengfeld, 2017). In light of the time period under consideration, we think that a focus on West German counties is reasonable in that it provides a cleaner, more homogeneous sample to study the questions at hand.

[7]In the early nineties there were still almost 400 counties in West Germany which were then merged and reshaped in various rounds of regional reforms. Election results and other control variables are according to those historic definitions of counties. To create a consistent panel based on the current shape of counties, we employ population weights which we obtained from the Federal Statistics Office.

Figure 3.1: Evolution of manufacturing share, robot penetration and ICT

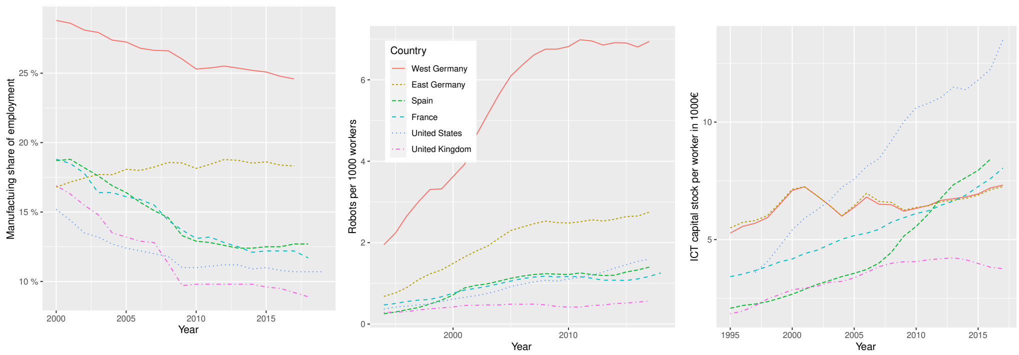

(a) Manufacturing share of employment

(b) Robots

(c) ICT

Note: The graph shows (a) the share of employees working in the manufacturing sector, (b) the number of robots per thousand employees and (c) the ICT capital stock per worker in 1000€. Compared to other advanced economies, West Germany still has a large manufacturing sector while robots are already playing an important role. Digitalization also plays an important role in West Germany. Sources: IFR, ILO, EUKLEMS, own calculations.

we normalize by the region's total employment in thousands. Finally, to account for the heavily skewed distribution of robots across regions, we apply a logarithmic scale. (The robustness section demonstrates that our results do not hinge on this transformation of the explanatory variable.)

$$\text{Robot intensity}_{r,t} = log\left(\frac{1}{E_r}\sum_j \frac{Robots_{j,t} * E_{j,r}}{E_j/1000}\right) \tag{3.1}$$

where E_r is the employment in region r, $E_{j,r}$ is the employment in industry j in region r, $Robots_{j,t}$ is the number of robots in industry j in year t and E_j is the total employment in industry j across all regions.

Information on local employment composition is derived from administrative data of the Institute for Employment Research (IAB).[8] In particular, we use employment records from a 2% sample randomly drawn from the universe of

[8]In constructing this regional measure of robot exposure, we closely follow Dauth et al. (2021). Unlike them, we do not have access to the universe of German employees but only to a (still very large) 2% sample of all employees provided to external researchers by the IAB. To get closer to the universe of employees, we take advantage of the fact that the IAB provides information on number of coworkers for all of the sampled workers. By counting all employees of their respective workplaces we increase the effective sample size drastically. Furthermore, we considered information from all years between 1984 and 1994 to get a clear estimate of regional employment composition. In Section C.6 of the appendix, we confirm that – despite the lack of access to the full employee sample – our proposed approach can successfully replicate the main results of Dauth et al. (2021).

German employees subject to social security (Antoni et al., 2019). For those, we have information on employment status, employer and occupation for any given day for the entire sampling period. An adjacent firm data set includes information on the firm's industry classification, its number of employees and geographic information. We aggregate information on all firms in a 10-year window prior to our sample period by region and industry to approximate local employment composition. Employment data is used from pre-sample period as later sectorial employment composition might be endogenous to the adoption of robots. In addition, IAB data also provides regional employment shares along various dimensions (e.g. by sector, main task or skill requirements). These time-varying, disaggregated employment shares allow us to carefully trace distributional implications on the regional level.

The measure constitutes a typical Bartik-style shift-share variable where an industry-level shock is apportioned across regions (Bartik and Doeringer, 1993).[9]

3.4.2 ICT investment

We use yearly changes in ICT capital stocks within industries to measure digitalization, drawing on the 2019 release of the EU-KLEMS dataset (Stehrer et al., 2019), which contains yearly measures of output, input and productivity for 40 industries in a wide range of countries, including Germany, and covers the period 1995 to 2017. The data is compiled using information from the national statistical offices and then harmonized to ensure comparability. Most importantly for our purposes, the database provides a breakdown of capital into ICT and non-ICT assets (O'Mahony and Timmer, 2009). We define the industry-level ICT capital stock as the capital stock in information technologies, communication technology and software and databases. Based on this, we create a time-varying measure of digitalization using a shift-share approach analog to our robot intensity measure. More specifically, we calculate the ICT capital stock per 1000€ in region r in year t as

$$\mathrm{ICT}_{r,t} = \frac{1}{E_r} \sum_j \frac{\mathrm{ICT}_{j,t} * E_{j,r}}{E_j} \tag{3.2}$$

where E_r is the employment in region r in the base year, $E_{j,r}$ is the employment in industry j in region r in the base year, $ICT_{j,t}$ is the industry ICT capital stock in 1000€ in industry j in year t and E_j is the total employment in industry j across all regions.

[9] Even though popular in the literature, this approach has also received criticism in recent years (for discussions see Goldsmith-Pinkham et al., 2020; Jaeger et al., 2018)

Figure 3.2: Regional distribution of new technologies

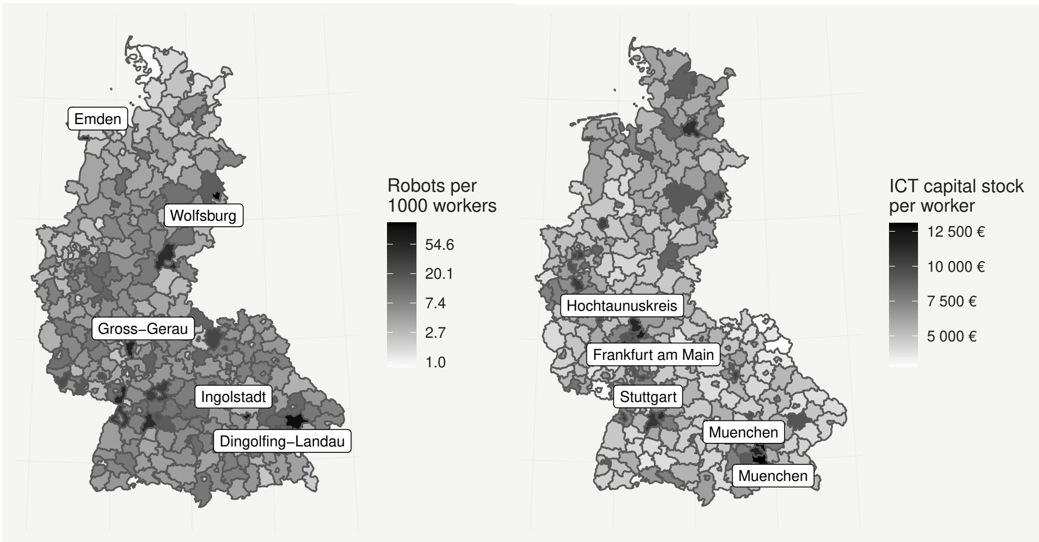

(a) Robot Intensity (b) ICT

Note: The graph shows (a) the estimated number of robots per thousand workers and (b) the ICT capital stock per worker for 324 West-German regions (*Kreise und kreisfreie Städte*) in 2017. Top 5 cities are labeled. Analogous to our measure of robot intensity in the main analysis, the color scale is in logs.

This allows for the creation of time-varying, industry-specific indicators of digitalization based on ICT stocks.

Figure 3.2 shows the spatial distribution of both measures of technological change per county for 2017. The left panel shows that most robots can be found in regions dominated by the automotive industry: For example, Volkswagen has its headquarters in Wolfsburg, Audi in Ingoldstadt, Opel in Gross-Gerau and Dingolfing-Landau and Emden are major production sites of BMW and Volkswagen respectively. Hence, face validity of our measure is high: All regions standing out due to their exceptionally high exposure to robotization can be clearly associated to (car) manufacturing hot-spots. The right panel shows that ICT is concentrated in the major service-sector business hubs of Munich, Frankfurt and Stuttgart. This shows that we capture two distinct forms of technological change. In fact, the correlation between the two measures is low (0.12).

3.4.3 Elections

For each county we gathered official election results for all Federal, State and European elections between 1994 and 2017 which yields 7 federal, 40 state elections and 5 European elections. If multiple election were held in the same year, we only consider one of them, preferring federal election over state election over EU election (order of voter turnout) which gives a total of 4277 county-election pairs. We consider all parties currently represented in national parliament: Grünen (greens), Linke (leftist), SPD (social democrats), FDP (pro market), CDU-CSU (christian democrats) and the Alternative fuer Deutschland (AfD, right-authoritarian). Since the AfD was only founded in 2013, we pool it with other right-authoritarian parties (NPD, DVU, Republikaner).

3.4.4 Empirical Approach

We employ a two-way fixed effect panel model (generalized diff-in-diff) to measure the effect of new technologies, measured as robotization or ICT investment, respectively, on economic and political outcomes:

$$Y_{r,t} = \beta_1 Technology_{r,t} + \mu_t + \eta_r + \epsilon_{r,t} \tag{3.3}$$

The dependent variable $Y_{r,t}$ is a party vote share or an employment outcome in region r in year t which is regressed on $Technology_{r,t}$ measured as (a) the number of log robots per 1000 workers or (b) the ICT capital stock per worker in 1000€. The model also includes region fixed effects η_r and year fixed effects μ_t. As robustness checks, we will further add a vector of control variables in later specifications.

3.5 Results

3.5.1 Political Outcomes

In line with our theoretical point of departure, we first turn our attention to political outcomes and look at ”reduced-form” specifications modelling the direct relationship between regional technological adoption and regional election outcomes. Figure 3.3 plots estimated marginal effect of regional robot intensity (see Panel 3.3a) and ICT investment (see Panel 3.3b) on regional electoral vote shares of all major German parties. The reported coefficients each stem from a separate regression. We first run a specification where only include one of the technological change measures (blue triangles) and secondly a specification including both measures of technological change simultaneously (red circles). Both specifications include a region and an election fixed effect.[10]

[10]See column (1) and (3) of Tables C.1-C.12 in the Appendix. Further note that election fixed effects differ from year fixed effects in the case of multiple state elections held in the same year.

The results show that regions exposed to more intense technology adoption generally shifted their electoral support to the progressive-left of the political spectrum. For ICT, the patterns are consistent and robust. We find that the green party *Die Grünen* and leftist party *Die Linke* were the parties that gained most votes in digitalizing regions. For the social-democratic SPD we find a positive but imprecisely estimated effect. On the other hand, the center-right CDU/CSU and the authoritarian-right party AfD received less support. The estimated effect for the pro-market party FDP is marginally negative. These findings are not affected when controlling for the effect of regional robotization. These reduced form models focusing on ICT investment hence provide evidence that the compositional effect, which favors progressive-left parties, seems to dominate the direct substition effect at the regional level.

For robotization, the overall pattern is similar but more noisy. When considering the effect of robotization in isolation, we find the same gradient across the political spectrum: progressive-left parties gain whereas conservative and authoritarian-right parties tend to receive less votes when a region adopts robots. However, only the effect of the progressive-left party *Die Grünen* is statistically significant. Moreover, when controlling for the parallel influence of ICT, the marginal effects of robotization hover around zero and none of them is significant. We interpret this as evidence that for robotization, the direct effect favoring authoritarian-right parties and the compositional effect favoring progressive left parties are on balance. Nevertheless, we consider this an important result as it contrasts with previous work claiming that robotization leads to an unambiguous shift towards the right of the political spectrum.

In terms of effect magnitude, our baseline models predict that a one standard deviation increase in the log number of robots per thousand workers (+30% more robots) is associated with an increase of the *Grünen* vote share of 0.15 percentage points. Taken as such, we interpret this as modest effects. However, considering that the average region increased its number of robots by 270% between 1994 and 2017, the accumulated effect for *Die Grünen* is an estimated increase of the vote share by 0.71 percentage points, which is important for a party which usually attracted a vote share of less than 10%.

Similarly, an increase of the ICT capital stock by one within-region standard deviation (+520€ per worker) is associated with an increase of the vote for *Die Grünen* by 0.19 percentage points. We resist the temptation to directly compare the magnitude of effects sizes of the two technologies for two reasons. First, we cannot directly compare the absolute change in robot intensity and ICT capital stocks as the former is measured in counts whereas the latter is measured in monetary terms. Secondly, the comparison would implicitly assume that we measure both concepts equally well. However, we have to be clear that both measures are only an approximation of the underlying concept and both suffer from measurement bias to some extent, which attenuates effect sizes.

Figure 3.3: Region-level exposure to technological change and Party Vote Shares

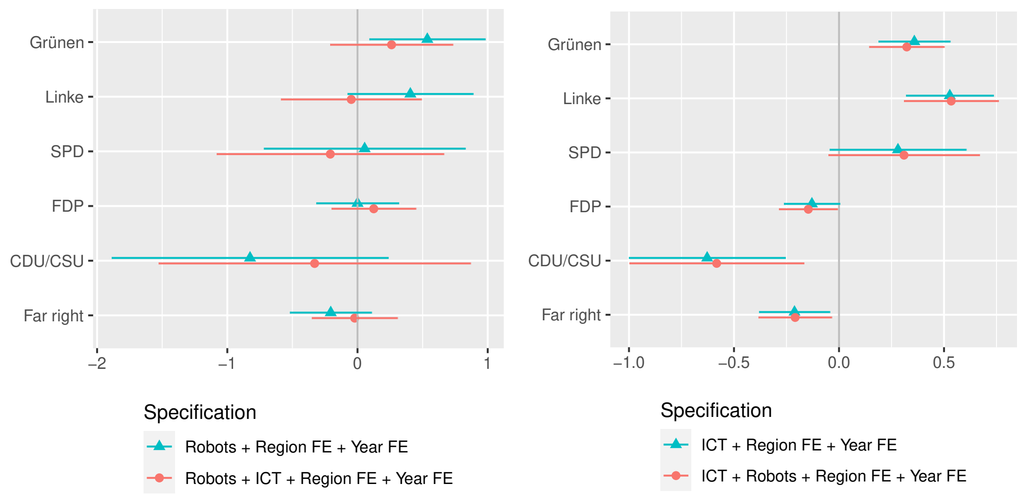

(a) Marginal Effect of Robot Intensity (b) Marginal Effect of ICT

Note: The graph shows estimated marginal effect of the (a) regional log number of robots per thousand workers and (b) the regional ICT capital stock per worker in 1000€ on regional party vote shares in percentage points (see Column (1) and (3) of Tables C.1-C.12). The sample consists of all federal, state and European elections between 1994 and 2017 measured on a county level (324 *Kreise und kreisfreie Städte*, NUTS-3) . AfD, DVU, Republikaner and NPD are coded as right-authoritarian parties.
Standard errors clustered at the county level. Bars represent 95% confidence intervals.

We run a series of robustness checks to increase confidence in our results (see Appendix C.2 for details). First, additional to the two-way fixed effects, we control for the trade exposure vis-à-vis China and Eastern Europe as an additional economic shock and GDP growth. We obtained data from the UN Comtrade database on industry level net-exports to construct another shift-share variable. Furthermore, we use an instrumental variable (IV) approach where we instrument technology adoption in Germany with values from other European countries.

