Public support for technology regulation: Evidence from 5 EU countries

Aina Gallego[*], Alexander Kuo[†]

August 29, 2023

Abstract

The acceleration of the fourth industrial revolution has led many experts and technological elites to demand government intervention to direct the path of technological development. This paper analyzes the bases of citizen support for policies that regulate and tax new technologies. We argue that policy preferences are shaped by core narratives about technology, principally related to the claim that technology favors economic growth and consumers (the 'pro-tech' narrative), or that it harms some workers and communities (the 'anti-tech' narrative). To examine how these core narratives affect public opinion, we develop new survey questions about six relevant policies (e.g. taxes on algorithms and robots) and experimentally manipulate the arguments provided to citizens about the consequences of such policies. We embed our experiments in large, representative samples from 5 EU countries (France, Germany, Italy, Poland, and Sweden) in which we also measure a theoretically motivated battery of objective and subjective technological risks, including risk of substitution by artificial intelligence (AI). We find that: a) there is considerable support for technology regulation; b) support is significantly reduced by arguing that regulation and taxation harms the economy and consumers; c) support is only modestly increased by appealing to the distributional consequences of technological disruption; d) objective indicators of technology risk neither predict attitudes nor moderate treatment effects. We conclude by discussing the implications of our results in light of growing concern about AI.

[*]Associate Professor, Department of Political Science, Constitutional Law and Philosophy of Law Universitat de Barcelona. E-mail: aina.gallego@ub.edu

[†]Associate Professor, Department of Politics and International Relations and Christ Church, University of Oxford . E-mail: alexander.kuo@politics.ox.ac.edu,

1 Introduction

Public debate about the impact of new technologies on society has until recently been characterized by unbridled techno-optimism. Accelerating digitalization has been one of the main goals of the Next Generation (NG) program of the European Union (EU), which states, "there is a broad consensus on the priorities for the European economic growth model, including the green and digital transitions" EC press release, March 2022). Technology entrepreneurs are known for their strong anti-regulatory sentiment (Broockman et al., 2019), underpinned by the belief that technological progress is overall beneficial. Political parties discuss mostly in positive terms about digitalization and mention cybersecurity or disinformation as possible risks of technological acceleration (König and Wenzelburger, 2019; Marenco and Seidl, 2021).

Yet this upbeat zeitgeist about the impacts of new technologies has been challenged since the eruption of large language models such as ChatGPT in late 2022.[1]On March 22, 2023, scientists and practitioners wrote in an open letter: "...we call on all AI labs to immediately pause for at least 6 months the training of AI systems more powerful than GPT-4. This pause should be public and verifiable, and include all key actors. If such a pause cannot be enacted quickly, governments should step in and institute a moratorium."[2] This call captures a cautionary regulatory sentiment that is also increasingly voiced by academics (Acemoglu, 2021).

To date, however, there is little research about where citizens stand in this debate. This paper considers the determinants of citizen views towards the regulation and taxation of new digital technologies. More specifically, we focus on policies that attempt to shape how new digital technologies are adopted in economic settings. The set of policies we consider corresponds to what others have called "technological steering" policies (Burgisser, 2023) —policies that intend to direct of slow down how new technologies are adopted in workplaces, production, or consumption settings and which are justified by

[1]Support for government regulation of new AI technologies was high among some AI developers and experts before 2023 (Zhang et al., 2021).

[2]Available at: https://futureoflife.org/open-letter/pause-giant-ai-experiments/. To date the letter has over 30,000 signatories.

the goal of softening the distributive or economic consequences of technological disruption. We exclude regulatory policies motivated by other types of reasons such as ethics (e.g. algorithmic discrimination), information quality (e.g. veracity stamps to prevent fake news), or health impacts (e.g. limits on video-game consumption for children).

This paper asks: Do citizens believe that governments should intervene in how new technologies are adopted in economic settings? If so, what explains positions on this issue? Which arguments or considerations are more likely to shape opinion? To address these questions, we identify two main types of arguments in academic and public debate in favor of and against the regulation and taxation of technology in workplaces and markets: a) arguments in favor of technological regulation that stress potential harms of and the need to protect the economic losers from technological disruption; b) arguments against regulation that focus on the negative economic effects of restraining technological change.

To examine the impact of both arguments, we provide new measurement of public support of six different regulatory policies, within a common survey experimental framework that randomizes these arguments. Our 2 x 2 design compares public support for technological regulation when given no frames about these policies; when arguments about how regulation protects economic losers from technology are made salient; when the argument that regulation reduces economic growth is made salient; and when both arguments are salient. In addition to the effects of arguments, we study how self-interest affects preferences over technological regulation. Our survey contains detailed questions about occupational risks and technology-oriented attitudes. We ask if citizens at higher objective and subjective risks of displacement by technology are more predisposed to support regulation and if they are more responsive to such arguments. For example, individuals whose jobs are more at risk from automation – or who have such concerns – should be more supportive of regulatory proposals when their impact on vulnerable workers is made salient.

We collected a wealth of new survey data gathered from February-April 2023 from large samples in five European countries (Germany, France, Italy, Poland, and Sweden). Our data measure support for six policy proposals, with a rich battery of variables to test a wide variety of occupational, attitudinal, and other demographic correlates of such policy views.

Overall, we find that arguments about the potential harm to the economy by regulation are more persuasive than arguments about the distributive implications of technological change that emphasize winners and losers. We also find that the correlation between self-interest, measured in a variety of ways, and support for technological regulation is weak, even when citizens are given information about who benefits from regulatory policies.

Our paper contributes to the growing academic literature that studies the determinants of technological change, digitalization, and AI and their labor market and political impacts. Recent empirical research has examined if citizens at risk of displacement by technology prefer redistribution, activation, basic income, or protection against change (Thewissen and Rueda, 2019; Busemeyer and Tober, 2023; Busemeyer et al., 2023; Gallego et al., 2022; Weisstanner, 2021; Gallego and Kurer, 2022). Results so far are mixed and do not explain why people exposed to technology-related labor market risk often fail to support policies that may help them insure against these risks and instead divert blame to other structural causes such as immigration or international trade (Wu, 2022; Mutz, 2021). Our study shows that the techno-optimist narrative that links technology to economic growth is much more persuasive than standard political economy arguments about the need to protect groups harmed by technological change. Individuals who are likely beneficiaries from such policies may find growth-based arguments against regulation more persuasive, and may be motivated to find alternative actors to blame for their economic decline.

We also contribute to a nascent literature that is critical of the current market-driven path of technological development (Johnson and Acemoglu, 2023). The deregulation of economic activity and the rise of digital technologies since the 1970s have resulted in a steady rise of income inequality, stagnant real incomes for broad sectors of the population, and an unprecedented concentration of economic and political power in a few giant technological firms. Reversing these trends requires government intervention to limit the power of large technological firms, steer technological development to welfareenhancing uses, and share the benefits of technological development. By examining public opinion about government regulation of technology, an understudied topic, our study sheds light on the potential popular support for policies that some scholars think would put democracies on a more inclusive path of development.

2 Technology, Regulation, and Preferences

In this section, we first summarize recent discussion on the current path of technological development and proposals for growing government regulation of technology. Then, we turn to the limited literature about what affects public support for technology regulation. We articulate hypotheses related to two types of explanations: a) standard political-economy accounts that focus on self-interest or group-based arguments about winners and losers of technological change; b) ideational or socio-tropic accounts that emphasize how narratives about technology shape attitudes.