Considering digitalization (ICT), effects are marginally attenuated when controlling for net exports or GDP growth. The IV results for ICT are in fact stronger than the OLS results: The progressive parties *Die Grünen* and *Die Linke* have larger positive coefficients whereas the socially-conservative and authoritarian-right end of the party spectrum has stronger negative effects. We conclude that the ICT results are very robust. Considering the automation in the form of robots, the results turn statistically insignificant when controlling for GDP or when considering the IV results. This again highlights that robotization may not have clear-cut political implications at the regional level. In additional robustness checks for the analysis on robots, we use an alternative specification of our main explanatory variable using the number of robots per thousand workers in levels rather than in logs. Using this alternative specification, the effects change substantially and in fact reverse to some extent, which seems to be a result of a few outlier regions with extreme robot concentration (Appendix C.2 provides details). It seems that different specification implying different relative weights of each region can be enough to change the interpretation of the results. The fact that results are quite unstable may be an indication that direct and compositional effect with respect to robotization are broadly on par but results should definitely be interpreted with caution.

3.5.2 Understanding Compositional Effects and Underlying Mechanisms

The remainder of the empirical exercise makes use of fine-grained individual and regional labor market data to trace underlying distributive implications of regional technology adoption. We first empirically confirm that the regional employment composition indeed shifts towards higher skilled and less routine occupations. Second, we show that the disappearing jobs are associated with conservative and authoritarian-right vote whereas the newly appearing jobs are associated with voting for more progressive parties. In sum, the analysis of intermediary distributive mechanisms on labor markets supports our conjecture that technological change results in a relative growth of occupations that are generally more supportive of progressive left parties.

Regional-Level Economic Outcomes

We first turn our attention to the economic effects of technology adoption by simply switching the dependent variable from voting results to labor market indicators. In line with much of the existing literature in labor economics (Michaels et al., 2014; Biagi and Falk, 2017; Dauth et al., 2021; Graetz and Michaels, 2018; Klenert et al., 2020; de Vries et al., 2020), we find that robot adoption and ICT investment affect the composition of the labor force but do not result in net employment loss. Both forms of technological innovation (if anything) marginally decreases manufacturing employment. Importantly, this decline in manufacturing is more than offset by an increase in the non-manufacturing (service) sector employment. The sum of both coefficients represents the effect of robot exposure on total employment relative to population. This hold considering each technology on its own or both jointly (see Figure 3.4 and Tables C.13-C.18).

Figure 3.4: Region-level exposure to robots and employment effects

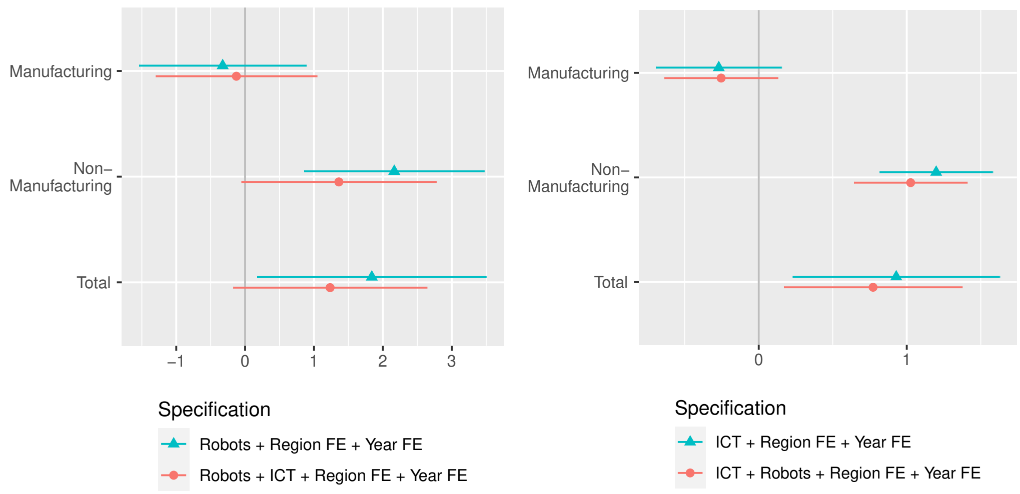

(a) Robot Intensity

(b) ICT

Note: Estimated coefficients of effect of log number of robots per thousand workers on employment to population ratios (in %) after controlling for region and year fixed effects. See column (1) of Table C.13 - C.15. Black bars represent 95% confidence intervals.

The point estimates imply that focusing on the within-region variation, a one standard deviation increase in robot exposure (+30% more robots) is associated with a decrease of the manufacturing employment to population ratio of -0.09 percentage points (not statistically significantly different from zero) and a statistically significant increase of the non-manufacturing employment to total population ratio of +0.65 percentage points. On the other hand, an increase of the ICT capital stock of one within-region standard deviation (+520€) is associated with a decrease of the manufacturing employment to population ratio of -0.14 percentage points (not statistically significantly different from zero) and a statistically significant increase of the non-manufacturing employment to total population ratio by +0.60 percentage points.

The main reason for an increase in aggregate employment is that the fall of routine jobs is often accompanied by disproportionate job growth in non-routine occupations (de Vries et al., 2020). Indeed, when looking at labor shares of task groups instead of sectors, we find that technology adoption increases non-routine cognitive jobs at the cost of routine jobs (see Figure 3.5). In line with our intuition, robots have a stronger replacement effect with respect to routine manual jobs whereas ICT investment substitutes in particular for routine cognitive occupations. The share of low-skilled manual non-routine jobs is not significantly affected by technology adoption in Germany.

This pattern of a ”polarized upgrading” (Oesch and Rodriguez-Menes, 2010) is largely confirmed when looking at labor shares by skill group. Technology-adopting regions experience a strong increase in the share of high-skilled jobs and stagnation or even decline in mid- and low-skill jobs (see Figure 3.6). ICT investment in particular seems to foster upskilling. Appendix Figure C.2 shows that the pattern looks similar but slightly more polarizing when looking at education requirements of a job rather than skill group.

Summing up, we show that regions with stronger exposure to technology adoption experience a polarized upgrading of labor markets. While overall employment is not negatively affected, the share (and numbers) of jobs in the semi-skilled and manufacturing domain decreases markedly. These findings are not in itself ground-breaking as they align with previous work on the labor market effects of automation. Nevertheless, they provide a vital first piece of evidence to strengthen our argumentation that compositional effects play an important role to understand how automation affects political preferences at a regional level.

Figure 3.5: Technological change and Regional Task Composition

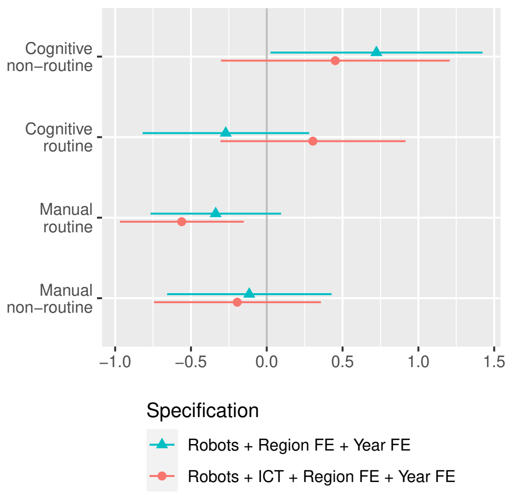

(a) Robot Intensity

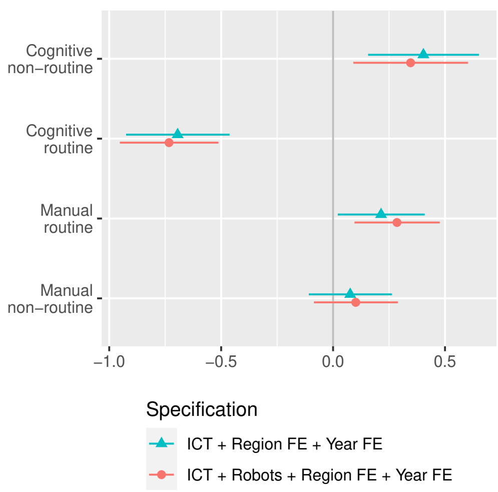

(b) ICT

Note: All variables are expressed as changes in regional employment shares in percentage points such that coefficients sum up to zero. Bars represent 95% confidence intervals where standard errors are clustered at the commuting zone-year level.

Regional-Level Relationship between Occupation and Vote Choice

To understand why technological change may shift the regional electoral landscape to the progressive left, it is important to analyze how the local labor

Figure 3.6: Technological change and Regional Skill Requirements

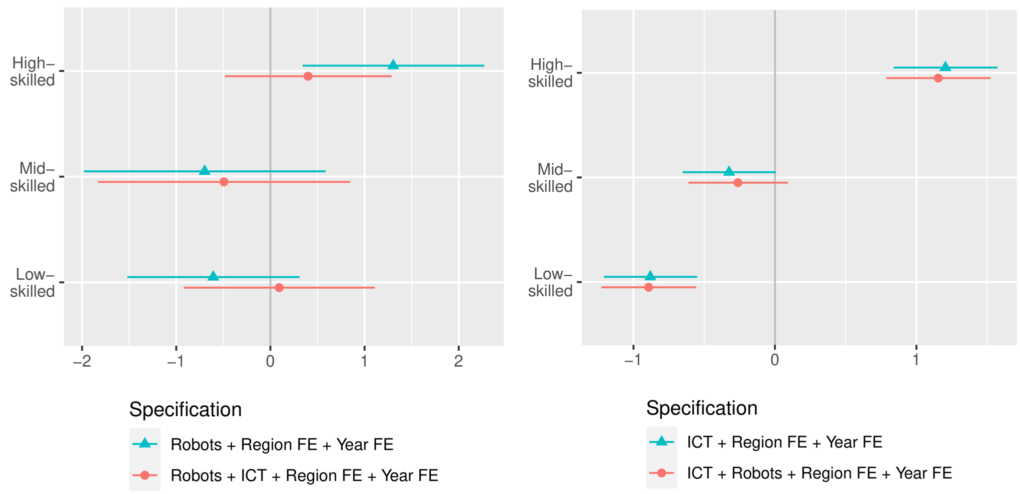

(a) Robot Intensity (b) ICT

Note: All variables are expressed as changes in regional employment shares in percentage points such that coefficients sum up to zero. Bars represent 95% confidence intervals where standard errors are clustered at the commuting zone-year level.

force composition affects voting outcomes. As we showed before, increased exposure to technology is associated with a shift from manufacturing to services, from (manual) routine occupations to non-routine (cognitive) occupations, from low- and mid-skilled jobs to high-skilled jobs and towards a more educated local workforce. According to the theory of occupational preference formation, all these changes in the labor market composition should shift political support more towards progressive parties. In order to corroborate these underlying expectations, the following analyses zoom in on the relationship between regional employment composition and party vote shares.

For this, we focus on the results of the 2017 Federal Elections (the last year in our sample) and regress the county-level party vote share on the local employment share as of 2017.[11] For each party p - employment share s (manufacturing share, routine worker share, etc.) pair we run a separate regression of the following kind:

$$VoteShare_r^p = \beta * EmploymentShare_r^s + \epsilon_r \tag{3.4}$$

where $VoteShare_r^p$ is the vote share of party p in region r which is regressed on the employment share of type s in region r.

The results presented in Figure 3.7 shows that a higher non-manufacturing (service) employment to population ratio is associated with more vote for

[11] Using previous election years leads to similar results.

progressive-left parties and a less support with conservatives and right-authoritarian parties. This closely resembles the effect of technological change on voting outcomes. On the other hand, conservatives and right-authoritarian parties perform particularly well where the manufacturing employment to population ratio is high (see Panel 3.7a).

Similarly, regional labor market characterised by a high share of cognitive non-routine occupations display more support for cosmopolitan-left parties less support for conservative and authoritarian-right parties. Conversely, regions with a large share of manual workers (both routine and non-routine) tend to be less support of the progressive left parties and more supportive of authoritarian right parties (see Panel 3.7b).

Furthermore, we find evidence that a high number of high-skilled workers is associated with more support for progressive left parties and less support for center-right and right-authoritarian parties. Conversely, it is mostly regions harboring more low-skilled and mid-skilled workers that are less supportive of progressive-left parties and more supportive of right-authoritarian parties. (see Panel 3.7c).

Finally, a higher share of highly educated workers is associated with more support for progressive workers and a lower support for conservative workers. The opposite is true looking at the share of workers with only intermediate levels of education. The share of High School drop-outs is weakly correlated with more support for authoritarian-right parties. However, the effects are imprecisely estimated (see Panel 3.7d).

Figure 3.7: Cross-sectional correlations of regional employment shares and party vote shares in 2017 Federal Elections

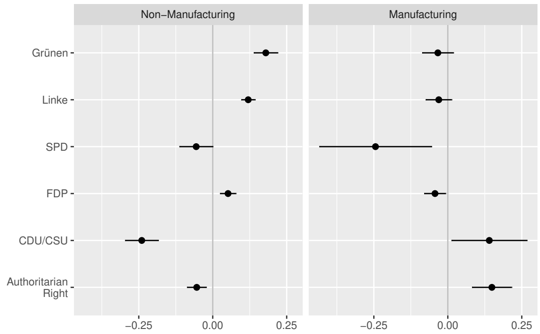

(a) Sector

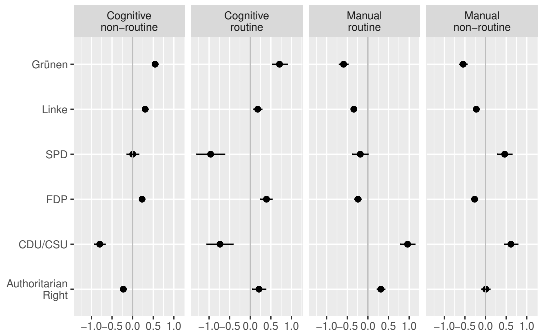

(b) Main task

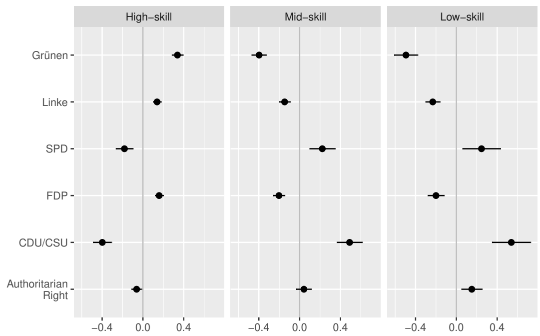

(c) Skill

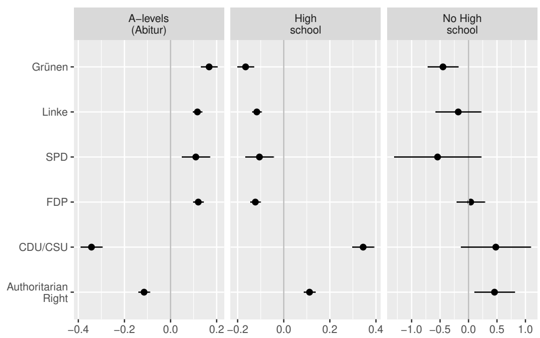

(d) Education

Note: Cross-sectional regression of regional party vote shares in 2017 federal elections on regional employment shares without controls (n=324 counties). The estimated coefficients are proportional to raw correlations. Bars represent 95% confidence intervals.

Individual-Level Relationship between Occupation and Vote Choice

To increase confidence in these patterns, we analyze the aggregate party preferences of different occupation groups using individual-level data from the German Socio-Economic Panel (SOEP). This allows us to test more directly how local labor market composition affects election results. Our aim here was to recreate the sectorial and occupational groups from the previous analysis as closely as possible to ensure comparability. Therefore, we considered all respondents between 18 and 65 for the years 1994 to 2018 (n=323000) and classified them into manufacturing and non-manufacturing, by main task (see Section 3.4 for details) and created three education groups ranging from low (High school drop-outs and basic secondary education, *Hauptschule*) over middle (intermediate secondary education, *Realschule*) to high (A-levels, *Abitur*). Figure 3.8 plots the party support of different occupational groups over time. To facilitate the visualization we grouped responses in 5-year intervals.

The findings confirm a few common priors of the relevant literature (e.g. Oesch, 2008; Kitschelt and Rehm, 2014). First, we find that the progressive-left party *Die Grünen* is mainly supported by non-manufacturing (service sector) workers whereas manufacturing workers became more and more supportive conservative and authoritarian-right parties over the last years (see panel 3.8a). Secondly, we observe the cognitive non-routine workers disproportionately support the progressive-left party *Die Grünen* whereas conservative parties are mainly supported by routine workers and authoritarian-right parties draw most support from manual occupations (both routine and non-routine) (see panel 3.8b). Finally, we find a strong education gradient. Highly educated workers are the core constituents of the green party (and the pro-market FDP) whereas conservative and far-right parties find most support among middle and low educated workers (see Panel 3.8c). This further corroborates the the idea that those occupational groups which expand due to technological change are more supportive of progressive-left parties whereas conservative and authoritarian-right parties find the size of occupational groups that mainly supported them to be in decline.A theory of occupational preference formation in tandem with a gradually changing composition of local labor markets hence provides a reasonable explanation of why technological innovation can shift the regional electoral landscape to the progressive left.

Figure 3.8: Party support of different segments of the workforce over time

(a) Sector

(b) Main Task

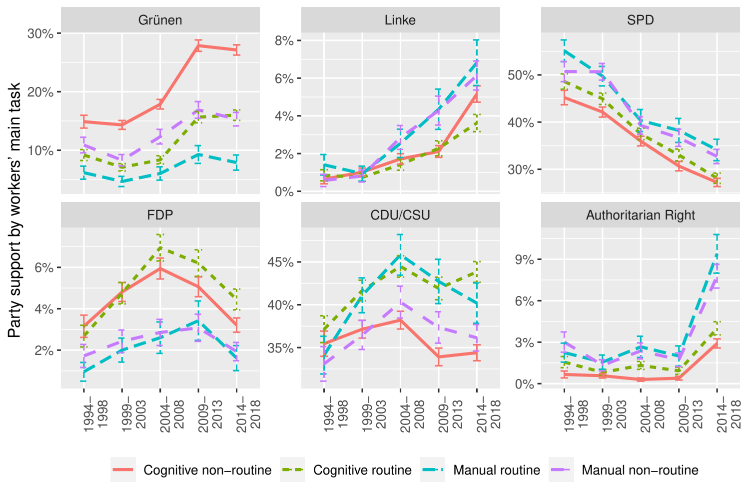

(c) Education

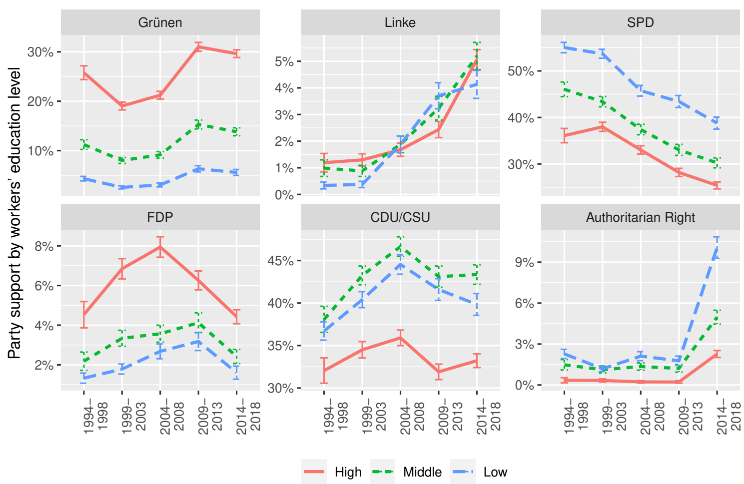

Note: Graphs show self-reported party support of different occupation groups over time (clustered into 5-year intervals). Bars represent 95% confidence intervals.

3.6 Discussion

In this paper, we demonstrate that, on average, technological innovation increased the *regional* vote shares of cosmopolitan left parties whereas right-authoritarian parties receive less votes in affected regions. The increased prevalence of robots and ICT changes the local labor market composition and shifts the employment structure from routine to non-routine jobs. This shift has important indirect consequences in that it opens more jobs for highly-educated, high skilled workers who often work on cognitive interactive tasks. Such ”children of digitalization” gravitate towards the cosmopolitan left whereas routine workers in manufacturing whose jobs were, as we show, partly replaced by robots, often feel attracted by the promises of right-wing populism. Hence, the common narrative that technological change and robotization will first and foremost result in political disruption may provide an incomplete perspective.

How can we reconcile our findings with previous work who showed evidence in favor of the populism narrative? Our study finds that *regions* exposed to robotization and digitalization tend to shift employment away from manufacturing and routine jobs, which in turn leads to less support for right-authoritarian parties. Hence, we would not expect that right-authoritarian parties make the strongest inroads in strongly technology-adopting regions. Here, the composition of local labor markets changes more substantially than in regions less exposed to technological change and economic modernization. And yet, it is important to repeat that we do *not* claim that technological innovation is unrelated to the recent surge in right-authoritarian and populist voting in Germany and elsewhere. The mounting evidence that automation increases right-authoritarian support among individuals or occupational groups that are imminently affected – or threatened – by automation is entirely plausible and convincing. However, we wish to highlight that the broader compositional changes in local labor markets work in the opposite direction and may well dominate the political response by those disaffected voters who lose out in the process of economic modernization.

Hence, we can resolve the apparent conflict by a conceptual differentiation between a compositional (regional) and a direct (individual) effect. This differentiation has important implications for future research, as it highlights the pros and cons of using a regional approach versus an occupational/individual level approach. The disadvantage of our regional analysis is its inability to isolate those workers directly threatened by technological innovation. Put differently, some disruptive political consequences of technological change ”might be masked [...] by the aggregate welfare gains brought about by automation” (Anelli et al., 2021, p. 4). This is exactly right: Our approach inherently bundles winners and losers within the unit of analysis. Depending on the workers’ skills and occupation, the adoption of technology can have either positive or negative effects, even if they live in the same region.

On the positive side, a regional approach allows us to capture the compositional effect of changing local labor markets, i.e. precisely the before-mentioned welfare *gains* in the aggregate. Recall that a focus on within-individual changes lets us focus on the direct effect but – by design – neglects the compositional effect. Positive indirect effects of technological innovation such as the creation of new jobs can only be captured by a regional approach. Also, the fact that new generations joining the labor market enter into different occupations and hold different political attitudes than previous generation is hidden when focusing on within-individual changes. The academic literature has shown that technological change mostly shapes employment composition through generational turnover rather than directly displacing affected workers. Hence, in the long term, the compositional effect may be considered more important and more consequential in political terms.

Appendix C

C.1 The Political Space in Germany

Figure C.1 shows that party positions in Germany are broadly aligned along one one dimension. They span from progressive-left (*Die Linke*) to authoritarian-right (*AFD* and other right-authoritarian parities). Notable exception is the pro-business party *FDP* which combines economic conservatism with social progressiveness. However, they do not play a central role in our analysis as their electoral support does not seem to be affected by robots adoption.

Figure C.1: Political Parties in the Two-Dimensional Space

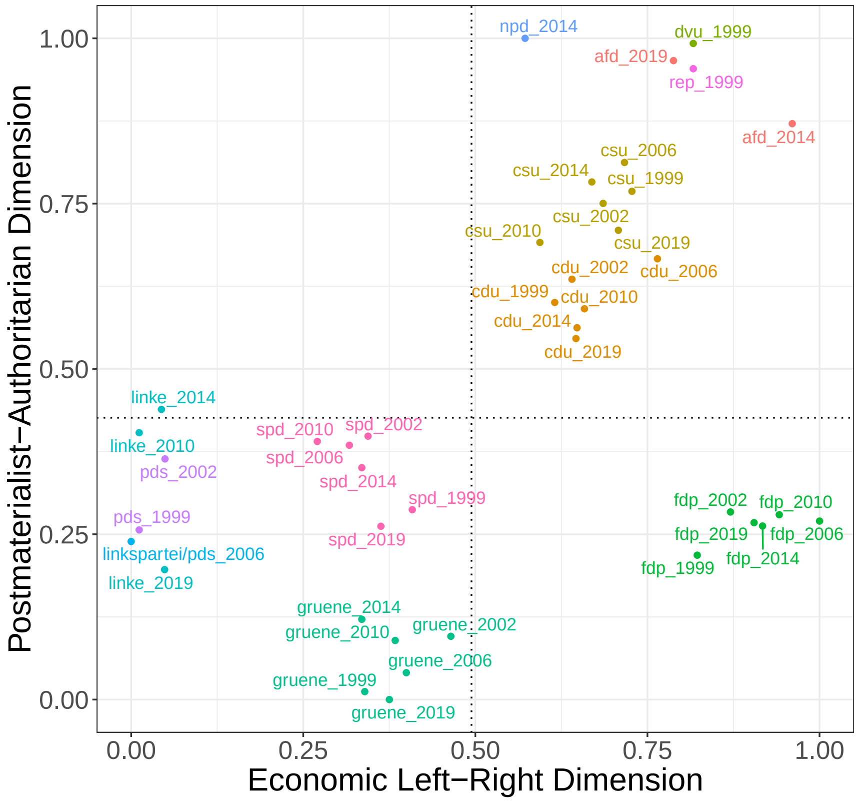

Note: Party positions based on Chapel Hill Expert Survey data between 1994 and 2019. Both lrecon and galtan dimensions are standardized between 0 and 1. The dotted lines show average values pooled over time weighted by party-seat share.

C.2 Robustness Checks

In this section we report in more detail on the robustness checks we briefly described in Section 3.5.1. We report one regression table for each economic and political outcome, once for robot adoption and once for ICT investment in the section that follows. In the first column of each table we present our baseline model which relies on county and year fixed effects.[1] Note that the two-way fixed effect specification is already quite demanding as it holds constant all factors that are either constant over time within a region (for example if a region belonged to the former GDR, an important factor to explain electoral differences in Germany) or common shocks to all regions in a given year (for example changing party platforms or external events that affect the general success of parties).

Next, we add economic shocks as control variable to rule out that our results suffer from omitted variable bias. In column (2) of each table, we control for the net trade balance of each region vis-à-vis China and Eastern Europe. This is important as thriving manufacturing regions, which adopt robots at a fast pace are likely to also be more involved in international trade. At the same time, it has been shown that trade exposure affects the political preferences of voters (Dippel et al., 2015; Colantone and Stanig, 2018). We find that this is not a major confounder as the unconditional correlation of net exports and robot intensity (0.04) or ICT (0.12) is low and also the estimated effect of regional robot intensity and regional ICT investment on regional election results and regional economic outcomes remain stable. Column (3) includes the other source of technological change as an additional control. Again, the concern is that it is an alternative economic shock is correlated with our technology shock.As noted before, the correlation between per worker ICT capital stocks and robot intensity is rather low (0.12). The effect of robotization on voting patterns virtually disappears after controlling for ICT. The effect of ICT on regional-level election outcomes on the other hand is not affected. As a third control, we include GDP per capita (column 4). This is important as robot adoption could be just one symptom of generally thriving regions (on the other hand, it could also be argued that GDP is a bad control as it part of the mechanism how technological change affects economic and political outcomes). Similar to controlling for the influence of the other technology, the point estimates of ICT on voting shares is not affected whereas there is no effect robotization on party support after controlling for GDP growth. Regarding the labor market consequences of technological change, it turns out that point estimates become more negative after controlling

[1]To be precise, we use election fixed effects for political outcomes. These differ from year fixed effects in the case of state elections as each state has its own fixed effect.

for GDP growth. This is intuitive as newly created job usually go hand in hand with economic growth.

Next, we use an instrumental variable approach where we instrument industry-level technology adoption in Germany with values from other European countries.[2] As argued before, the pace of robot adoption or ICT investment might be influenced by surrounding labor market institutions. In Germany, workers councils and trade unions have been shown to affect the process how companies digitalize (Genz et al., 2019). Simultaneously, labor unions have strong linkages to leftist and social democratic parties which could create an omitted variable bias in our OLS estimates. Using the speed of adoption in other European countries as a valid instrument implies the exclusion restriction that specific labor market and political institution in Germany do not affect industry level decision to adopt new technologies abroad. Instead, it is assumed to be driven by a technological frontier. In a second panel of each table we replicate all specifications using a 2SLS estimator. We find that labor market outcomes are comparable to the OLS estimates when considering ICT. Again, for robotization the result are less stable. Concerning the case of robots, it has been noted that despite the strong first stage, using other Western countries as an instrument might be problematic in the case of Germany as it precedes other Western countries when it comes to adopting robots. Nevertheless, we included the instrumental variable analysis to facilitate the comparison to previous research.

Finally, we use the number of robots per thousand workers in levels (not in logs) as main explanatory variable (third panel). This gives more weight to outlier region (recall that a few manufacturing hot-spots attracted the bulk of new robots). The voting pattern results completely change and this analysis suggests that automation is associated with less support for progressive-left parties and more support for conservative and authoritarian-right parties. However, as is shown in the last panel of each table, this pattern reverts if we exclude the top ten regions in terms of robot intensity. The estimated labor market consequences of both specifications are similar and in line with the results described previously. This suggests that the general distributive effects are captured with either approach. However, voting results depend on the specification. We interpret this as further evidence that here, the compositional and the treatment are of similar strength.

Summing up, we find stable results for ICT with respect to voting and labor market outcomes. Regarding robotization, the labor market effects are relatively robust, the political consequences are robotization are not robust.

[2]For robotization, we use data on all European countries included in the IFR database: Sweden, Denmark, Italy, Belgium, Netherlands, Austria, Slovenia, Spain, Slovakia, France Finland, Czech Republic. For ICT, we use data from all other EU member state countries (EU28 including the UK).