2.1 The case for technology regulation

Daron Acemoglu and coauthors (Acemoglu and Restrepo, 2018; Acemoglu et al., 2020, 2022a,b; Acemoglu and Lensman, 2023; Acemoglu and Restrepo, 2019; Acemoglu, 2021; Acemoglu and Restrepo, 2022) provide an emerging framework and strong normative case in favor of government oversight of technology development. This research agenda was initially motivated by evidence that the third industrial revolution along with a deregulation of economic activity has produced job polarization and inequality in advanced industrial economies, since the 1970s (e.g. Goos et al., 2009, 2014; Autor and Dorn, 2013; Acemoglu and Restrepo, 2022), has produced stagnant or declining real wages for men without college degrees (Autor, 2019), and has concentrated gains and power in a few superstar firms Autor et al. (2020).[3]

[3] The trend towards growing concentration of economic power has intensified in recent years. The stock market value of the companies known as "Big Five" —Apple, Microsoft, Alphabet, Amazon and Meta —together with two new entrants —Nvidia and Tesla —has soared 60% to in 2023 $11 trillion fuelled by optimism about AI. In 2022, USA's GDP was $25 trillion.

Collectively, these findings cast doubt on the claim of a "productivity bandwagon," by which technological development creates economic growth that eventually trickles down and benefits everyone in society (Johnson and Acemoglu, 2023). Several reasons may prevent economic growth driven by technology from resulting in gains for a majority of the population, including the oligopolistic power of large companies, innovation that is excessively oriented towards automating work and substituting workers rather than raising marginal productivity and creating new tasks, or because firms do not share rents with workers.

Accordingly, economists are reevaluating canonical theoretical models in which technological development is always beneficial as excessively techno-optimistic. More recent task-based theoretical frameworks permit technological change to have net *negative* impacts on employment and wages; this occurs if new technologies that automate tasks previously done by humans do not create enough new tasks, directly or indirectly, to countervail the substitution effect (Acemoglu and Restrepo, 2018). In sum, economic analysis is shifting towards a more cautionary stance about the aggregate impact of market-driven technological change on welfare. Scenarios in which technological change do not produce net aggregate gains or shared prosperity are not just a theoretical possibility. Johnson and Acemoglu (2023) amass historical examples of technologies that *reduced* living standards for a majority of the population, including early agriculture, water- and windmills in medieval times, and the early industrial revolution. In these contexts, all gains were concentrated in a small elite comprised by less than 10% of the population.

Correspondingly, when technological development does not follow a path that is socially optimal or when it reduces aggregate welfare, there is a clear normative justification for government intervention to affect its development (Acemoglu and Lensman, 2023). Scholars in other disciplines are increasingly raising concerns about the capacity of new technologies, and specifically about AI, to undermine democracy through changes in informational environments (Boix, 2021; Jungherr, 2023). Although this is not the focus of this work, these arguments provide additional reasons to support regulation.

Public views about the regulation of technology remain poorly understood. This lacuna is critical because in democratic contexts, public opinion is an important determinant of or barrier to policy adoption. Despite the obvious importance of recent accelerated technological change to occupational stability and overall growth, we are lacking studies that analyze preferences about technology regulation policies that could have strong steering effects. The next sections provide the key theoretical arguments about policy preferences in this arena and how they might be shaped.

2.2 Technological Winners and Losers

The most frequent theoretical framework applied in the literature about technological change and political attitudes comes from standard political economy accounts where policy preferences follow from income maximizing individuals who pursue their self-interest. In this perspective, individuals consider their labor market situation and how technological change is likely to affect it in the future. They calculate the likelihood that technology-related changes will affect them positively or negatively and form preferences on public policies depending that follow from these calculations. This framework has been mainly developed to address the question of how technological change affects attitudes towards redistribution; in this section we consider how this general logic applies to preferences towards the regulation of technology.

The growing literature about technological change in political behavior takes as the main point of theoretical departure some version of the claim that occupation-based "displacement" risk shapes preferences for social policies (e.g. Rehm, 2009).[4] In this standard account, workers at higher objective labor market risk due to technological change are at some level aware of this risk and support more generous compensation policies to insure against income loss. The core empirical approach in this literature has been to code occupations or tasks of occupations that can be substituted by technology and link this risk to compensation support. In one of the first systematic analyses, Thewissen and Rueda (2019) find with observational European evidence that routine-task intensity (RTI) is cor-

[4]This approach draws on the larger literature on the role of unemployment risks and support for welfare as insurance.

related with support for compensatory redistribution. In relation to technological change, scholars have studied how individual exposure affects attitudes towards policies related to labor market risks, including redistribution, active labor market policies, and basic income (for reviews see Weisstanner, 2021; Gallego and Kurer, 2022; Busemeyer et al., 2022).

In applying this framework to policy preferences about the regulation and taxation of technology, the main theoretical expectation is that people who stand to lose from technological development – for instance because they are at risk of job displacement by automation – are more likely to support government regulation of technology. Conversely, winners from technological change should be less likely to demand government regulation than losers.

Following previous literature, we distinguish between *objective* and *subjective* technology-related risks.[5] We conceptualize objective risks as derived from demographic features that are not attitudes of the respondent. Citizens at a relatively high objective risk of technological disruption can be identified by using socio-demographic proxies of risk —individual characteristics that have been theorized by previous research in labor economics as being less likely to benefit from the adoption of new technologies, such as older age, lower education, and rural residency. More directly, they can be identified with measures of the risk estimated in their current occupations (we discuss this point in more detail below). We conceptualize subjective risk as a suite of beliefs that people will not benefit or will economically lose out from technological changes. Subjective risk can include broad concern about job substitution from technology, more specific versions of this concern such as how digital technologies allow outsourcing white-collar jobs,[6] stress about learning skills to manage technology (technostress), or distrust in large technological companies. Subjective risks are by construction attitudes and we discuss the correlation between objective and subjective risks later in the paper and in the SI.

[5] "Subjective" concern (or "subjective risk") is a key micro-foundation linking objective risk and policy preferences (see Walter (2010), for an application of this logic to the case of international trade).

[6] We use the term "globotics", which refers to the risk that one's job can be performed by a worker living in another country with lower salaries and integrated in the workplace through digital technologies (Baldwin, 2019).

We hypothesize the following:

H1a: Individuals at higher objective risk of being negatively affected by technological change are more likely to support government regulation of technology.

H1b: Individuals at higher subjective risk of being negatively affected by technological change are more likely to support government regulation of technology.

While the standard framework is theoretically straightforward, one caveat is that people at high objective risk of displacement may not be aware of such risks, and thus may fail to develop preferences about policies aligned with their self-interest. There is growing evidence that people are generally unaware of technology-related labor-market risks and often misattribute these types of risks to other factors such as international trade and immigration (Wu, 2022, 2023; Mutz, 2021). Empirical research about the question if workers at high objective risk of being displaced by automation support various compensatory policies has found mixed results, which may depend on the sample in question, estimation procedures and controls, and the form of redistribution policy (Gallego and Kurer, 2022; Weisstanner, 2021). For example, recent evidence indicates that workers concerned about automation are *less* likely to support active labor market policies (Busemeyer et al., 2023; Busemeyer and Sahm, 2022; Kurer and Hausermann, 2022), or universal basic income (Weisstanner, 2021). As mentioned above, one possible explanation is that people at a high objective risk of substitution are unaware of that risk. Empirical research finds very small correlations between objective and subjective risks (Gallego et al., 2022).