C.3 Regression Tables

C.3.1 Robots & Election Outcomes

Table C.1: Fixed-Effects Estimation of robot exposure on support for *Die Grünen*

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	0.536**	0.571**	0.261	0.364	0.278
	(0.227)	(0.230)	(0.240)	(0.237)	(0.248)
Net Exports		−0.036			−0.045
		(0.033)			(0.031)
ICT			0.323***		0.216**
			(0.091)		(0.101)
GDP per capita				0.038***	0.034***
				(0.012)	(0.012)
2SLS					
Robots	−0.204	−0.182	−0.270	−0.239	−0.213
	(0.327)	(0.330)	(0.321)	(0.355)	(0.353)
First-stage F-stat	252.44	124.72	153.61	146.35	78.66
Non-logged robots					
Robots	0.0004	0.001	−0.014*	−0.020***	−0.026***
	(0.007)	(0.007)	(0.007)	(0.007)	(0.008)
Non-logged robots exclude outliers					
Robots	0.024	0.026	0.007	−0.004	−0.007
	(0.025)	(0.025)	(0.023)	(0.025)	(0.023)
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.937	0.937	0.937	0.937	0.938

Note: Fixed-effects regressions of party vote share (in %) on log number of robots per 1000 workers for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*), once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*) and once using robots in levels but excluding 10 outlier counties (*Non-logged robots exclude outliers*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county. *p<0.1; **p<0.05; ***p<0.01

Table C.2: Fixed-Effects Estimation of robot exposure on support for *Die Linke*

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	0.406*	0.382	−0.048	0.258	−0.105
	(0.246)	(0.244)	(0.275)	(0.252)	(0.271)
Net Exports		0.024			0.021
		(0.022)			(0.022)
ICT			0.534***		0.561***
			(0.115)		(0.116)
GDP per capita				0.011	−0.004
				(0.008)	(0.008)
2SLS					
Robots	0.388	0.365	0.315	0.265	0.232
	(0.351)	(0.351)	(0.371)	(0.352)	(0.371)
First-stage F-stat	236.7	116.74	141.53	142.26	74.08
Non-logged robots					
Robots	0.011	0.010	−0.009	0.002	−0.012
	(0.009)	(0.009)	(0.010)	(0.012)	(0.012)
Non-logged robots exclude outliers					
Robots	0.057*	0.055*	0.035	0.042	0.028
	(0.030)	(0.029)	(0.030)	(0.028)	(0.028)
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	3,792	3,792	3,792	3,651	3,651
Adjusted R^2	0.888	0.888	0.890	0.892	0.894

Note: Fixed-effects regressions of party vote share (in %) on log number of robots per 1000 workers for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*) and once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county. *p<0.1; **p<0.05; ***p<0.01

Table C.3: Fixed-Effects Estimation of robot exposure on support for SPD

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	0.055	0.082	−0.209	−0.261	−0.442
	(0.394)	(0.400)	(0.445)	(0.416)	(0.447)
Net Exports		−0.028			−0.026
		(0.044)			(0.044)
ICT			0.309*		0.331*
			(0.183)		(0.197)
GDP per capita				0.016	0.007
				(0.019)	(0.019)
2SLS					
Robots	−0.215	−0.193	−0.268	−0.783	−0.784
	(0.684)	(0.694)	(0.698)	(0.591)	(0.599)
First-stage F-stat	252.44	124.72	153.61	146.35	78.66
Non-logged robots					
Robots	−0.010	−0.010	−0.023	−0.033*	−0.041**
	(0.012)	(0.012)	(0.014)	(0.019)	(0.020)
Non-logged robots exclude outliers					
Robots	−0.026	−0.025	−0.041	−0.062*	−0.068*
	(0.038)	(0.038)	(0.039)	(0.037)	(0.038)
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.962	0.962	0.963	0.962	0.962

Note: Fixed-effects regressions of party vote share (in %) on log number of robots per 1000 workers for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*) and once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*). All models include region and year fixed effects. Standard errors reported in parenthesis are clustered by county. *p<0.1; **p<0.05; ***p<0.01

Table C.4: Fixed-Effects Estimation of robot exposure on support for FDP

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	0.0002	0.038	0.125	−0.070	0.092
	(0.161)	(0.162)	(0.165)	(0.169)	(0.168)
Net Exports		−0.039**			−0.042**
		(0.019)			(0.019)
ICT			−0.146**		−0.183**
			(0.071)		(0.074)
GDP per capita				0.011	0.017**
				(0.007)	(0.006)
2SLS					
Robots	−0.060	−0.026	−0.040	−0.097	−0.038
	(0.253)	(0.255)	(0.258)	(0.250)	(0.253)
First-stage F-stat	252.44	124.72	153.61	146.35	78.66
Non-logged robots					
Robots	0.003	0.003	0.008	0.002	0.007
	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)
Non-logged robots exclude outliers					
Robots	0.004	0.007	0.010	−0.005	0.001
	(0.014)	(0.014)	(0.013)	(0.014)	(0.013)
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.917	0.917	0.917	0.917	0.918

Note: Fixed-effects regressions of party vote share (in %) on log number of robots per 1000 workers for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*) and once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county. *p<0.1; **p<0.05; ***p<0.01

Table C.5: Fixed-Effects Estimation of robot exposure on support for CDU / CSU

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	−0.825	−0.926*	−0.329	−0.035	0.022
	(0.542)	(0.543)	(0.611)	(0.586)	(0.612)
Net Exports		0.103			0.121
		(0.078)			(0.076)
ICT			−0.582***		−0.302
			(0.212)		(0.229)
GDP per capita				−0.096***	−0.091**
				(0.036)	(0.037)
2SLS					
Robots	0.073	−0.008	0.182	0.765	0.660
	(0.949)	(0.957)	(0.965)	(0.859)	(0.867)
First-stage F-stat	252.44	124.72	153.61	146.35	78.66
Non-logged robots					
Robots	−0.017	−0.018	0.006	0.036	0.042
	(0.019)	(0.019)	(0.021)	(0.027)	(0.027)
Non-logged robots exclude outliers					
Robots	−0.048	−0.053	−0.023	0.043	0.043
	(0.058)	(0.058)	(0.058)	(0.059)	(0.059)
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.924	0.924	0.924	0.926	0.926

Note: Fixed-effects regressions of party vote share (in %) on log number of robots per 1000 workers for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*) and once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county. *p<0.1; **p<0.05; ***p<0.01

Table C.6: Fixed-Effects Estimation of robot exposure on support for right-authoritarian Parties

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	−0.205	−0.216	−0.022	−0.295	−0.118
	(0.160)	(0.163)	(0.168)	(0.180)	(0.191)
Net Exports		0.013			0.013
		(0.029)			(0.029)
ICT			−0.209**		−0.306***
			(0.089)		(0.095)
GDP per capita				0.006	0.014**
				(0.007)	(0.007)
2SLS					
Robots	−0.294	−0.304	−0.265	−0.373	−0.381
	(0.209)	(0.212)	(0.212)	(0.245)	(0.247)
First-stage F-stat	234.07	116.32	139.11	142.37	73.7
Non-logged robots					
Robots	0.008*	0.008*	0.019***	0.006	0.015**
	(0.005)	(0.005)	(0.005)	(0.007)	(0.007)
Non-logged robots exclude outliers					
Robots	−0.022	−0.022	−0.009	−0.019	−0.011
	(0.014)	(0.014)	(0.013)	(0.016)	(0.014)
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	3,197	3,197	3,197	3,056	3,056
Adjusted R^2	0.925	0.925	0.926	0.924	0.925

Note: Fixed-effects regressions of party vote share (in %) on log number of robots per 1000 workers for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*) and once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county. *p<0.1; **p<0.05; ***p<0.01

C.3.2 ICT & Election Outcomes

Table C.7: Fixed-Effects Estimation of ICT on support for *Die Grünen*

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	0.359***	0.363***	0.323***	0.240**	0.216**
	(0.087)	(0.087)	(0.091)	(0.099)	(0.101)
Net Exports		−0.032			−0.045
		(0.031)			(0.031)
Robots			0.261		0.278
			(0.240)		(0.248)
GDP per capita				0.033***	0.034***
				(0.012)	(0.012)
2SLS					
ICT	0.732***	0.734***	0.742***	0.660***	0.659***
	(0.152)	(0.152)	(0.164)	(0.169)	(0.179)
First-stage F-stat	296.46	147.08	127.71	137.83	71.04
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.937	0.937	0.937	0.937	0.938

Note: Fixed-effects regressions of county-level party vote share (in %) on ICT capital stocks per worker (in 1000 €) for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per 1000 workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.8: Fixed-Effects Estimation of ICT on support for *Die Linke*

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	0.527***	0.524***	0.534***	0.552***	0.561***
	(0.106)	(0.106)	(0.115)	(0.112)	(0.116)
Net Exports		0.022			0.021
		(0.022)			(0.022)
Robots			−0.048		−0.105
			(0.275)		(0.271)
GDP per capita				−0.004	−0.004
				(0.008)	(0.008)
2SLS					
ICT	0.691***	0.689***	0.711***	0.722***	0.740***
	(0.164)	(0.164)	(0.180)	(0.171)	(0.179)
First-stage F-stat	246.95	122.56	105.71	112.41	58.55
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	3,792	3,792	3,792	3,651	3,651
Adjusted R^2	0.890	0.890	0.890	0.894	0.894

Note: Fixed-effects regressions of county-level party vote share (in %) on ICT capital stocks per worker (in 1000 €) for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per 1000 workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.9: Fixed-Effects Estimation of ICT on support for SPD

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	0.281*	0.285*	0.309*	0.282	0.331*
	(0.166)	(0.166)	(0.183)	(0.191)	(0.197)
Net Exports		−0.031			−0.026
		(0.043)			(0.044)
Robots			−0.209		−0.442
			(0.445)		(0.447)
GDP per capita				0.005	0.007
				(0.019)	(0.019)
2SLS					
ICT	0.078	0.079	0.079	0.059	0.084
	(0.253)	(0.254)	(0.279)	(0.283)	(0.296)
First-stage F-stat	296.46	147.08	127.71	137.83	71.04
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.963	0.963	0.963	0.962	0.962

Note: Fixed-effects regressions of county-level party vote share (in %) on ICT capital stocks per worker (in 1000 €) for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per 1000 workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.10: Fixed-Effects Estimation of ICT on support for FDP

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	−0.129*	−0.124*	−0.146**	−0.177**	−0.183**
	(0.068)	(0.068)	(0.071)	(0.073)	(0.074)
Net Exports		−0.036**			−0.042**
		(0.018)			(0.019)
Robots			0.125		0.092
			(0.165)		(0.168)
GDP per capita				0.016**	0.017**
				(0.007)	(0.006)
2SLS					
ICT	−0.051	−0.049	−0.056	−0.091	−0.094
	(0.115)	(0.115)	(0.125)	(0.126)	(0.131)
First-stage F-stat	296.46	147.08	127.71	137.83	71.04
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.917	0.917	0.917	0.918	0.918

Note: Fixed-effects regressions of county-level party vote share (in %) on ICT capital stocks per worker (in 1000 €) for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per 1000 workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.11: Fixed-Effects Estimation of ICT on support for CDU / CSU

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	-0.628^{***}	-0.642^{***}	-0.582^{***}	-0.288	-0.302
	(0.190)	(0.193)	(0.212)	(0.221)	(0.229)
Net Exports		0.097			0.121
		(0.076)			(0.076)
Robots			-0.329		0.022
			(0.611)		(0.612)
GDP per capita				-0.087^{**}	-0.091^{**}
				(0.037)	(0.037)
2SLS					
ICT	-0.835^{**}	-0.840^{**}	-0.822^{**}	-0.528	-0.542
	(0.324)	(0.327)	(0.345)	(0.378)	(0.391)
First-stage F-stat	296.46	147.08	127.71	137.83	71.04
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	4,276	4,276	4,276	4,135	4,135
Adjusted R^2	0.924	0.925	0.924	0.926	0.926

Note: Fixed-effects regressions of county-level party vote share (in %) on ICT capital stocks per worker (in 1000 €) for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per 1000 workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.12: Fixed-Effects Estimation of ICT on support for right-authoritarian Parties

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	−0.212**	−0.214**	−0.209**	−0.317***	−0.306***
	(0.086)	(0.086)	(0.089)	(0.093)	(0.095)
Net Exports		0.013			0.013
		(0.028)			(0.029)
Robots			−0.022		−0.118
			(0.168)		(0.191)
GDP per capita				0.014*	0.014**
				(0.007)	(0.007)
2SLS					
ICT	−0.382***	−0.384***	−0.398***	−0.475***	−0.475***
	(0.122)	(0.123)	(0.131)	(0.138)	(0.142)
First-stage F-stat	210.8	104.69	88.26	93.46	48.56
Region FE	X	X	X	X	X
Election FE	X	X	X	X	X
Observations	3,197	3,197	3,197	3,056	3,056
Adjusted R^2	0.926	0.926	0.926	0.925	0.925

Note: Fixed-effects regressions of county-level party vote share (in %) on ICT capital stocks per worker (in 1000 €) for federal, state and European Elections. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per 1000 workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and election fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

C.3.3 Robots & Labor Market Composition

Table C.13: Fixed-Effects Estimation of robot exposure on total employment

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	1.840**	1.811**	1.234*	−0.264	−0.192
	(0.849)	(0.849)	(0.717)	(0.471)	(0.454)
Net Exports		0.029			0.013
		(0.063)			(0.053)
ICT			0.772**		−0.165
			(0.307)		(0.196)
GDP per capita				0.238***	0.242***
				(0.023)	(0.025)
2SLS					
Robots	0.991	0.947	0.946	0.291	0.268
	(0.832)	(0.833)	(0.802)	(0.578)	(0.585)
First-stage F-stat	202.29	98.9	121.66	141.65	76.81
Non-logged robots					
Robots	0.138***	0.138***	0.127**	0.003	0.006
	(0.049)	(0.049)	(0.049)	(0.018)	(0.017)
Non-logged robots exclude outliers					
Robots	0.155**	0.152**	0.143**	0.030	0.034
	(0.067)	(0.066)	(0.067)	(0.046)	(0.044)
Region FE	X	X	X	X	X
Year FE	X	X	X	X	X
Observations	7,774	7,774	7,774	7,492	7,492
Adjusted R^2	0.978	0.978	0.978	0.985	0.985

Note: Fixed-effects regressions of total employment to population ratio (in %) on log number of robots per 1000 workers. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*), once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*) and once using robots in levels but excluding 10 outlier counties (*Non-logged robots exclude outliers*). All models include region and year fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.14: Fixed-Effects Estimation of robot exposure on manufacturing employment

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	−0.326	−0.339	−0.127	−1.219**	−0.867
	(0.619)	(0.623)	(0.597)	(0.559)	(0.628)
Net Exports		0.013			0.009
		(0.062)			(0.065)
ICT			−0.254		−0.684***
			(0.196)		(0.172)
GDP per capita				0.081***	0.099***
				(0.020)	(0.021)
2SLS					
Robots	0.026	0.019	0.042	−0.292	−0.347
	(0.596)	(0.603)	(0.599)	(0.660)	(0.684)
First-stage F-stat	202.29	98.9	121.66	141.65	76.81
Non-logged robots					
Robots	0.045	0.045	0.061**	−0.009	0.005
	(0.029)	(0.030)	(0.029)	(0.020)	(0.018)
Non-logged robots exclude outliers					
Robots	0.010	0.009	0.028	−0.017	−0.007
	(0.057)	(0.058)	(0.056)	(0.053)	(0.050)
Region FE	X	X	X	X	X
Year FE	X	X	X	X	X
Observations	7,774	7,774	7,774	7,492	7,492
Adjusted R^2	0.958	0.958	0.959	0.965	0.966

Note: Fixed-effects regressions of manufacturing employment to population ratio (in %) on log number of robots per 1000 workers. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*), once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*) and once using robots in levels but excluding 10 outlier counties (*Non-logged robots exclude outliers*). All models include region and year fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.15: Fixed-Effects Estimation of robot exposure on non-manufacturing employment

	(1)	(2)	(3)	(4)	(5)
OLS					
Robots	2.166***	2.149***	1.361*	0.955*	0.675
	(0.667)	(0.673)	(0.722)	(0.515)	(0.587)
Net Exports		0.016			0.004
		(0.041)			(0.038)
ICT			1.026***		0.519***
			(0.195)		(0.164)
GDP per capita				0.157***	0.143***
				(0.027)	(0.027)
2SLS					
Robots	0.964	0.928	0.904	0.583	0.615
	(0.745)	(0.755)	(0.724)	(0.599)	(0.623)
First-stage F-stat	202.29	98.9	121.66	141.65	76.81
Non-logged robots					
Robots	0.093***	0.092***	0.065***	0.012	0.001
	(0.023)	(0.023)	(0.024)	(0.014)	(0.014)
Non-logged robots exclude outliers					
Robots	0.145***	0.142***	0.115***	0.048	0.041
	(0.045)	(0.045)	(0.044)	(0.036)	(0.036)
Region FE	X	X	X	X	X
Year FE	X	X	X	X	X
Observations	7,774	7,774	7,774	7,492	7,492
Adjusted R^2	0.979	0.979	0.980	0.984	0.984