A complementary and perhaps more realistic logic is that when people are given direct information about the groups that stand to win and to lose from technology and corresponding regulation thereof, they may be better able to align their preferences with their self-interest. Thus, we straightforwardly hypothesize the following:

H2: Individuals at higher objective and subjective risks from technology are more supportive of technological regulation when policy beneficiaries are made salient.

We flag that previous empirical work also casts doubt on this hypothesis. Experimental findings indicate that *even when they are informed about technology-related labor market risks*, people often fail to change their preferences (Zhang et al., 2021; Jeffrey, 2021). For this reason, we also consider an additional argument related to the political economy theoretical tradition but more anchored in psychologically realistic models of political behaviour as those advocated by Achen and Bartels (2017). Attitudes toward the regulation of technology can follow from considerations about the types of people who are likely to benefit or lose out from technological change, and the attitudes that respondents have towards these groups, such as considering them deserving or not, or liking these groups or not. Recent research demonstrated that appeals to social groups can shape political behavior, independently on whether people belong to these groups or not (Thau, 2021; Huber, 2022). This mechanism only requires having general attitudes towards the groups being considered and could hold even if respondents do not directly connect self-interest to preferences. Individuals might be motivated to shelter particular groups from the negative effects of technological change (independent of the individual's own risk), simply because people may wish to protect vulnerable or liked groups. For instance, if a particular technology constitutes a labor market risk for older workers, attitudes related to regulation could follow from considerations about the harms to this type of workers and whether the person likes this group or not. Thus, we hypothesize that:

H3: Making salient individuals or communities that are harmed by technological change increases overall support for government regulation of technology.

2.3 The Core Techno-optimist Narrative

The findings discussed above suggest that self-interest may have a limited capacity to shape preferences towards technology regulation, even when citizens are directly aware or informed of who benefits from regulation. Because of these empirically grounded doubts about the capacity of *self-interest oriented* standard political economy theories to explain attitudes towards regulation, we turn to an alternative explanation. We call it the "core techno-optimist narrative" and base our argument on ideational and socio-tropic accounts.

This alternative hypothesis is partially motivated by Johnson and Acemoglu (2023) who identify the power of ideas as a critical factor that explains population acquiescence with particular paths of development. Persuasive narratives are often promulgated by technology and political elites and justify why a particular path of development is chosen by appealing to the common good. Acemoglu and Johnson pithily summarize the appeal of egotropic arguments in policy framing: "if you tell others to follow what is blatantly good for you, they will balk, seeing it as a crude attempt to get what you want. For an idea to be successful, you need to articulate a broader viewpoint that transcends your interests or, at the very least, appears to do so" (Johnson and Acemoglu, 2023, 81). Socio-tropic narratives justifying technological disruption have historically been common, and were prevalent even prior to democratization (when there was more coercive capacity or technologies where themselves coercive). For instance, during the early industrial revolution, narratives that extolled the inventiveness and entrepreneurial spirit of engineers proliferated (Johnson and Acemoglu, 2023). Persuasion through ideas should be even more crucial in democratic societies, where there are more constraints to coercion.

The claim that core narratives about technology that connect innovation to a common good are strongly capable of shaping attitudes towards the regulation of technology relates to a large body of work in political behavior that finds that voters are generally more responsive to sociotropic concerns about the national economy than to egotropic concerns based on their personal economic situation (Lewis-Beck and Paldam, 2000). This theoretical claim is also well-aligned with established theories in social psychology (Jost et al., 2004) and behavioral economics (Benabou and Tirole, 2006) about how system-justifying narratives produce conformity and legitimate inequality. More generally, it is also consistent with recent interest about the causal power of ideas and narratives (Djourelova, 2023) as well as with classical work about framing effects (Chong and Druckman, 2007).

The contemporary version of a system-justifying narrative that protects technological elites from government intervention through appeals to the common good takes the following form: technological change is the key driver of economic growth; the gains from such growth will eventually benefit or trickle down to the whole population; by potentially halting or slowing down innovation, government regulation of new technology reduces economic growth; technological elites should be trusted to choose the path of technological development that is best for societies. Such arguments are commonly made by such elites defending creative destruction, as exemplified by Facebook's initial motto of "move fast and break things." Work in political science documents that technological elites are distinctively skeptical of government regulation (Broockman et al., 2019) and that delaying the adoption and enforcement of regulation has been a key component in the strategy of gig economy firms such as Uber (Mazur and Serafin, 2023). This core narrative is supported by policy-makers when they point to digitalization as the main driver of economic growth.

We expect that appeals to the core techno-optimist narrative that links technology to economic growth may be more relevant to preference formation about the regulation of technology than redistributive appeals to winners and losers. To date, however, empirical studies about the persuasive power of this core narrative about technology in current democracies are lacking, particularly vis-a-vis other self-interested political economy arguments. This paper attempts to fill this gap. Our final hypothesis is thus:

H4: Making salient the potential negative impact of technological regulation on economic growth reduces support for such policies.

3 Data and Research Design

To test our hypotheses about how self-interest, group appeals, and sociotropic narratives shape preferences for the regulation of technology, we fielded online surveys in spring 2023 (from late February until April 2023) in five EU countries (France, Germany, Italy, Poland, and Sweden), with a sample size of 3,500 per country.The design was preregistered and deposited in a repository.[7] We first discuss how we measure preferences regarding technological regulation and then explain our experimental design.

[7]See **insert link here**. The surveys were fielded by Respondi-Bilendi. The sample was stratified by population quotas of female vs. male gender, five standard age groups, education (university attendance versus not), and NUTS-1 or broad geographical regions.

3.1 Measuring support for technological regulation

Support for the regulation of technology is a concept by postulation (Saris and Gallhofer, 2014). Rather than asking citizens about support for regulation in the abstract, our approach is to identify a set of specific policies, and ask respondents if they support such policies. Based on a review of current debates in many post-industrial countries, we develop six questions that cover different spheres. Three of our questions refer to policies explicitly designed to protect workers from job displacement due to automation. First, we consider support for taxes on companies that use algorithms to replace workers. This is an operationalization in a survey context of proposals about 'robot taxes' from prominent individuals such as Bill Gates. This policy is especially relevant regarding AI applications, and in this policy (see design section) we mention white and blue-collar jobs. Second, we measure support for strengthening labor institutions (such as unions) to have more power regarding the integration of technology in the workplace. Previous literature finds that such institutions can greatly affect how technology is incorporated and thus affect employment and wages of existing workers (for a rich study of the case of robots in Germany see Dauth et al. (2021). Third, we measure support for regulations to make it difficult for firms to replace workers with machines or technology.

Our fourth measure assesses support for a proposal to increase stricter regulations on firms that use such monitoring technologies. An increasing source of concern is the capacity of companies to use technology to monitor workers more carefully in unprecedented ways. Examples of monitoring include tracking what workers do in their work hours or indicate the appropriate action in each moment. These policies reduce worker autonomy and may significantly harm job satisfaction and mental health. We thus think it normatively important to measure citizen support for policies to regulate monitoring.

Finally, we consider two policies that more directly affect consumers. The first one is the regulation of the platform economy, which of course is both a worker-protection issue. The second one is related to consumer services and the issues include job protections for such workers and fairness towards incumbent companies (such as taxis), largely trading off against price competitiveness and convenience for consumers. The sixth policy is taxation of large Internet companies such as Amazon; such firms use algorithms and technological practices that permit both price competition against incumbents and consumer benefits. Taxation or punitive measures can be viewed as a general policy of reducing technological 'incursion' into everyday consumer affairs and as a way of preserving existing communities (e.g. shops in town centers).[8]

Specifically, we ask whether individuals agree or disagree with these six policies (wordings below). While these policies are not exhaustive, they capture a wide range of regulatory and taxation policies that have been proposed in some capacity, with different instruments, and potential beneficiaries or costs.

3.2 Experimental design

Our control group asks about support for these six policies and provides a baseline measurement of support. Digitalization is currently not a salient political issue,[9] for most policies and in most countries. We are aware that for many people, we may be measuring non-crystalized attitudes (Converse, 1964). Thus, we interpret responses to the baseline condition as general sentiment towards technological regulation in a context in which this issue has low political salience.

The treatment conditions allow us to examine how attitudes change when more information about the implications of technology regulation is given, as would be the case if the issue became more salient. We compare policy support in the control condition with support in three experimental treatments. In the first treatment, we tell respondents that technological change can harm some citizens or communities and that the policy can protect these groups. We call this condition "Winners and Losers" or "Distributional" argument. In the second treatment, we use some policy-specific version of the argument that technology produces growth and conversely that regulation harms the economy or consumers. We call this condition the "Regulation harms growth" or "Economic growth"

[8] Amazon in particular has been criticized as well for issues related to job security and monitoring of workers.

[9] Exceptions to this include specific debates about the role of platform economies, which affect more salient incumbent industries (such as transportation and tourism), and the role of large technological and Internet companies in terms of being potentially unfairly competitive.

argument. Our final treatment condition combines these two claims and we refer to this condition as "Both". Summarizing, the conditions are:

A) Control (no further statement after policy proposal)

B) Arguments about winners and losers highlighting distributive consequences and how regulation of technology can benefit particular groups of citizens

C) Arguments about the negative socio-tropic economic consequences or regulation highlighting how the regulation of technology harms the economy or consumers

D) Both arguments B and C

Table 1 displays the wording in the control group and the treatments for each policy proposal. The first column provides the question asked in the control group, with no additional information. The second column provides the additional wording mentioning specific groups as winners or losers of the policy. For respondents who belong to the groups mentioned, this allows us to test self-interest based arguments. The third column provides, for each policy, arguments that connect to the core techno-optimist narrative and claims in various forms that technology regulation harms the economy or consumers. Response options were “strongly disagree,” “disagree,” “Neither agree nor disagree,” “agree,” and “strongly agree.”

We slightly adjust some of the specific beneficiaries and the sociotropic costs of regulation across the treatment groups, to account for realistic variations in impact across the policies. We discuss here the justification for the treatment statements we chose. For the three policy proposals that aim to protect workers, we explicitly mention beneficiaries as being groups that are relatively clearly labeled as such by labor economists. For both policy questions 1-2, the beneficiaries of the policy are “. . . workers whose jobs are more threatened by technology, such as the older and less educated.” For policy question 3 (taxing robots and arguments), as the policy question explicitly gives examples of white and blue-collar jobs that are threatened by algorithms and software, the wording frames the policy as “workers that currently perform tasks such as translators, accountants or sales people.” Regarding the treatment about costs of regulation, we directly mention economic growth when this is plausible given the content of the policy, as in the two first

Table 1: Question wording and treatments for the six policy areas

Nr.	Policy	Winners and losers	Regulation harms growth
1	*The government should give trade unions or workers more power to decide whether new technologies are adopted at work and how they are implemented.*	*This policy could protect workers whose jobs are more threatened by technology, such as the older and less educated.*	*This policy could reduce economic growth and make [COUNTRY] less economically competitive compared to other countries.*
2	*The government should make it harder for firms to adopt new technologies or machines, if they reduce salaries or jobs.*	*This policy could protect workers whose jobs are more threatened by technology, such as the older and less educated.*	*This policy could reduce economic growth and make [COUNTRY] less economically competitive compared to other countries.*
3	*The government should increase taxes and regulations on firms that adopt software, robots, or algorithms that do the work that their workers do (for instance text translation, accountants, checkout machines, and customer service chats).*	*This policy could protect workers that currently perform tasks such as translators, accountants or sales people.*	*This policy could raise the prices of these services.*
4	*The government should adopt more regulations about how companies use digital technologies to monitor what people do at work.*	*to protect workers in these companies.*	*even if this reduces service quality for customers or clients.*
5	*The government should more strongly regulate "platform" companies (like Uber, Airbnb or Deliveroo).*	*to protect workers in these companies or workers in the competing sectors.*	*even if this increases prices for consumers.*
6	*The government should increase taxes on larger Internet retailers like Amazon.*	*protect smaller businesses that compete with these retailers.*	*if this could raise the prices of goods sold online.*

policies. For the third labor-protecting policy, the cost is framed as increasing the prices of these services (which is a plausible direct, observable cost from this policy).

For policy proposal 4 on worker monitoring, which captures a non-economic form of labor protection, the "winners and losers" condition straightforwardly lists the beneficiaries as workers who are being monitored. The "regulation harms the economy" condition is framed as reducing service quality for customers or clients, a frequent justification given for worker monitoring by employers.

For the last two policies that have clearer consumer implications, taxing large Internet retailers and the platform economy (though they have consequences for worker welfare as well), the beneficiaries are workers in the relevant sectors. For question 5 (gig economy) on regulation of the platform economy, beneficiaries are listed as "workers in these companies or workers in the competing sectors,"; for question 6 (online retail) on taxation of large Internet retailers, beneficiaries are the "smaller businesses that compete with these retailers." For both of these policies, the most obvious costs were higher prices of the provided goods.

The order in which the respondent read the six policies was randomized. Importantly, we maintained the same "treatment category" for all six policies to minimize contamination across question consideration. For example, if a respondent was assigned to the control group, then she read the "control-group" wording of all six policy questions.

3.3 Measurement: risks, correlates, and moderators

In this part of the design section, we present the rich set of additional variables collected, focusing on objective and subjective variables that are expected to both positively correlate with support for the protectionist policies in the control group (hypothesis 1), and to moderate the treatment effect of the distributional frame (hypothesis 2). Unless otherwise specified, all expectations discussed after the variable descriptions refer to the control group. We present the following five categories of variables.