Note: Fixed-effects regressions of non-manufacturing employment to population ratio (in %) on log number of robots per 1000 workers. Column (2) adds net exports per worker (in 1000 €), column (3) adds ICT capital stocks per worker (in 1000 €), column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest same specifications as above. Once instrumenting robot adoption in Germany with values from other EU countries (*2SLS*), once using the number of robots per 1000 workers in levels instead of logs (*Non-logged robots*) and once using robots in levels but excluding 10 outlier counties (*Non-logged robots exclude outliers*). All models include region and year fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

C.3.4 ICT & Labor Market Composition

Table C.16: Fixed-Effects Estimation of ICT on total employment

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	0.929***	0.922**	0.772**	−0.179	−0.165
	(0.357)	(0.357)	(0.307)	(0.202)	(0.196)
Net Exports		0.042			0.013
		(0.060)			(0.053)
Robots			1.234*		−0.192
			(0.717)		(0.454)
GDP per capita				0.242***	0.242***
				(0.024)	(0.025)
2SLS					
ICT	0.034	0.031	−0.175	−0.746***	−0.756***
	(0.323)	(0.322)	(0.361)	(0.284)	(0.287)
First-stage F-stat	217.32	108.13	96.27	97.68	52.87
Region FE	X	X	X	X	X
Year FE	X	X	X	X	X
Observations	7,774	7,774	7,774	7,492	7,492
Adjusted R^2	0.978	0.978	0.978	0.985	0.985

Note: Fixed-effects regressions of total employment to population ratio (in %) on log number of robots per 1000 workers. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per thousand workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest using the same specification as above while instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and year fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.17: Fixed-Effects Estimation of ICT on manufacturing employment

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	−0.270	−0.272	−0.254	−0.759***	−0.684***
	(0.216)	(0.218)	(0.196)	(0.161)	(0.172)
Net Exports		0.011			0.009
		(0.063)			(0.065)
Robots			−0.127		−0.867
			(0.597)		(0.628)
GDP per capita				0.095***	0.099***
				(0.023)	(0.021)
2SLS					
ICT	−1.067***	−1.068***	−1.125***	−1.445***	−1.402***
	(0.263)	(0.263)	(0.275)	(0.318)	(0.281)
First-stage F-stat	217.32	108.13	96.27	97.68	52.87
Region FE	X	X	X	X	X
Year FE	X	X	X	X	X
Observations	7,774	7,774	7,774	7,492	7,492
Adjusted R^2	0.959	0.959	0.959	0.966	0.966

Note: Fixed-effects regressions of total employment to population ratio (in %) on log number of robots per 1000 workers. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per thousand workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest using the same specification as above while instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and year fixed effects. Standard errors reported in parenthesis are clustered by county: *p<0.1; **p<0.05; ***p<0.01.

Table C.18: Fixed-Effects Estimation of ICT on non-manufacturing employment

	(1)	(2)	(3)	(4)	(5)
OLS					
ICT	1.199***	1.194***	1.026***	0.580***	0.519***
	(0.195)	(0.195)	(0.195)	(0.150)	(0.164)
Net Exports		0.031			0.004
		(0.041)			(0.038)
Robots			1.361*		0.675
			(0.722)		(0.587)
GDP per capita				0.147***	0.143***
				(0.028)	(0.027)
2SLS					
ICT	1.101***	1.099***	0.950***	0.700**	0.646**
	(0.377)	(0.376)	(0.331)	(0.313)	(0.287)
First-stage F-stat	217.32	108.13	96.27	97.68	52.87
Region FE	X	X	X	X	X
Year FE	X	X	X	X	X
Observations	7,774	7,774	7,774	7,492	7,492
Adjusted R^2	0.980	0.980	0.980	0.984	0.984

Note: Fixed-effects regressions of total employment to population ratio (in %) on log number of robots per 1000 workers. Column (2) adds net exports per worker (in 1000 €), column (3) adds log number of robots per thousand workers, column (4) adds GDP per capita (in 1000 €). Column (5) adds all three controls jointly. Below are reported the estimates for our variable of interest using the same specification as above while instrumenting ICT capital stocks in Germany with values from other EU countries (*2SLS*). All models include region and year fixed effects. Standard errors reported in parenthesis are clustered by county: *$p<0.1$; **$p<0.05$; ***$p<0.01$.

Figure C.2: Technological change and Regional Education Levels

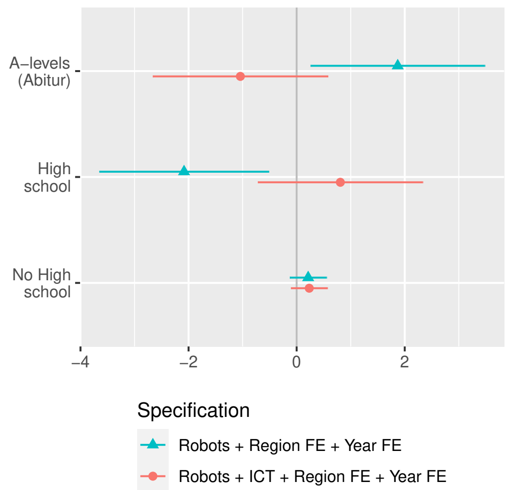

(a) Robot Intensity

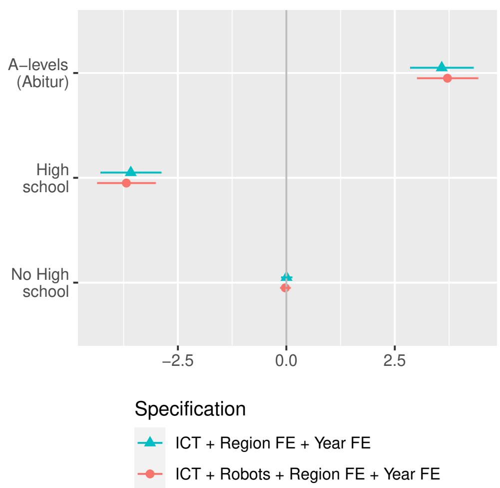

(b) ICT

Note: All variables are expressed as changes in regional employment shares in percentage points such that coefficients sum up to zero. Bars represent 95% confidence intervals where standard errors are clustered at the commuting zone-year level.

C.4 Mechanisms: Skill Requirements instead of Education

Similar to our findings presented in the main body of the text, we also find that education requirements are changing in technology-exposed regions with a more pronounced polarization of within the workforce (see Figure C.2). Investment in robots or ICT increases the share of workers with at least a university entrance degree (Abitur) but decreases the share of workers with only High school degrees. Interestingly, with respect to education requirements, we find some evidence of polarizing labor markets in the sense that technology adoption does not reduce the the share of workers who did not finish secondary school. These workers presumably find jobs in low-skilled services which are created due to positive spillover effects of technology adoption (see Figure C.2). The described patterns are generally robust to controlling for the other type of technology adapted. Only the effect of robotization on the education composition of the labor force changes markedly. This again supports the conjecture that ICT has a stronger impact on the overall labor force than robotization, a reasonable finding in light of the strong concentration of robots in a few highly-exposed sectors.

C.5 Reconciling Conflicting Evidence

Our findings conflict with more gloomy projections put forward in the media and the public debate but also in important scholarly work, in which especially robotization has been related to the rise of populism and the success of right-authoritarian parties (Frey et al., 2018; Im et al., 2019; Anelli et al., 2019; Kurer, 2020; Milner, 2021).

We argued before that the differences partially stem from different level of analysis. Studies with individual level data focus on what we dubbed the *direct effect*: those threatened by technological innovation (normally measured through robotization) become more supportive of authoritarian right parties.

However, some studies also use regional data which includes the *compositional effect* and still find a shift towards authoritarian right parties. This is at odds with our theory and empirical findings and in the following, we will try to determine why this is the case.

For this, we replicate the work of Anelli et al. (2019) (henceforth ACS) which also inspired the work of Milner (2021). Their regional analysis is the study most similar to our setting. ACS find that Western European regions with increasing robot exposure became more supportive of right-authoritarian parties. They argue that displaced or economically threatened manufacturing workers turned to right-authoritarian parties as they felt left behind.

The aim of this replication exercise is to determine where conflicting results stem from. We consider three potential explanations. First, it could be that West Germany is a special case, which deviates from the general patterns across Western European democracies presented in ACS. Recall the exceptional importance of robots in West German manufacturing or the fact that, at least partly for historical reasons, no right-authoritarian party was represented in German national parliament until 2017. Second, the competing results could arise from differences in how the data set is constructed. Our studies differ in the level of regional disaggregation, the sample period, and which election types are used. In particular, we have much more fine-grained levels of regional disaggregation (NUTS-3) than the data ACS use to analyze the German case (NUTS-1). Third, it could be that distinct modelling approaches make the difference. While we employ a fixed effect panel model, ACS rely on a repeated short difference specification.

We present the details of our replication exercise in the following sections. In short, we demonstrate that the conflicting evidence is a consequence of different modelling approaches. In our understanding, the reason for diverging results lies in the difference between levels and changes. What the ACS approach captures is that new robots are installed where the *level* of manufacturing employment is high. This is intuitive as industrial robots are most needed in manufacturing hot-spots. As we showed before, manufacturing workers (who are also concentrated in manufacturing hot-spots) feel attracted by right-authoritarian parties and hence there is a positive correlation between the change in robot exposure, the level of the manufacturing share and the level of right-authoritarian support. Yet, this does not prove that regional growth in robot exposure leads to positive *change* in regional right-authoritarian support. Our approach using fixed effects instead captures how the *change* in robot exposure affects the *change* in partisan support. To be sure, regions initially specialized in manufacturing adopted more robots and were generally more supportive of right-authoritarian parties. Yet, their support of right-authoritarian parties grew slower than in regions without robots due to the compositional effect: robots increased the local share of, for example, socio-cultural professionals in sectors with a strong emphasis on interpersonal interactions, which in turn limited the appeal of right-authoritarian parties.[3]

C.5.1 Replication of Anelli et al 2019

ACS regress party vote shares of right-authoritarian parties (in levels) on robot exposure (in changes) as a repeated cross section. They define regional robot exposure as the change of regional robot intensity (robots per thousand workers) in the two years prior to the elections (short-difference approach). Since they do not dispose over the data on employment composition at the county level (NUTS-3), they calculate robot exposure at the broader state level (NUTS-1 for Germany). Election results are measured at the more fine-grained county level. The estimated model is:

$$Y_{r,s,t} = \beta_1(Robots_{s,t-1} - Robots_{s,t-n}) + \mu_{election} + \epsilon_{r,s,t} \qquad \text{(C.1)}$$

where $Y_{r,s,t}$ is the electoral outcome in region r located in state s in year t. The difference between $Robots_{s,t-1}$ and $Robots_{s,t-n}$ (number or robots per 1000 employees on state level) expresses their measure of robot exposure. $\mu_{election}$ is an election specific fixed effect. Contrary to our model, no geographic fixed effects are used.

[3]Our rich data set allows to provide further empirical support for this line of argumentation by studying intermediary economic outcomes of robotization relying on ACS' empirical strategy. This empirical exercise shows why the mix of changes in the independent variable and levels in the dependent variables might be problematic. In fact, strictly applying ACS's modelling approach suggests that robotization leads to *more* manufacturing, routine and mid-skilled jobs (see Appendix C.5). Hence, there seems to be a mismatch between the theoretically hypothesized mechanism (robots threaten the jobs of manufacturing and routine workers) and the empirical reality resulting from the applied modelling approach (robots increase manufacturing and routine employment).

Figure C.3: Replication ACS

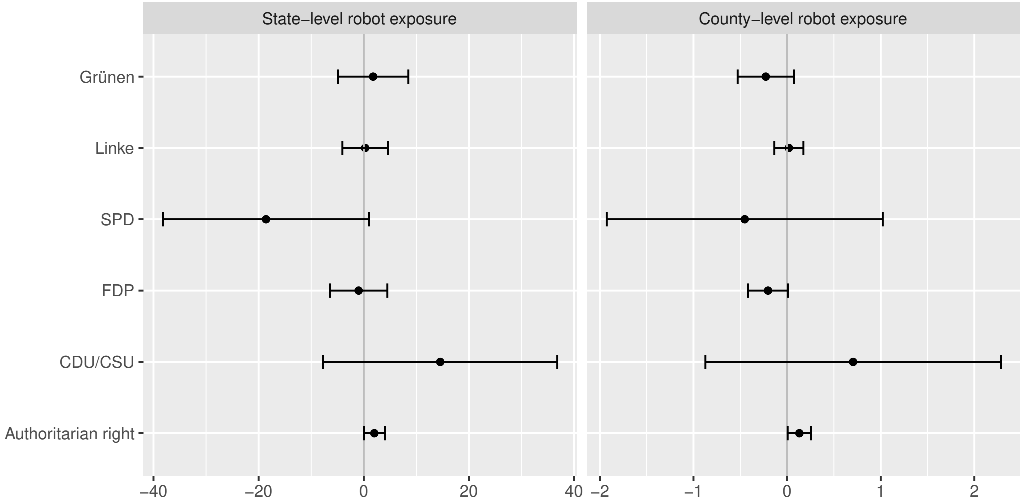

Note: The graph shows the effect of state level robot exposure on county-level vote shares (left panel) and county-level robot exposure on county-level vote share (right panel. Robot exposure defined as change in the number of robots per thousand workers in a two year window prior to the election. The left panel shows a replication of the specification of ACS which measures robot exposure at the state level (NUTS-1) and includes election fixed effects. The right panel measures robot exposure at the region level (401 *Kreise und kreisfreie Städte*, NUTS-3) and adds region fixed effects. Standard errors are clustered at the state-election level (left panel) or at the commuting-zone election level (right panel). Bars represent 95% confidence intervals.

C.5.2 Political Outcomes

In general, we can replicate their result that robotization is associated with more right-authoritarian support if we use ACS's modelling approach even though they jointly analyzed several European countries and we only have data on Germany.

Similar to ACS, we find that one standard deviation increase in robot exposure (+0.25 robots / 1000 workers) is associated with a significant increase in the vote share of right-authoritarian parties by 0.54 percentage points when we apply their statistical model to our data.[4] These results remain stable when using an instrumental variable approach (panel B of Table C.19). We thus conclude differences in observed results do not stem from the uniqueness of the German case.

[4]More generally, we find that increased exposure to robots shift party support to the right. Besides right-authoritarian parties, this modelling approach implies that Germany's Christian Democrats CDU has the largest point estimate, even though imprecisely estimated. On the other hand, the results show that according to this modelling approach, left and liberal parties lose support in affected areas.The only significant result is for leftist party *Die Linke* (see left panel of Figure C.3 and panel A of Table C.19).