Demographic data. We collect a rich set of demographic variables. These include respondent gender (male, female, other), region of residence (NUTS 2 and 3 level), age, highest education completed, current labor-market situation, labor-contract, after-tax income, native of country, type of area of residence (small town, city, suburb, country village, country home). Of these, we consider lower education (non-university), higher age still in the labor force (45-64), less secure labor contract, and residence in the countryside or a small town as objective indicators of risk of displacement by automation. We expect these demographic indicators to be correlated with support for worker protection policies (hypothesis 1a, applied to policies 1-3), and to be more sensitive to treatment conditions emphasizing worker protection.[10]

Subjective attitudes related to technology-specific risks. To test hypotheses about the role of subjective concerns in shaping policy preferences, we collected much information on attitudes about technology. In some cases, these risks were intentionally measured to assess their correlation with the relevant policy choice as described above, and are applicable to those in the workforce. Such measures include: higher subjective concern

[10]We do not have such clear expectations for the other policies (policies 4-6), but we collect the indicators that are more relevant to the groups mentioned in these questions: workers in the gig economy, own experience of workplace monitoring, and use of online retail.

about technology in the workplace, belief about the percent of workplace tasks that could be automated in the near future (0-100%), concern about being monitored by technology at the workplace, concern about their job being substituted by technology abroad, overall sentiment towards the five most salient technology companies, and finally, stress about learning new technology at work ("technostress"). We expect that these subjective concerns should correlate with support for all six regulatory policies.

Respondent economic situation. We also collect data on other relevant economic concerns; we measure concern about job loss and belief in the ability to find a new job if needed, and job precariousness. We expect that concern about job loss positively correlates with support for regulatory policies.

Political attitudes. We measure respondent identification with their country, occupation and gender, trust in economic and political institutions, recent national turnout, policy attitudes (redistribution and others), and vote intention in next national election. Left-wing ideology, proxied by left-wing party vote intention or by responses to redistributive policy items, should correlate with support for all six policies.

Occupational and consumer behaviors. We also measure a battery of occupation and consumer-specific behaviors, for completeness, and in the expectation that they may correlate with relevant policies. On the worker side, we measure if the respondent attains any income from platform work (excluding online surveys), and if so, what percent of income is derived from such work. We additionally assess how much technology the respondent uses for her job. On the consumer side, we measure use of apps as a consumer and use of online retailing.

'Objective' occupational indicators based on occupation and tasks. For those in the workforce, we use the inputted 4-digit ISCO occupation code to calculate using a crosswalk the imputed 'AI' risk in terms of probability that the job can be substituted with AI (calculated from occupation data). We collect various indicators of AI risk: Webb (2019), Brynjolfsson et al. (2018), Felten et al. (2021). In addition, we use the occupation codes to calculate routine-task intensity (RTI), a standard occupation-based predictor of automation-based substitution risk (Autor et al., 2003). We expect that these objective

measures of risk should correlate with support for all six policies, with stronger positive correlations for the first three worker protection policies.

4 Results

This section presents the findings most relevant to test the hypotheses while additional results are presented in the Appendix. The first subsection describes the data. Hypotheses 1a and 1b are correlational, and we test them focusing on respondents in the control group. We use the 2x2 experimental design to test hypotheses 2, 3, and 4. The distributional frames are designed to allow a direct test of hypothesis 3 about the impact of providing arguments about winners and losers of regulatory policies, regardless of the respondents' status and risks. When combined with information about the individual characteristics of respondents, this treatment condition also allows us to test hypothesis 2 about how arguments about which particular groups stand to win from a regulatory policy makes the beneficiaries of the policy more likely to support it. Finally, the economic growth condition allows a direct test of hypothesis 4 about the impact of narratives that link technology to economic growth and the regulation of technology to negative sociotropic consequences.

4.1 Baseline support for technological regulation

We first discuss baseline support for the six policies across the five countries, using only respondents in the control group, for whom no additional information was provided about the implications of these policies. Figure 1 shows that, for all six regulatory policies, there is moderate support in the population in all countries.

Figure 1: Distribution of responses in the control condition

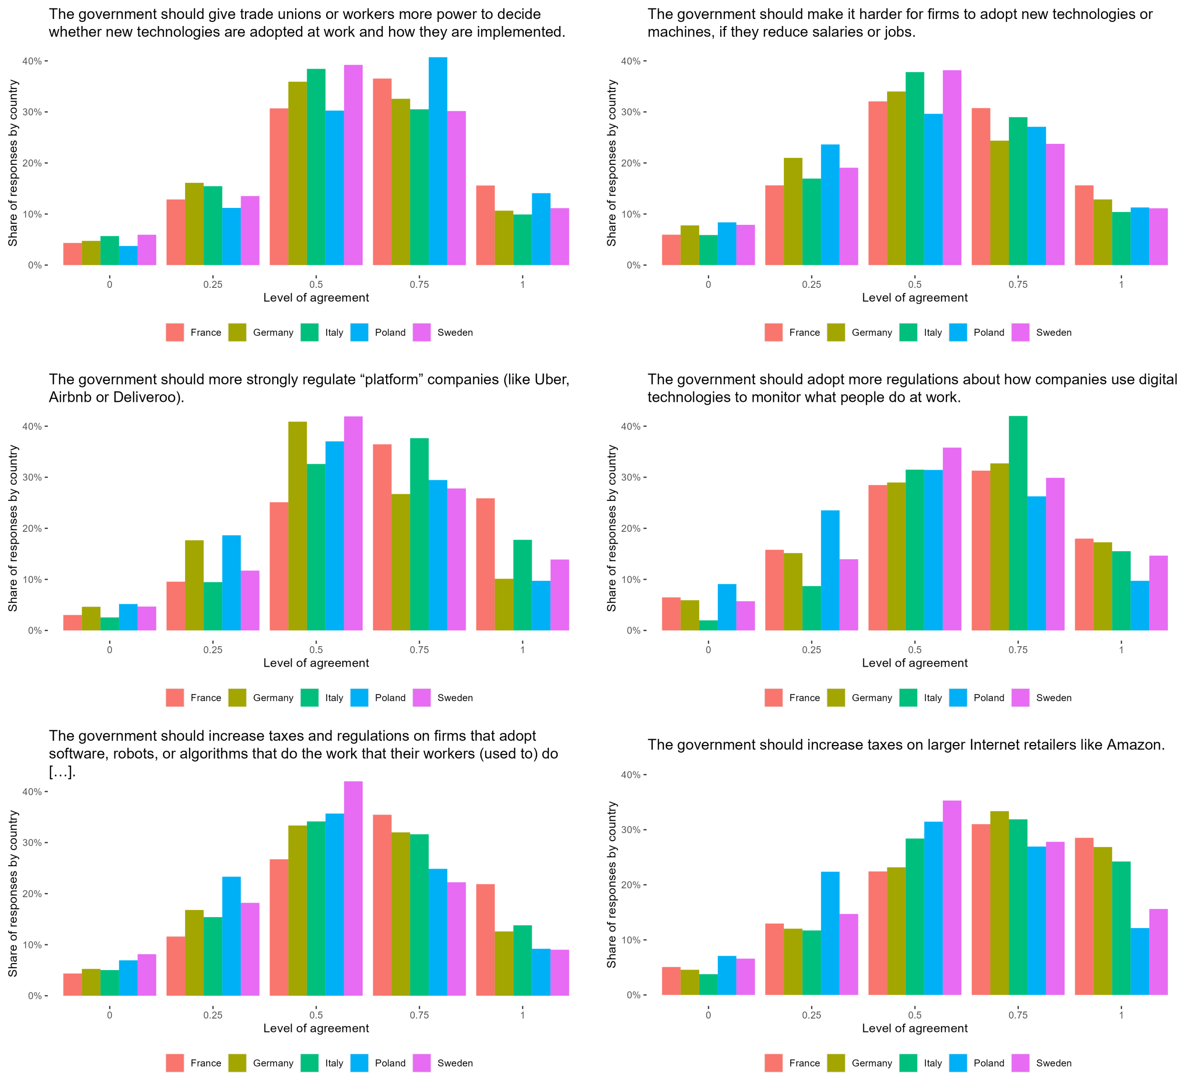

Note: The figure reports the distribution of responses to the six policy questions about technological steering for all five countries. Only respondents in the control condition are included.

Only a minority disagrees with such policies. A plurality in all countries neither agrees nor disagrees with the policies, consistent with the possibility that attitudes are not crystallized. But among those who state an opinion, support for regulation is more likely than opposition. Taxes on online retailers is the policy that commands more support, except in Poland. While there are some cross-country differences, overall, results are notably similar across the five countries.

4.2 Winners and losers, and support for technology regulation

To test hypotheses 1a and 1b, we examine if our objective and subjective measures of risks correlate with support for technology regulation.

4.2.1 Objective measures of risk

First, to test hypothesis 1a, we present the results of simple regressions of preferences towards regulation on objective indicators of risk of displacement due to automation:

$$Support\ for\ regulation_{id} = \alpha + \beta \mathbf{X}_{id} + \eta_d Country_d + \epsilon_{id} \tag{1}$$

where support for a regulatory policy by individual i in country d is a function of a vector $\mathbf{X_i}$ of socio-demographic factors (age, education, place of residence). $Region_d$ denote country fixed effects.

We focus on respondents in the control condition, who were not given any information about the implications of the regulatory policies. As discussed, one approach to measuring objective labor market risks related to automation is based on socio-demographic variables while a different one measures risk of substitution based on occupational codes. Table 2 first presents the results of the correlational analyses using socio-demographic indicators (age, education and place of residence) as objective measures that proxy for risk of substitution. Unless otherwise noted, all results are based on linear models with variables rescaled 0-1 to ease interpretation or entered as categories with the 'lowest' or intuitive category as a baseline.

Table 2: Correlates of support. Objective measures of risk in the control condition

	Unions	Governments	Gig economy	Monitoring	Amazon	Algorithms tax
Age: 45-64	-0.035***	0.569***	0.558***	0.590***	0.049***	0.540***
	(0.008)	(0.010)	(0.009)	(0.010)	(0.009)	(0.010)
Age: 65+	-0.069***	-0.009	0.055***	-0.002	0.098***	0.013
	(0.011)	(0.009)	(0.008)	(0.009)	(0.012)	(0.009)
Education: Vocational	0.024**	-0.054***	0.093***	-0.000	0.048***	0.026*
	(0.009)	(0.012)	(0.011)	(0.012)	(0.010)	(0.012)
Education: University	-0.008	0.017	0.075***	0.049***	0.050***	0.066***
	(0.010)	(0.010)	(0.009)	(0.010)	(0.011)	(0.010)
Education: Post-graduate	-0.035	-0.026*	0.035***	0.009	0.101**	0.016
	(0.029)	(0.012)	(0.010)	(0.012)	(0.033)	(0.011)
Place: A country village	-0.013	-0.076 *	0.142***	0.053	0.022	-0.005
	(0.011)	(0.034)	(0.030)	(0.033)	(0.013)	(0.032)
Place: Countryside	-0.005	0.007	0.008	0.007	-0.072***	0.022
	(0.017)	(0.013)	(0.012)	(0.013)	(0.019)	(0.013)
Place: A town or a small city	-0.010	-0.010	-0.047**	-0.063***	0.017	-0.025
	(0.009)	(0.019)	(0.017)	(0.019)	(0.011)	(0.019)
Place: The suburbs of big city	-0.012	0.007	-0.003	0.009	0.032*	0.001
	(0.013)	(0.011)	(0.010)	(0.011)	(0.014)	(0.011)
(Intercept)	0.647***	0.014	0.005	0.025	0.570***	0.024
	(0.009)	(0.015)	(0.013)	(0.015)	(0.010)	(0.014)

Note: The table presents the results of regressing support for technology regulation with objective variables that are proxies of risk by substitution. Only respondents in the control condition are included.

We find that older people at a working age (those aged 45 to 64) are more favorable to regulation than younger people for five of the six policies. The only exception is opinion on whether trade unions should be given more say over the adoption of technology, a policy slightly more supported among younger people. The differences between younger respondents and the oldest age group (aged 65 and more) are smaller than those between younger and older but working-age respondents. The only exception to this general pattern is taxes on internet retailers, a policy that is most supported among citizens aged 65 and older. The more common curvilinear pattern in which support for regulation is highest among age group 45-64 is aligned with a self-interest interpretation in which labor market considerations drive preferences for regulation.

By contrast, we find less definitive patterns regarding education. Respondents with a university education are more likely than less educated citizens to support the regulation of technology, even though it is the less educated who are generally considered as having been lagged behind by technological development in the last decades.

Similarly, we do not find that people living in big cities, who are generally considered as beneficiaries of technological change, exhibit lower support for regulation than people who live outside large cities, as would be expected from a self-interest perspective.

Next, we focus on the role of AI-risk, which is the most up-to-date measure of risk vulnerability in occupations. To the best of our knowledge, this paper is the first to link a recent measure of risk of substitution by AI proposed by Webb (2019) to survey data in political science. We find that most respondents are in occupations in which risk of substitution by AI is relatively limited, as shown in Appendix A. Out of space constraints, this table is not shown, but we find a weak correlation between this measure of objective risk and subjective risk (such as perceptions of technology) once basic controls are accounted for. Our overall takeaway is that many objective occupational measures of risk do not correlate with support for technology regulation, while older working-age, a much coarser measure of overall risk, seems to play a stronger role.

4.2.2 Subjective measures of risk

We now turn to testing hypothesis 1b and examine whether attitudes towards regulation are correlated with various subjective concerns about the labor market implications of technological change, or other attitudes broadly related to this topic. Again, in these correlational analyses, we only include respondents in the control condition. We examine a wide variety of attitudes, capturing different aspects of subjective risk.

Our first two subjective risk indicators are specifically about labor market risks and have been used in previous work, such as concern about the overall impact of technology at work on job prospects (Gallego et al., 2022) and the percentage of tasks that can be substituted by new technologies as estimated by the respondent (Kurer and Hausermann, 2022). Our third measure asks if respondents are concerned about the possibility that their job can be performed by telecommuting workers in other countries with lower salaries (globotics). We also measure and test for the role for three related indicators of techno-stress used in the management literature, as such stress should affect policy views (Ayyagari et al., 2011): how frequently they need to learn how to use new technologies in their jobs, whether respondents are concerned about being outperformed or replaced with workers with better or newer digital skills, whether they find it stressful to learn new digital skills. Our seventh measure is general sentiment about the Big Five tech firms (Apple, Alphabet, Microsoft, Meta and Amazon). Our eighth measure is based on self-reports about the intensity of use of new technologies at work. The two final measures are not specifically about technology but give a sense about respondents' concern about losing their current job in the next five years and the perceived outside options, as measured by the believe that they could find a similar or better job if they lost the current one.