Next we turn our attention to the question if differences in the geographic disaggregation of the robot exposure, the sample period or the types of elections considered could be explaining the different results. Recall that our approach uses county-level variation (NUTS-3) whereas ACS's approach uses state-level variation (NUTS-1 for Germany). Therefore, we now want to apply their modelling approach to county level rather than state-level variation in robot exposure:

$$Y_{r,t} = \beta_1(Robots_{r,t-1} - Robots_{r,t-3}) + \mu_{election} + \epsilon_{r,t} \tag{C.2}$$

We now regress the electoral outcome in a region $Y_{r,t}$ directly on the regional robot exposure measured as the difference in robot intensity in the two years prior to the elections ($Robots_{r,t-1} - Robots_{r,t-n}$). Otherwise we use the same specifications as before, namely a plane OLS (see right panel of Figure C.3 and panel A of Table C.20) and a 2SLS specification (panel B). .[5]

The general pattern remains that increases in robot exposure are associated with a shift of political support more to the right of the political spectrum. This is remarkable as we note that point estimates are reduced dramatically if we use regional instead of state variation in the measure of robot exposure.

The positive effect of robotization on right-authoritarian support also remain stable if we restrict our sample period to the one used by ACS (they only look at the years 1999-2015), or if we remove European or state elections (both not reported).

Hence, we conclude that different results are also not driven by differences in data set construction.

C.5.3 Implied Economic Outcomes

Next, we want to ask the question if the two approaches (i.e. our two-way fixed panel model approach and ACS's short-difference repeated cross-section) indeed capture the same labor market transformation. For this, we analyze the effect of robot exposure on the same economic outcomes as used previously in Section 3.5.2 but now using their modelling approach. ACS do not report own labor market results but instead refer to previous literature which showed that robotization decreased manufacturing employment. It is tacitly assumed that their approach would lead to the same results.

[5] The remaining difference is that ACS's approach uses the change in number of robots in a time window two years prior to the elections as main explanatory variable. Our approach instead directly uses the log-number of robots per thousand workers in the year of the elections. As argued before, their approach mixes levels and changes whereas our approach uses levels for LHS and RHS variables. Note that due to the region and year fixed effect, our approach is equivalent to using a difference estimator (differences on the LHS and RHS).

Figure C.4: Robots and Employment: ACS approach

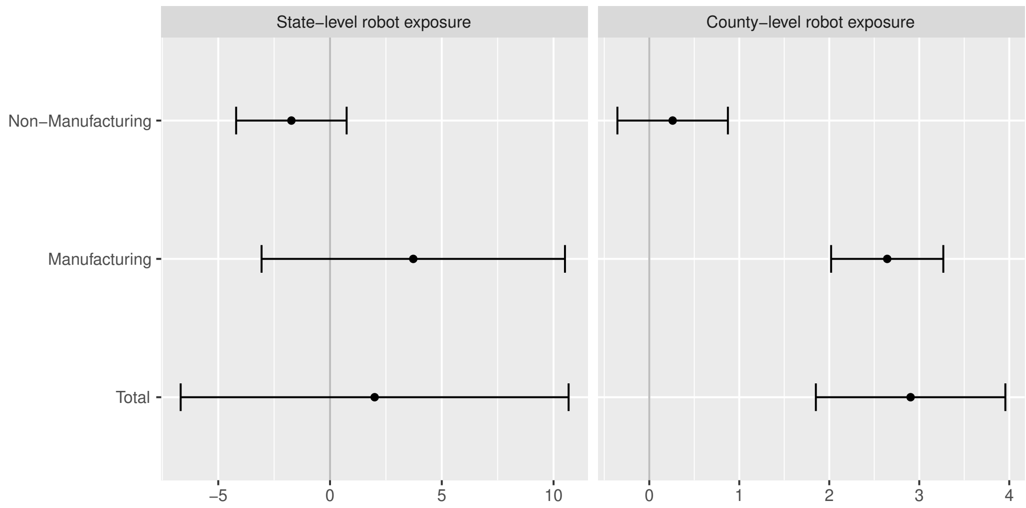

Note: The graph shows the effect of regional robot exposure on regional employment relative to population in %. Robot exposure defined as change in the number of robots per thousand workers in a two year window prior to the election. The left panel shows a replication of the specification of ACS which measures robot exposure at the state level (NUTS-1) and includes election fixed effects. The right panel measures robot exposure at the county level NUTS-3). Standard errors are clustered at the state-election level (left panel) or at the county level (right panel). Bars represent 95% confidence intervals.

However, Figure C.4 shows that their modelling approach suggests that increased robot adoption leads to *more* manufacturing employment. This holds using either the replication of their approach (left side) or using a model employing more fine-grained regional variation of robot exposure (right side). Table **??** repeats the region-level analysis using a plain OLS model (panel A) and a 2SLS model (panel B). Under all specifications, the results indicate that increased exposure to robotization is associated with a larger fraction of workers being employed in the manufacturing sector. Additionally, we analyze what changing job requirements are implied by the approach of ACS. Figure C.5 shows that the repeated short-difference approach suggests that regions adopting more robots create more manual-routine jobs and using state-level variation in robot adoption, this modelling approach does not replicate well the often described 'hollowing out of the middle class'. These results stand in contrast to what we found before (see Section 3.5.2). Furthermore, they do not square well with the hypothesized mechanisms put forward by ACS. Rather than left behind, this approach suggests that semi-skilled routine workers in manufacturing are doing well in the face of increased robotization.

More generally, these patterns also do not align well with the RBTC paradigm. Both theoretical and empirical studies on the matter agree that semi-skill routine jobs are taken over by robots and diminish in numbers if a region is more exposed to automation (Acemoglu and Restrepo, 2020; Dauth et al., 2021). Note however, that our proposed compositional story fits the economic and political results proposed by this approach. Automation affects party support mainly through changing occupational structures and regions who still harbor a large group of semi-skilled routine workers are the ones who are most supportive of conservative and authoritarian-right parties. Since this approach predicts growing manufacturing employment with routine jobs, etc. it does not come as a surprise that this approach concludes that robotization is associated with more support for right-authoritarian parties.

As mentioned before, we believe that the results stem from mixing the *change* in robot penetration with *levels* of employment shares and party vote shares in the modelling approach (see Equation C.1). However, since the number of robots grows most where there is a large manufacturing sector and routine work, this approach implicitly correlates the size of the manufacturing sector or the number of routine workers with right-authoritarian support. As we have shown in Section 3.5.2, larger shares of routine, mid-skilled manufacturing workers are associated with political support for right-authoritarian parties.

Figure C.5: Changing Job requirements: ACS approach

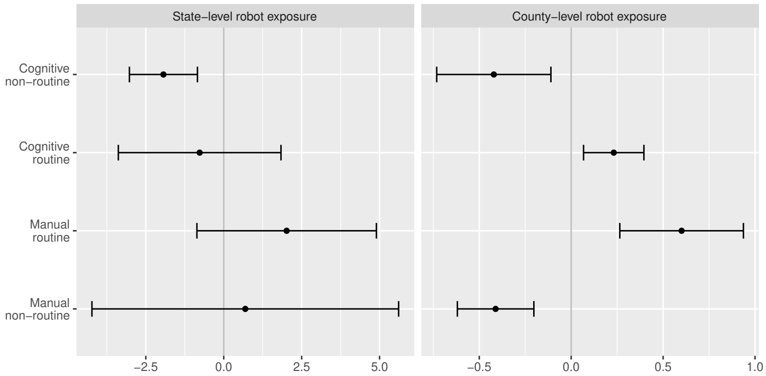

(a) Main task

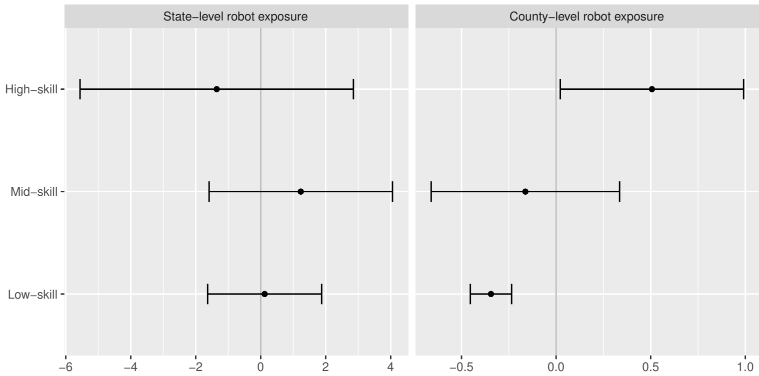

(b) Skill

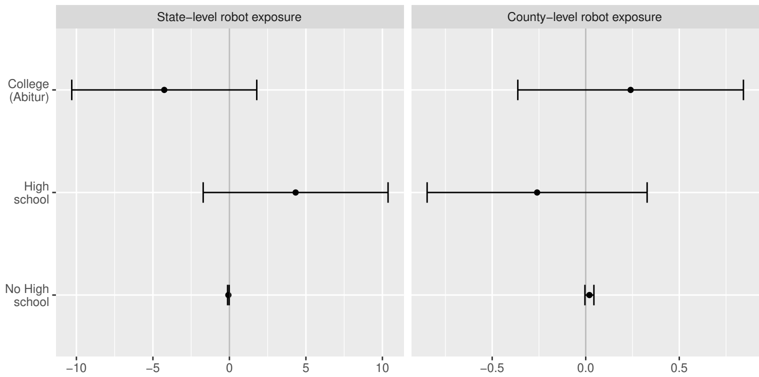

(c) Education

Note: The graph shows estimated effect of the robot exposure (change of number of robots per thousand workers over previous two years) on regional employment outcomes including year fixed effects (similar to election fixed effects used by AVS. The dependent variable in panel (a) is the main task of regional occupation composition. Panel (b) show the effect on regional jobs by skill requirement. Panel (c) shows the effect of robotization on regional employment composition by education level.).

All variables are expressed as share of regional employment in % such that coefficients sum up to zero. Black bars represent 95% confidence intervals.

C.5.4 Regression Tables

Table C.19: State-level Robot Exposure and Party Vote Shares

(A) OLS

	Dependent variable:					
	Grünen	Linke	SPD	FDP	CDU/CSU	Authoritarian right
State-Level Robot Exposure	1.784	0.267	−18.588*	−0.969	14.558	2.016**
	(3.424)	(2.207)	(9.982)	(2.790)	(11.363)	(1.015)
Election FE	X	X	X	X	X	X
Observations	1,619	1,619	1,619	1,619	1,619	1,619
R^2	0.166	0.710	0.431	0.745	0.123	0.753

(B) 2SLS

	Dependent variable:					
	Grünen	Linke	SPD	FDP	CDU/CSU	Authoritarian right
State-Level Robot Exposure	1.164	−0.578	−8.726	−1.815	7.835	1.679
	(4.770)	(2.113)	(20.367)	(3.803)	(16.756)	(1.464)
Election FE	X	X	X	X	X	X
First-stage F-stat	10.69	10.69	10.69	10.69	10.69	10.69
Observations	1,619	1,619	1,619	1,619	1,619	1,619
R^2	0.166	0.709	0.420	0.745	0.117	0.753

Note: Regressions of regional party vote share (in %) on robot exposure (change in number of robots per thousand workers over 2 years prior to the elections) measured at state level. Panel (A) shows plain OLS with election fixed effects. Panel (B) instruments robot exposure with values from other European countries. Panel (C) adds region fixed effects.
Replication of Table 1 from Anelli et al. (2019).
Standard errors reported in parenthesis are clustered at the state-election level. *p<0.1; **p<0.05; ***p<0.01

Table C.20: County-level Robot Exposure and Party Vote Shares

(A) OLS

	Dependent variable:					
	Grünen	Linke	SPD	FDP	CDU/CSU	Authoritarian right
Regional Robot Exposure	−0.228	0.019	−0.453	−0.204*	0.705	0.131**
	(0.153)	(0.079)	(0.752)	(0.109)	(0.805)	(0.064)
Election FE	X	X	X	X	X	X
Observations	1,619	1,619	1,619	1,619	1,619	1,619
R^2	0.164	0.710	0.393	0.745	0.097	0.741

(B) 2SLS

	Dependent variable:					
	Grünen	Linke	SPD	FDP	CDU/CSU	Authoritarian right
Regioanal Robot Exposure	−0.493*	0.012	0.342	−0.423**	0.340	0.173*
	(0.256)	(0.113)	(1.302)	(0.196)	(1.367)	(0.095)
Election FE	X	X	X	X	X	X
First-stage F-stat	35.04	35.04	35.04	35.04	35.04	35.04
Observations	1,619	1,619	1,619	1,619	1,619	1,619
R^2	0.162	0.710	0.392	0.744	0.097	0.741

Note: Regressions of regional party vote share (in %) on robot exposure (change in number of robots per thousand workers over 2 years prior to the elections) measured at region level. Panel (A) shows plain OLS with election fixed effects. Panel (B) instruments robot exposure with values from other European countries. Standard errors reported in parenthesis are clustered at the state-election level. *p<0.1; **p<0.05; ***p<0.01

C.6 Replication of Dauth, Findeisen, Südekum and Wössner (2021)

In this section, we replicate a set of basic results of Dauth et al. (2021) (henceforth DFSW) which is the study most closely related to the economic aspect of this paper. This serves mainly as a validation of our approach how to measure robot intensity. As explained in Section 3.4, DFSW also use IAB data to study labor market consequence of robot adoption in Germany. However, they have access to a more encompassing data set to construct their measure of robot intensity and we want to show that using our measure leads to comparable results. In contrast to our analysis, they also include Eastern Germany. We will follow their sample choice for the replication.

DFSW use a long difference approach where they analyze how increases in the robot exposure between 1994 and 2014 changed the labor market composition over the same time period. Their main findings are that while overall employment is not affected by the adoption of new robots, there are distributional consequences. Manufacturing jobs disappear but this is compensated by the creation of jobs in the non-manufacturing sector. Additionally leveraging on individual level data, they can show that incumbent

Figure C.6: Region-level exposure to robots and employment growth.

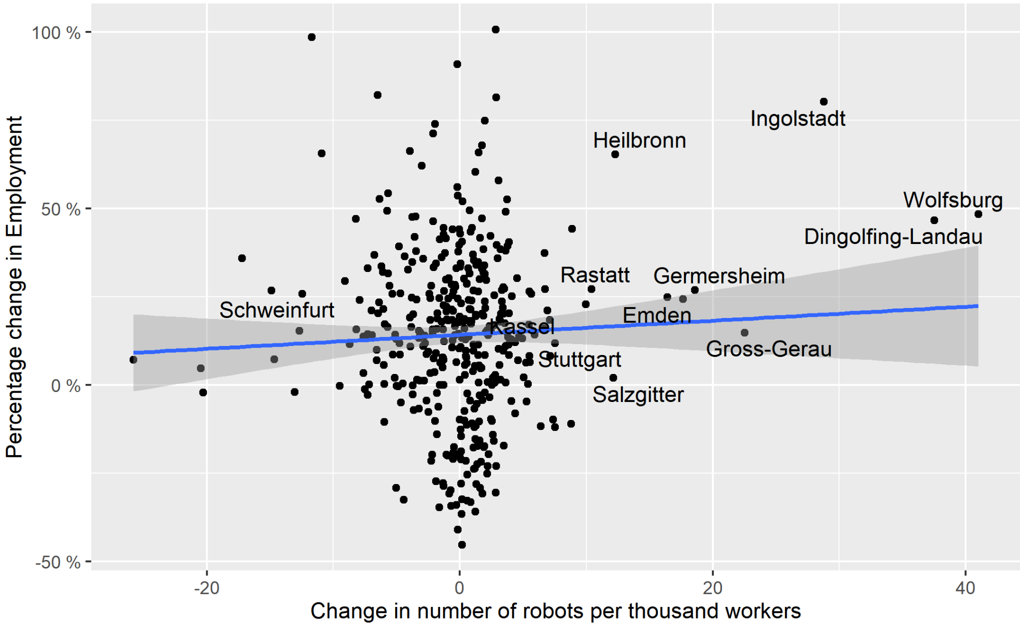

Note: The graph plots the change in estimated number of robots per thousand workers (conditional on regional employment shares in nine broad industry groups and federal state dummies) for 401 German counties (*Kreise und kreisfreie Städte*) and the growth rate of total employment between between 1994 and 2014.

workers are not displaced. New generations have to cope with changing labor demand by either obtaining a college education and to move into non-routine cognitive jobs or settling with rather precarious low skilled service sector jobs. Furthermore, they find evidence for skill polarization.

We replicate the part of their study focusing on regional employment effects of robotization. We successfully replicate the main figure and two main tables of the previous study using specifications as close as possible to DFSW. Figure C.6 plots the relationship between regional robot adoption and employment change. The x-axis shows that change in the number of robots per thousand workers (conditional on employment shares in broad industry groups and federal state dummies). The y-axis displays the regional employment growth. The correlation is slightly positive but not statistically significant. DFSW's graph shows similar results.

These (null-)findings are validated in a long-difference regression analysis where we regress a region's employment growth between 1994 and 2014 on changes of the region's robot exposure (see Table C.21). Again following DFSW, we additionally control for regional employment composition in the base year (employment shares of nine industry groups, share of high-, mid- and low-skilled workers, share of workers above fifty, share of female workers, share of foreign workers, 4 broad economic region dummies). Furthermore, we successively add changes in the region's trade exposure and changes in ICT capital stocks as addition controls.

Table C.21 which is an exact replication of Table 2 of DFSW shows that robots do not have an effect on total employment. The point estimate is always small and insignificant. The change of the number of robots per worker between 1994 and 2014 does not predict employment changes over the same time period. This effect holds controlling for a wide range of demographic characteristics of the region (column 2) and controlling for other economic shocks such as changes in trade exposure (column 3) and investments in ICT (column 4).[6] These results are very close to the ones DFSW find.

Table C.22 distinguishes between the manufacturing and the non-manufacturing sector. Again following DFSW, we use the same specifications as column 2-4 from Table C.21. However, the dependent variable is now either the growth of manufacturing employment or the growth of non-manufacturing manufacturing employment. While the effect of robots on manufacturing employment is slightly negative (column 2-4), the effect of robots on non-manufacturing employment is positive (column 5-7).

The coefficient we find are smaller in size (maybe due to higher measurement error) but the general pattern is close to DFSW. Note that we were not able to exactly reconstruct their ICT measure and instead use changes in the regional ICT capital stock per worker.[7]

Finally, we analyze how robots affect employment composition relative to the region's population. As Table C.23 shows, we also find a shift away from manufacturing employment towards non-manufacturing employment. However the overall trend is slightly more positive. The estimated effect of increasing robot exposure on manufacturing employment hovers around zero (column 2-4) while employment in the non-manufacturing sector is increasing (column 7-9).

What is interesting about this specification is that it allows us to calculate the effect a of single robot. We find that each robot affects manufacturing between -0.3 and +0.1 jobs depending on the specification. At the same time an additional robot is associated with the creation of between .9 and 1.5 non-manufacturing jobs. These numbers smaller than the results of DFSW who find that each robot replaces between 1.6-1.8 manufacturing jobs while it creates additional 1.4-1.8

[6] We use the changes in the capital stock in information technology, communication technology and software and databases normalized by employment from the EUKLEMS database. Since the time series only starts in 1995 we use the difference 1995-2014 in the long-difference approach.

[7] See Gallego et al. (2021) for a detailed description of the construction of our ICT measure. Note that EUKLEMS data only starts in 1995. Therefore, we use the difference between 1995-2014.

non-manufacturing jobs.[8]

[8] We calculate the absolute number effects of one additional robot similar to Acemoglu and Restrepo (2020):

$$\frac{Employment_2}{Population_2} = \frac{Employment_1}{Population_1} + \beta \frac{\Delta robots}{\frac{Employment_1}{1000}}$$

Assuming a constant population ($population_2 = population_1$), dividing the point estimate as it was in percentage points and rearranging:

$$Employment_2 - Employment_1 = \frac{\hat{\beta}}{100} \frac{\Delta robots}{\frac{Employment_1/Population_1}{1000}}$$

The average employment to population ratio across all regions in our base year ($Employment_{1994}/Population_{1994}$) is 0.301. Hence, each additional robot affects employment as:

$$\Delta Employment = 10 \frac{\hat{\beta}}{0.301} \Delta robots$$

Table C.21: Robot Exposure and Employment

	Dependent variable:			
	% change in total employment between 1994 and 2014			
	(1)	(2)	(3)	(4)
Δ robots per 1000 workers	0.132	0.023	0.050	−0.189
	(0.105)	(0.119)	(0.126)	(0.139)
% manufacturing	−0.217*			
	(0.123)			
% food products		2.575***	2.517***	2.518***
		(0.383)	(0.372)	(0.386)
% consumer goods		0.439	0.493	0.419
		(0.308)	(0.316)	(0.327)
% industrial goods		0.475**	0.412*	0.419**
		(0.207)	(0.215)	(0.212)
% capital goods		0.884***	0.825***	0.753***
		(0.248)	(0.257)	(0.257)
% construction		1.179***	1.116***	1.046***
		(0.307)	(0.317)	(0.322)
% services		0.260	0.252	−0.294
		(0.244)	(0.246)	(0.325)
% public sector		0.656***	0.635**	0.546**
		(0.250)	(0.250)	(0.255)
Δ net exports			0.588	0.422
			(0.407)	(0.446)
Δ ICT capital stock				6.050***
				(2.051)
Observations	401	401	401	401
R^2	0.469	0.556	0.558	0.567

Note: Replication of Table 2 from Dauth et al. (2021). Regressions of total employment growth (in %) on the change in robot exposure between 1994 and 2014. All specifications include a constant, broad region dummies indicating if the region is located in the north, west, south, or east of Germany and demographic control variables, measured in the base year 1994. The demographic control variables are the employment shares of female, foreign, age > 50, medium skilled (*fachliche Tätigkeit*), and high skilled (*komplexe Spezialistentätigkeit, hochkomplexe Expertentätigkein*) workers relative to total employment (reference category: *Helfertätigkeit*). In column 1, we control for the manufacturing share in total employment. In columns 2-4, we instead include broad industry shares to control better for regional industry patterns. Industry shares cover the percentage of workers in eight broad industry groups (agriculture (reference); food products; consumer goods; industrial goods; capital goods; construction; services; public sector) in the base year 1994. Columns 3 and 4 successively take into account the change in German net exports vis-à-vis China and Eastern Europe (in 1000 € per worker), and the change in ICT capital stock (in 1000 € per worker), both between 1994 and 2014.

Standard errors reported in parenthesis are clustered at the level of 50 commuting zones. *p<0.1; **p<0.05; ***p<0.01

Table C.22: Composition Effects - Employment Growth

	Dependent variable:						
	Employment growth %						
	Total	Manufacturing	Manufacturing	Manufacturing	Non-manufacturing	Non-manufacturing	Non-manufacturing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ robots per 1000 workers	−0.189	−0.081	−0.010	−0.173	0.373**	0.389**	0.220
	(0.139)	(0.163)	(0.174)	(0.264)	(0.185)	(0.191)	(0.212)
% manufacturing	2.518***	2.016**	1.863**	1.863**	3.012***	2.978***	2.979***
	(0.386)	(0.786)	(0.777)	(0.767)	(0.426)	(0.414)	(0.426)
% food products	0.419	−0.361	−0.217	−0.268	1.032***	1.064***	1.012***
	(0.327)	(0.672)	(0.650)	(0.632)	(0.344)	(0.345)	(0.364)
% consumer goods	0.419**	−0.291	−0.456	−0.451	1.035***	0.999***	1.004***
	(0.212)	(0.765)	(0.754)	(0.751)	(0.220)	(0.241)	(0.241)
% industrial goods	0.753***	0.437	0.283	0.233	1.117***	1.083***	1.032***
	(0.257)	(0.714)	(0.702)	(0.679)	(0.268)	(0.289)	(0.299)
% capital goods	1.046***	0.371	0.204	0.157	1.318***	1.281***	1.232***
	(0.322)	(0.903)	(0.876)	(0.866)	(0.401)	(0.429)	(0.439)
% construction	−0.294	−1.045	−1.068	−1.441*	0.541***	0.536**	0.151
	(0.325)	(0.844)	(0.827)	(0.750)	(0.209)	(0.217)	(0.344)
% services	0.546**	0.084	0.031	−0.031	0.716***	0.705***	0.641**
	(0.255)	(0.697)	(0.686)	(0.668)	(0.257)	(0.260)	(0.275)
% public sector	0.422		1.555***	1.442**		0.340	0.224
	(0.446)		(0.533)	(0.565)		(0.596)	(0.631)
Δ net exports	6.050***			4.134			4.265
	(2.051)			(3.234)			(2.927)
Observations	401	401	401	401	401	401	401
R^2	0.567	0.352	0.361	0.364	0.644	0.645	0.647

Note: Replication of Table 3 Panel A from Dauth et al. (2021). Regressions of employment growth (in %) on the change in robot exposure between 1994 and 2014 for different sectors. See Table C.21 for further details.

Standard errors reported in parenthesis are clustered at the level of 50 commuting zones. *p<0.1; **p<0.05; ***p<0.01

Table C.23: Composition Effects - Employment to Population Ratio

	Dependent variable:						
	Employment growth %						
	Total	Manufacturing	Manufacturing	Manufacturing	Non-manufacturing	Non-manufacturing	Non-manufacturing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Δ robots per 1000 workers	0.017	0.002	0.006	−0.010	0.044*	0.041*	0.027
	(0.056)	(0.027)	(0.026)	(0.033)	(0.023)	(0.023)	(0.034)
% manufacturing	0.202**	0.064	0.055	0.055	0.140**	0.147**	0.147**
	(0.095)	(0.047)	(0.046)	(0.046)	(0.071)	(0.070)	(0.070)
% food products	−0.053	−0.092**	−0.083**	−0.088***	0.047	0.040	0.036
	(0.083)	(0.036)	(0.035)	(0.034)	(0.056)	(0.055)	(0.057)
% consumer goods	−0.040	−0.069	−0.079*	−0.079*	0.031	0.038	0.039
	(0.083)	(0.048)	(0.047)	(0.046)	(0.051)	(0.052)	(0.053)
% industrial goods	0.077	0.039	0.030	0.025	0.049	0.056	0.052
	(0.085)	(0.044)	(0.043)	(0.041)	(0.053)	(0.054)	(0.056)
% capital goods	−0.096	−0.00002	−0.010	−0.015	−0.084	−0.077	−0.081
	(0.098)	(0.058)	(0.057)	(0.056)	(0.068)	(0.070)	(0.072)
% construction	−0.169**	−0.113**	−0.115**	−0.153***	0.015	0.016	−0.016
	(0.084)	(0.054)	(0.053)	(0.047)	(0.042)	(0.043)	(0.061)
% services	−0.003	−0.025	−0.028	−0.034	0.034	0.037	0.031
	(0.095)	(0.038)	(0.038)	(0.036)	(0.071)	(0.072)	(0.073)
% public sector	0.001		0.092	0.081		−0.070	−0.080
	(0.113)		(0.072)	(0.074)		(0.071)	(0.072)
Δ net exports	0.776			0.419			0.357
	(0.581)			(0.275)			(0.503)
Observations	401	401	401	401	401	401	401
R^2	0.499	0.383	0.386	0.389	0.653	0.653	0.654

Note: Replication of Table 3 Panel B from Dauth et al. (2021). Regressions of change in the employment to population ratio on the change in robot exposure between 1994 and 2014 for different sectors. See Table C.21 for further details.

Standard errors reported in parenthesis are clustered at the level of 50 commuting zones. *p<0.1; **p<0.05; ***p<0.01

Bibliography

Acemoglu, D. and Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30.

Acemoglu, D. and Restrepo, P. (2020). Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244.

Anelli, M., Colantone, I., and Stanig, P. (2019). We Were The Robots: Automation and Voting Behavior in Western Europe. *IZA Discussion Paper Series*, No. 12485.

Anelli, M., Colantone, I., and Stanig, P. (2021). Individual vulnerability to industrial robot adoption increases support for the radical right. page 45.

Antoni, M., Schmucker, A., Seth, S., and Berge, P. V. (2019). *Sample of integrated labour market biographies (SIAB) 1975-2017*. Institut für Arbeitsmarkt-und Berufsforschung (IAB), Nürnberg.

Autor, D., Dorn, D., Hanson, G., and Majlesi, K. (2016). Importing Political Polarization? *NBER Working Paper No. 22637*.

Autor, D. H. and Dorn, D. (2013). The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China Syndrome: Local Labor Market Effects of Import Competition in the United States. *American Economic Review*, 103(6):2121–2168.

Autor, D. H., Dorn, D., and Hanson, G. H. (2015). Untangling Trade and Technology: Evidence from Local Labour Markets. *The Economic Journal*, 125(584):621–646.

Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.

Barber, B. V., Beramendi, P., and Wibbels, E. (2013). The behavioral foundations of social politics: Evidence from surveys and a laboratory democracy. *Comparative Political Studies*, 46(10):1155–1189.

Barberá, P., Casas, A., Nagler, J., Egan, P. J., Bonneau, R., Jost, J. T., and Tucker, J. A. (2019). Who Leads? Who Follows? Measuring Issue Attention and Agenda Setting by Legislators and the Mass Public Using Social Media Data. *American Political Science Review*, 113(4):883–901.

Bartik, T. J. and Doeringer, P. B. (1993). Who Benefits from State and Local Economic Development Policies? *Journal of Policy Analysis and Management*, 12(2):411.

Biagi, F. and Falk, M. (2017). The impact of ICT and e-commerce on employment in Europe. *Journal of Policy Modeling*, 39(1):1–18.

Boix, C. (2019). *Democratic capitalism at the crossroads: technological change and the future of politics*. Princeton University Press, Princeton, NJ.

Bornschier, S. (2010). *Cleavage Politics and the Populist Right The New Cultural Conflict in Western Europe*. Temple University Press,U.S.

Brambor, T., Clark, W. R., and Golder, M. (2006). Understanding interaction models: Improving empirical analyses. *Political Analysis*, 14(1):63–82.

Broockman, D. E. (2014). Distorted Communication, Unequal Representation: Constituents Communicate Less to Representatives Not of Their Race. *American Journal of Political Science*, 58(2):307–321.

Broockman, D. E., Ferenstein, G., and Malhotra, N. (2019). Predispositions and the Political Behavior of American Economic Elites: Evidence from Technology Entrepreneurs. *American Journal of Political Science*, 63(1):212–233.

Buck, N. and McFall, S. (2011). Understanding Society: design overview. *Longitudinal and Life Course Studies*, 3(1):5–17.

Bush, S. S. and Prather, L. (2020). Islam, gender segregation, and political engagement: evidence from an experiment in tunisia. *Political Science Research and Methods*, pages 1–17.

Butler, D. M., Nickerson, D. W., et al. (2011). Can learning constituency opinion affect how legislators vote? results from a field experiment. *Quarterly Journal of Political Science*, 6(1):55–83.

Cagé, J., Hervé, N., and Viaud, M.-L. (2020). The production of information in an online world. *The Review of Economic Studies*, 87(5):2126–2164.

Campillo, I. (2019). 'If we stop, the world stops': the 2018 feminist strike in Spain. *Social Movement Studies*, 18(2):252–258.

Clayton, A. (2021). How do electoral gender quotas affect policy? *Annual Review of Political Science*, 24:235–252.