We run separate regression models for each subjective variable and each policy:

$$Support\ for\ regulation_{id} = \alpha + \beta \mathbf{X}_{id} + \gamma Subjective\ attitude_{id} + \eta_d Country_d + \epsilon_{id} \quad (2)$$

where support for a regulatory policy by individual i in country d is a function of the vector $\mathbf{X_i}$ of socio-demographic factors included in Table 2 (age, education, place of residence)and one subjective variable of interest at a time. Again, $Region_d$ denote country fixed effects. Thus, we do not include control for other subjective risk variables.

Table 3 presents the results. For brevity we only display the coefficient (γ) of the relevant risk variable. Two of the subjective-risk variables, belief in technology at work as positive and attitudes towards the Big Five tech firms, are coded such that higher values indicate less 'risk,'; the rest are coded such that higher values proxy for greater risk.

We first begin by analysing results of the two most established indicators of subjective technological risk in the literature. In both cases, we only find partial support for our expectations. People who are optimistic about the impact of technology at work are less likely to support technology regulation, but only for two of the six measures examined. Similarly, people who estimate that a higher percentage of tasks in their occupation can be substituted by technology are actually less likely to support regulation for two out of six measures, and this risk does not correlate with support for the other four policies. Where the estimated coefficients are in the expected directions, but they are weak and frequently do not reach conventional levels of statistical significance.

In the case of globotics, we find that people who are concerned that their job can be performed by telecommuting workers in other countries for cheaper salaries are more

Table 3: Correlates of support. Subjective measures of risk in the control condition

	Unions	Governments	Gig economy	Monitoring	Amazon	Algorithms tax
Impact of technology at work	0.029	-0.082 **	-0.042	-0.019	-0.009	-0.132 ***
is positive	(0.024)	(0.028)	(0.024)	(0.028)	(0.027)	(0.027)
% of tasks that can be	-0.020	-0.075 **	-0.028	-0.000	-0.052 *	-0.038
substituted by technology	(0.020)	(0.023)	(0.020)	(0.023)	(0.022)	(0.022)
Worry about substitution by	0.063 **	0.080 ***	0.012	-0.008	-0.025	0.075 ***
workers in other countries	(0.020)	(0.023)	(0.020)	(0.023)	(0.023)	(0.022)
Frequently needs to learn	-0.331	-0.213	1.014	1.097	1.391	0.385
new technologies	(0.770)	(0.900)	(0.781)	(0.881)	(0.867)	(0.862)
Concern about being	-0.262	-0.146	0.988	1.093	1.433	0.426
outperformed or replaced	(0.764)	(0.893)	(0.776)	(0.874)	(0.861)	(0.856)
Stressful to learn new	0.044 *	0.038	-0.015	0.000	-0.013	0.046 *
digital skills	(0.019)	(0.022)	(0.019)	(0.022)	(0.021)	(0.021)
Attitude towards big	0.053 *	-0.063 *	-0.109 ***	-0.072 **	-0.259 ***	-0.066 *
tech firms	(0.024)	(0.028)	(0.024)	(0.027)	(0.027)	(0.027)
Use of new technologies	-0.029	0.019	0.003	0.011	0.027	-0.005
at work	(0.017)	(0.020)	(0.017)	(0.020)	(0.019)	(0.019)
Perceived probability of losing	0.049 **	0.059 **	0.054 **	0.015	0.033	0.027
job in the next 5 years	(0.017)	(0.020)	(0.017)	(0.020)	(0.019)	(0.019)
Perceived employment	0.023	0.027	0.005	0.036	-0.017	-0.028
options if job loss	(0.017)	(0.020)	(0.017)	(0.019)	(0.019)	(0.019)
N	2515	2515	2515	2515	2515	2515
R2	0.048	0.045	0.055	0.022	0.076	0.063
logLik	47.348	-344.879	10.426	-290.289	-251.780	-237.086
AIC	-48.695	735.757	25.149	626.579	549.561	520.172

Note: The table presents the results of regressing support for technology regulation on attitudes related to how technology affects labor markets, technostress, attitudes towards big tech firrms, self-reported use of technology at work and general perceived labor market risks. Only respondents in the control condition are included.

likely to support regulation in three of the six areas studied: giving unions more capacity to decide about the introduction of technology in workplaces, support for governments intervention to make it more difficult that firms substitute workers by technology, and a tax on robots and algorithms.

By contrast, we do not find clear evidence that people who are more techno-stressed, as measured by our three indicators, have different attitudes towards technology regulation.

Respondents who hold more positive attitudes towards the Big Five tech firms are less likely to support regulation in five of the policy areas we consider, but, surprisingly, they are more likely to support giving unions more say in decisions about technology adoption. Turning to our measure of how frequently respondents use new technologies at work, we find no correlation with preferences for the regulation of technology.

Finally, we find evidence that workers who believe that they may lose their job in the next years are more likely to support regulatory measures for three of the indicators considered (trade unions, government intervention and the gig economy). This positive impact seems limited to precariousness. Our last indicator, on perceived outside options in case of job loss, does not correlate with preferences for regulation.

With this large battery of subjective-risk measures, overall, we do find some evidence that subjective concerns about technology and about labor market risks predict to some extent attitudes towards regulation, but the results are often inconsistent across policy areas. We conclude that subjective concerns are generally weak predictors of preferences about regulation and that we find only partial support for hypothesis 1b.

4.3 Experimental results

We now assess if the arguments we provided (appealing to the claim that certain groups can benefit from regulation and the narrative that technology fosters growth) affected support for the six policy questions.[11] Recall that we hypothesized that individuals appealing to how particular groups stand to win from a regulatory policy makes people from groups that stand to win from it become more likely to support regulation (hypothesis 2); that appealing to the benefits of regulation for individuals or communities that are harmed by technological change increases support for regulation (hypothesis 3); and appealing to the link between technological change and economic growth reduces support for regulation (hypothesis 4). Hypotheses 3 and 4 refer to the average (unconditional) effects of our treatments, and we start by testing them. Hypothesis 2 is about the interaction between selected characteristics of respondents and the treatment that emphasizes implications for some particular groups.

We report the results of the average treatment effects relative to the control group in Figure A-3 below. As we show in Appendix C, the results are identical when the models include basic socio-demographic controls.

[11] We report the results of balance tests in Appendix B.