Colantone, I. and Stanig, P. (2018). Global competition and Brexit. *American Political Science Review*, 112(2):201–218.

Conover, M. D., Ratkiewicz, J., and Francisco, M. (2011). Political polarization on twitter. *Icwsm*, 133(26):89–96.

Cortes, G. M. (2016). Where have the middle-wage workers gone? A study of polarization using panel data. *Journal of Labor Economics*, 34(1):63–105.

Cyert, R. M. and March, J. G. (1963). *A behavioral theory of the firm*, volume 2.

Dauth, W., Findeisen, S., Südekum, J., and Wössner, N. (2017). German Robots – The Impact of Industrial Robots on Workers. Technical report, IAB Discussion Paper Nr. 30.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The Adjustment of Labor Markets to Robots. *Journal of the European Economic Association*.

de Vries, G. J., Gentile, E., Miroudot, S., and Wacker, K. M. (2020). The rise of robots and the fall of routine jobs. *Labour Economics*, 66.

Denrell, J. (2005). Why most people disapprove of Me: Experience sampling in impression formation. *Psychological Review*, 112(4):951–978.

Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. (2018). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. Technical report.

Dippel, C., Gold, R., and Heblich, S. (2015). Globalization and Its (Dis-)Content: Trade Shocks and Voting Behavior.

Doherty, D., Gerber, A. S., and Green, D. P. (2006). Personal Income and Attitudes toward Redistribution: A Study of Lottery Winners. *Political Psychology*, 27(3):441–458.

Dolan, K. (2010). The impact of gender stereotyped evaluations on support for women candidates. *Political Behavior*, 32(1):69–88.

Dolan, K. (2011). Do women and men know different things? measuring gender differences in political knowledge. *The Journal of Politics*, 73(1):97–107.

Dolan, K. A. (2014). *When does gender matter?: Women candidates and gender stereotypes in American elections*. Oxford University Press, USA.

Druckman, J. N. and Jacobs, L. R. (2006). Lumpers and splitters: The public opinion information that politicians collect and use. *International Journal of Public Opinion Quarterly*, 70(4):453–476.

Eagly, A. H., Wood, W., and Diekman, A. B. (2000). Social role theory of sex differences and similarities: A current appraisal. *The developmental social psychology of gender*, 12:174.

Ennser-Jedenastik, L. (2018). Welfare Chauvinism in Populist Radical Right Platforms: The Role of Redistributive Justice Principles: Welfare Chauvinism in Populist Radical Right Platforms: The Role of Redistributive Justice Principles. *Social Policy & Administration*, 52(1):293–314.

Fenger, M. (2018). The social policy agendas of populist radical right parties in comparative perspective. *Journal of International and Comparative Social Policy*, 34(3):188–209.

Fenno, R. F. (1977). Us house members in their constituencies: An exploration. *American Political Science Review*, 71(3):883–917.

Fiedler, K. (2012). Meta-cognitive myopia and the dilemmas of inductive-statistical inference. In *Psychology of Learning and Motivation*, volume 57, pages 1–55. Elsevier.

Frey, C. B., Berger, T., and Chen, C. (2018). Political machinery: did robots swing the 2016 US presidential election? *Oxford Review of Economic Policy*, 34(3):418–442.

Galesic, M., Olsson, H., and Rieskamp, J. (2018). A sampling model of social judgment. *Psychological Review*, 125(3):363.

Gallego, A., Kurer, T., and Schöll, N. (2021). Neither Left Behind nor Superstar: Ordinary Winners of Digitalization at the Ballot Box. *Journal of Politics*, Forthcoming.

Gay, C. (2007). Spirals of Trust? The Effect of Descriptive Representation on the Relationship between Citizens and Their Government. *American Journal of Political Science*, 46(4):717.

Genz, S., Bellmann, L., and Matthes, B. (2019). Do German Works Councils Counter or Foster the Implementation of Digital Technologies? : First Evidence from the IAB-Establishment Panel. *Jahrbucher fur Nationalokonomie und Statistik*, 239(3):523–564.

Gingrich, J. and Häusermann, S. (2015). The decline of the working-class vote, the reconfiguration of the welfare support coalition and consequences for the welfare state. *Journal of European Social Policy*, 25(1):50–75.

Goldin, C. D. and Katz, L. F. (2009). *The Race between Education and Technology*. Harvard University Press, Cambridge, MA.

Goldsmith-Pinkham, P., Sorkin, I., and Swift, H. (2020). Bartik instruments: What, when, why, and how. *American Economic Review*.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *The American Economic Review*, 104(8):2509–2526.

Graetz, G. (2020). Labor Demand in the Past, Present and Future.

Graetz, G. and Michaels, G. (2018). Robots at Work. *The Review of Economics and Statistics*, 100(5):753–768.

Grimmer, J. and Stewart, B. M. (2013). Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political analysis*, 21(3):267–297.

Hannan, M. T., Le Mens, G., Hsu, G., Kovács, B., Negro, G., Pólos, L., Pontikes, E., and Sharkey, A. J. (2019). *Concepts and categories: Foundations for sociological and cultural analysis*. Columbia University Press.

Hardy, W., Keister, R., and Lewandowski, P. (2018). Educational upgrading, structural change and the task composition of jobs in Europe. *Economics of Transition*, 26(2):201–231.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *Springer Series in Statistics*, volume 27.

Haynes, E. (1997). Women and legislative communication. *Harvard University. Typescript*.

Heath, O. (2018). Policy Alienation, Social Alienation and Working-Class Abstention in Britain, 1964–2010. *British Journal of Political Science*, 48(4):1053–1073.

Henderson, G., Hertel-Fernandez, A., Mildenberger, M., and Stokes, L. C. (2021). Conducting the heavenly chorus: Constituent contact and provoked petitioning in congress. *Perspectives on Politics*, pages 1–18.

Holland, J. H. (1992). *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press, Cambridge, Massachusetts.

Hope, D. and Martelli, A. (2019). The Transition to the Knowledge Economy, Labor Market Institutions, and Income Inequality in Advanced Democracies. *World Politics*, 71(2):236–288.

Huddy, L. and Terkildsen, N. (1993). Gender stereotypes and the perception of male and female candidates. *American journal of political science*, pages 119–147.

IFR (2016). Executive summary world robotics 2016 industrial robots.

Im, Z. J., Mayer, N., Palier, B., and Rovny, J. (2019). The "losers of automation": A reservoir of votes for the radical right? *Research & Politics*, 6(1).

Inglehart, R. (1977). *The Silent Revolution: Changing Values and Political Styles Among Western Publics*. Princeton University Press, Princeton, NJ.

Iversen, T. and Soskice, D. (2006). Electoral Institutions and the Politics of Coalitions: Why Some Democracies Redistribute More Than Others. *American Political Science Review*, 100(02):165–181.

Iversen, T. and Soskice, D. (2019). *Democracy and Prosperity. Reinventing Capitalism through a Turbulent Century.* Princeton University Press, Princeton, NJ.

Jaeger, D., Ruist, J., and Stuhler, J. (2018). Shift-Share Instruments and the Impact of Immigration. *National Bureau of Economic Research Working Paper Series*.

Jaeger, K. (2016). EU-KLEMS growth and productivity accounts 2017 release-description of methodology and general notes. *The Conference Board Europe*.

Jensen, J. B., Quinn, D. P., and Weymouth, S. (2017). Winners and Losers in International Trade: The Effects on US Presidential Voting. *International Organization*, 71(3):423–457.

Jungherr, A. (2016). Twitter use in election campaigns: A systematic literature review. *Journal of information technology & politics*, 13(1):72–91.

Jungherr, A., Rivero, G., and Gayo-Avello, D. (2020). *Retooling politics: How digital media are shaping democracy*. Cambridge University Press.

Kitschelt, H. (1994). *The transformation of European Social Democracy*. Cambridge studies in comparative politics. Cambridge University Press, Cambridge etc.

Kitschelt, H. and Rehm, P. (2014). Occupations as a Site of Political Preference Formation. *Comparative Political Studies*, 47(12).

Klenert, D., Fernandez-Macias, E., and Antón, J.-I. (2020). Do robots really destroy jobs? Evidence from Europe. Technical Report 2020/01, European Commission, Seville.

Koch, M., Manuylov, O., and Smolka, M. (2019). Robots and Firms. *CESinfo Working Paper*.

König, P. D. and Wenzelburger, G. (2018). Why parties take up digitization in their manifestos An empirical analysis of eight Western European economies. *Journal of European Public Policy*, 0(0):1–18.

Kurer, T. (2020). The Declining Middle: Occupational Change, Social Status, and the Populist Right. *Comparative Political Studies*, OnlineFirst.

Kurer, T. and Gallego, A. (2019). Distributional consequences of technological change: Worker-level evidence. *Research & Politics*, 6(1).

Kurer, T. and Palier, B. (2019). Shrinking and shouting: the political revolt of the declining middle in times of employment polarization. *Research & Politics*, 6(1).

Landis, J. R. and Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biomedics*, pages 159–174.

Lawless, J. L. (2015). Female candidates and legislators. *Annual Review of Political Science*, 18:349–366.

Le Mens, G. and Denrell, J. (2011). Rational learning and information sampling: On the “naivety” assumption in sampling explanations of judgment biases. *Psychological Review*, 118(2):379–392.

Le Mens, G., Denrell, J., Kovács, B., and Karaman, H. (2019). Information sampling, judgment, and the environment: Application to the effect of popularity on evaluations. *Topics in Cognitive Science*, 11(2):358–373.

Lengfeld, H. (2017). Die „Alternative für Deutschland“: eine Partei für Modernisierungsverlierer? *KZfSS Kölner Zeitschrift für Soziologie und Sozialpsychologie*, 69(2):209–232.

Lovenduski, J. and Norris, P. (2003). Westminster women: The politics of presence. *Political Studies*, 51(1):84–102.

Lowande, K., Ritchie, M., and Lauterbach, E. (2019). Descriptive and substantive representation in congress: Evidence from 80,000 congressional inquiries. *American Journal of Political Science*, 63(3):644–659.

Mansbridge, J. (1999). Should blacks represent blacks and women represent women? a contingent” yes”. *The Journal of politics*, 61(3):628–657.

Mansbridge, J. (2003). Rethinking representation. *American political science review*, pages 515–528.

March, J. G. and Shapira, Z. (1992). Variable risk preferences and the focus of attention. *Psychological review*, 99(1):172.

Margalit, Y. (2011). Costly Jobs: Trade-related Layoffs, Government Compensation, and Voting in U.S. Elections. *American Political Science Review*, 105(01):166–188.

Margalit, Y. (2013). Explaining Social Policy Preferences: Evidence from the Great Recession. *American Political Science Review*, 107(01):80–103.

Margalit, Y. (2019). Political Responses to Economic Shocks. *Annual Review of Political Science*, 22(1):277–295.

Marx, P. and Nguyen, C. (2016). Are the Unemployed Less Politically Involved? A Comparative Study of Internal Political Efficacy. *European Sociological Review*, 32(5):634–648.

Metaxas, P. T., Mustafaraj, E., Wong, K., Zeng, L., O'Keefe, M., and Finn, S. (2015). What do retweets indicate? results from user survey and meta-review of research. In *ICWSM*, pages 658–661.

Michaels, G., Natraj, A., and Van Reenen, J. (2014). Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *Review of Economics and Statistics*, 96(1):60–77.

Miller, W. E. and Stokes, D. E. (1963). Constituency influence in congress. *American political science review*, 57(1):45–56.

Milner, H. V. (2021). Voting for Populism in Europe: Globalization, Technological Change, and the Extreme Right. *Comparative Political Studies*, page 001041402199717.

Mokyr, J., Vickers, C., and Ziebarth, N. L. (2015). The History of Technological Anxiety and the Future of Economic Growth: Is This Time Different? *Journal of Economic Perspectives*, 29(3):31–50.

Oesch, D. (2006). *Redrawing the class map: stratification and institutions in Britain, Germany, Sweden and Switzerland*. Palgrave Macmillan, Basingstoke [England]; New York.

Oesch, D. (2008). Explaining Workers' Support for Right-Wing Populist Parties in Western Europe: Evidence from Austria, Belgium, France, Norway, and Switzerland. *International Political Science Review*, 29(3):349–373.

Oesch, D. and Rodriguez-Menes, J. (2010). Upgrading or polarization? Occupational change in Britain, Germany, Spain and Switzerland, 1990–2008. *Socio-Economic Review*, page mwq029.

O'Mahony, M. and Timmer, M. P. (2009). Output, input and productivity measures at the industry level: the EU KLEMS database. *The Economic Journal*, 119(538):374–403.

Pereira, M. M. (2021). Understanding and reducing biases in elite beliefs about the electorate. *American Political Science Review*, pages 1–17.

Phillips, A. (1995). *The politics of presence*. Oxford University Press.

Powell, W. W. and Snellman, K. (2004). The Knowledge Economy. *Annual Review of Sociology*, 30(1):199–220.

Prior, M. (2007). *Post-broadcast democracy: How media choice increases inequality in political involvement and polarizes elections*. Cambridge University Press.

Rama, J., Zanotti, L., Turnbull-Dugarte, S. J., and Santana, A. (2021). *VOX: The rise of the Spanish populist radical right*. Routledge.

Rehm, P., Hacker, J. S., and Schlesinger, M. (2012). Insecure alliances: Risk, inequality, and support for the welfare state. *American Political Science Review*, 106(02):386–406.

Rodden, J. (2019). *Why Cities Lose. The Deep Roots of the Urban-Rural Political Divide*. Basic Books, New York.

Rosenstone, S. J. (1982). Economic adversity and voter turnout. *American Journal of Political Science*, 26(1):25–46.

Soroka, S. N. and Wlezien, C. (2010). *Degrees of democracy: Politics, public opinion, and policy*. Cambridge University Press.

Stehrer, R., Bykova, A., Jäger, K., Reiter, O., and Schwarzhappel, M. (2019). Industry Level Growth and Productivity Data with Special Focus on Intangible Assets. *Report on methodologies and data construction for the EU KLEMS Release 2019*, (2018).

Stimson, J. A., Mackuen, M. B., and Erikson, R. S. (1995). Dynamic Representation. *American Political Science Review*, 89(3):543–565.

Sunstein, C. R. (2018). *# Republic: Divided democracy in the age of social media*. Princeton University Press.

Sutton, R. and Barto, A. (2018). *Reinforcement learning: An introduction*. MIT Press.

Thewissen, S. and Rueda, D. (2017). Automation and the Welfare State: Technological Change as a Determinant of Redistribution Preferences. *Comparative Political Studies*, 52(2):171–208.

Thorndike, E. L. (1927). The law of effect. *The American journal of psychology*.

Tversky, A. and Kahneman, D. (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and uncertainty*, 5(4):297–323.

Verba, S., Schlozman, K. L., and Brady, H. E. (1995). *Voice and equality: Civic voluntarism in American politics*. Harvard University Press.

Wängnerud, L. (2009). Women in parliaments: Descriptive and substantive representation. *Annual Review of Political Science*, 12:51–69.

Wlezien, C. (1995). The public as thermostat: Dynamics of preferences for spending. *American journal of political science*, pages 981–1000.

Wlezien, C. (2004). Patterns of representation: Dynamics of public preferences and policy. *The Journal of Politics*, 66(1):1–24.

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., and Brew, J. (2019). Huggingface's transformers: State-of-the-art natural language processing. *ArXiv*, abs/1910.03771.

Zhuravskaya, E., Petrova, M., and Enikolopov, R. (2020). Political Effects of the Internet and Social Media. *Anual Review of Economics*.