Figure 2: *Main Experimental results (no controls)*

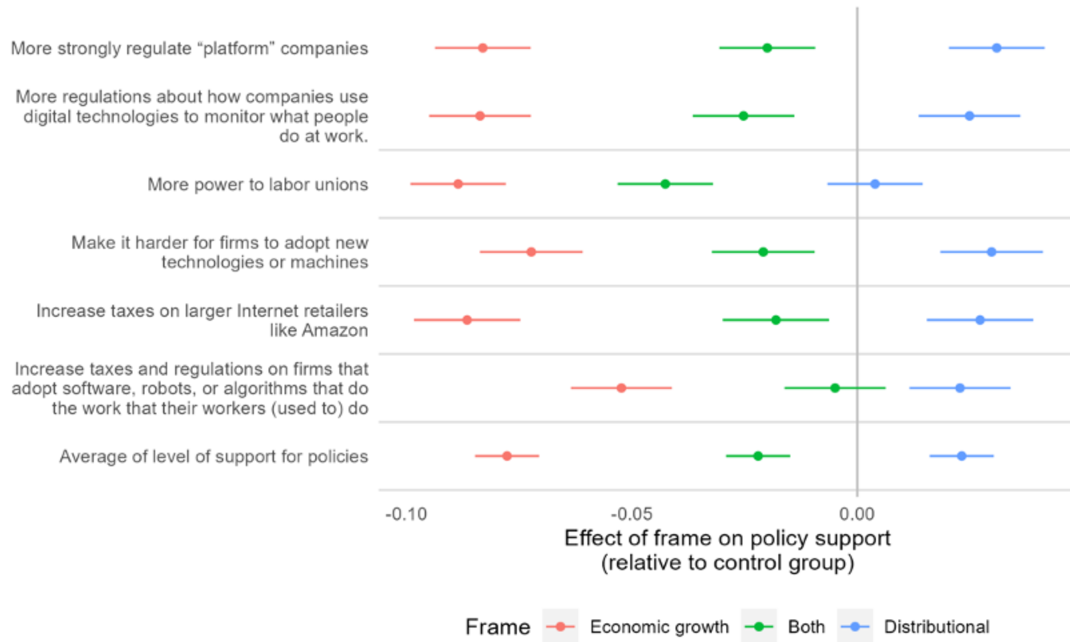

Note: The figure reports the differences in support for the six policy questions about the regulation of technology among each of the three experimental conditions, relative to the control group.

We find support for hypothesis 3. Arguments that appeal to winners and losers and point to the potential beneficiaries of regulatory policies (indicated by the blue dots) have a moderate capacity to affect opinion, relative to the baseline. Mentioning the beneficiaries of policies increases support for them by a small extent, a little under five percentage points across countries and policies.

We find much stronger support for hypothesis 4: arguments that link regulation to less economic growth or costs for consumers in general strongly reduce support for policy proposals (the red dots indicating assignment to this condition are far to the left of the baseline control, indicated by the vertical gray line). Clearly, the frame used by techno-optimist elites so far on the reasons why technology should not be regulated or taxed is effective, as in our sample it reduces support for such policies by about 10 percentage points fairly consistently across countries.

The impact of simultaneously providing both arguments is additive: in the treatment group that provides both arguments, support for the policies are slightly reduced relative to the baseline, which suggests that in a higher information environment, the socio-tropic argument is more persuasive than the distributional argument.

4.3.1 Conditional impact of moderators

Next, we turn to testing the more specific version of standard political economy models discussed in our hypotheses. Specifically, we test if individuals who are are mentioned as beneficiaries of regulation across the multiple policies are more likely to support government regulation of technology; we do so by assessing if there is a precisely estimated interaction effect between treatment assignment and relevant moderator. Figures 3 and 4 display the predicted probabilities of policy support across each treatment condition. Unless otherwise noted, all figures come from the same linear model from the previous figure, except for inclusion of an interaction term between treatment assignment and the relevant moderator. We focus on a select sub-set of the most relevant moderators that are actually mentioned within the treatment text and present additional results about moderators in Appendix A-5. For simplicity, we dichotomize all moderators.

For transparency, we display the impact of each moderator for each policy, but we are particularly interested in the question if the hypothesized moderators strengthen the effects of the distributional prime, which flagged the beneficiaries of the policies. For instance, we expect that less educated respondents are more likely to increase support for a regulatory policy when provided with information that this policy is particularly helpful for less educated people. Thus, the main feature of interest is whether the predicted probabilities of supporting a regulatory policy substantially vary at different values of the relevant moderators (indicated by triangles vs dots) in the distributional (light blue) condition. According to self-interest oriented political economy theories, respondents who are likely beneficiaries of a policy should be more likely to support such policy, and support should increase more when given information about the type of beneficiaries of the policy.

Figure 3: Experimental results by age, education, and risk of substitution by AI

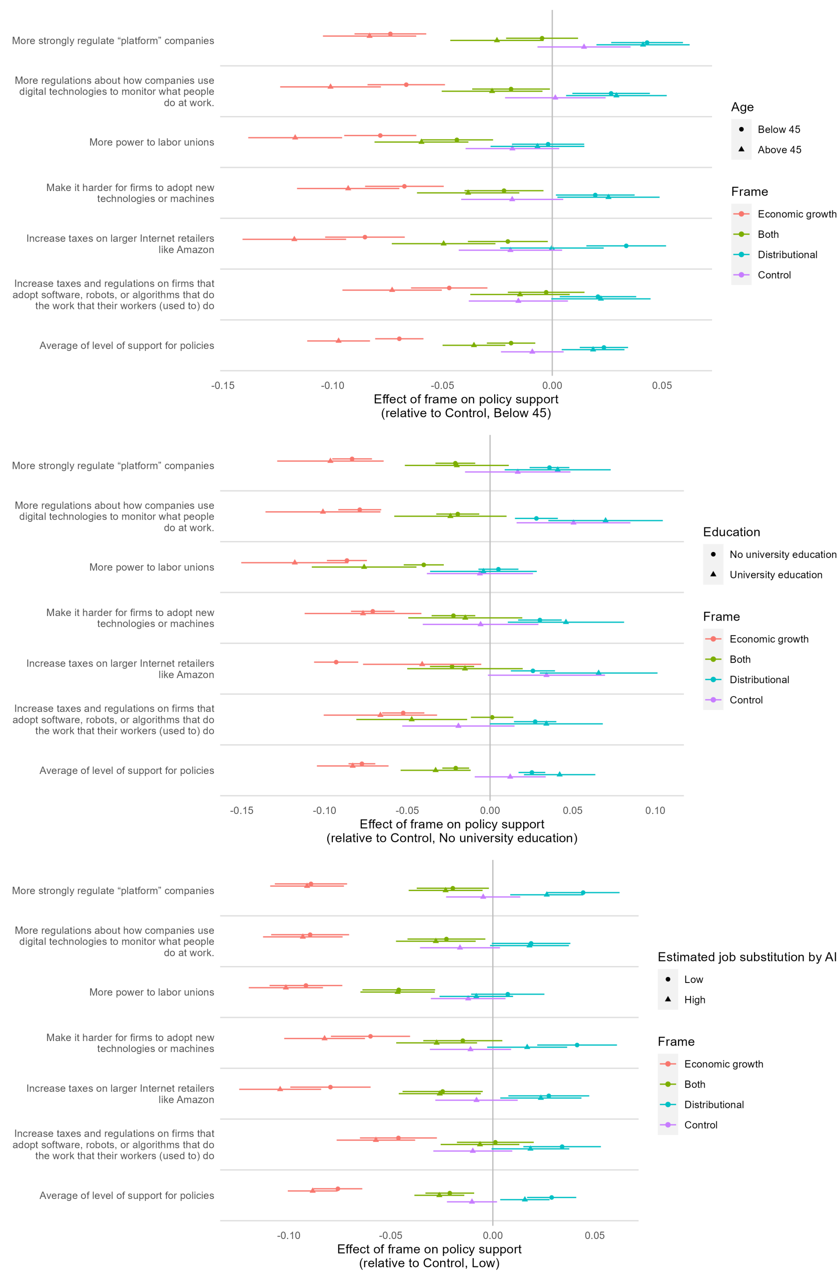

Note: The figure reports the treatment effects by age, education and risk of substitution by AI.

Figure 4: Experimental results by participation the gig economy and monitoring

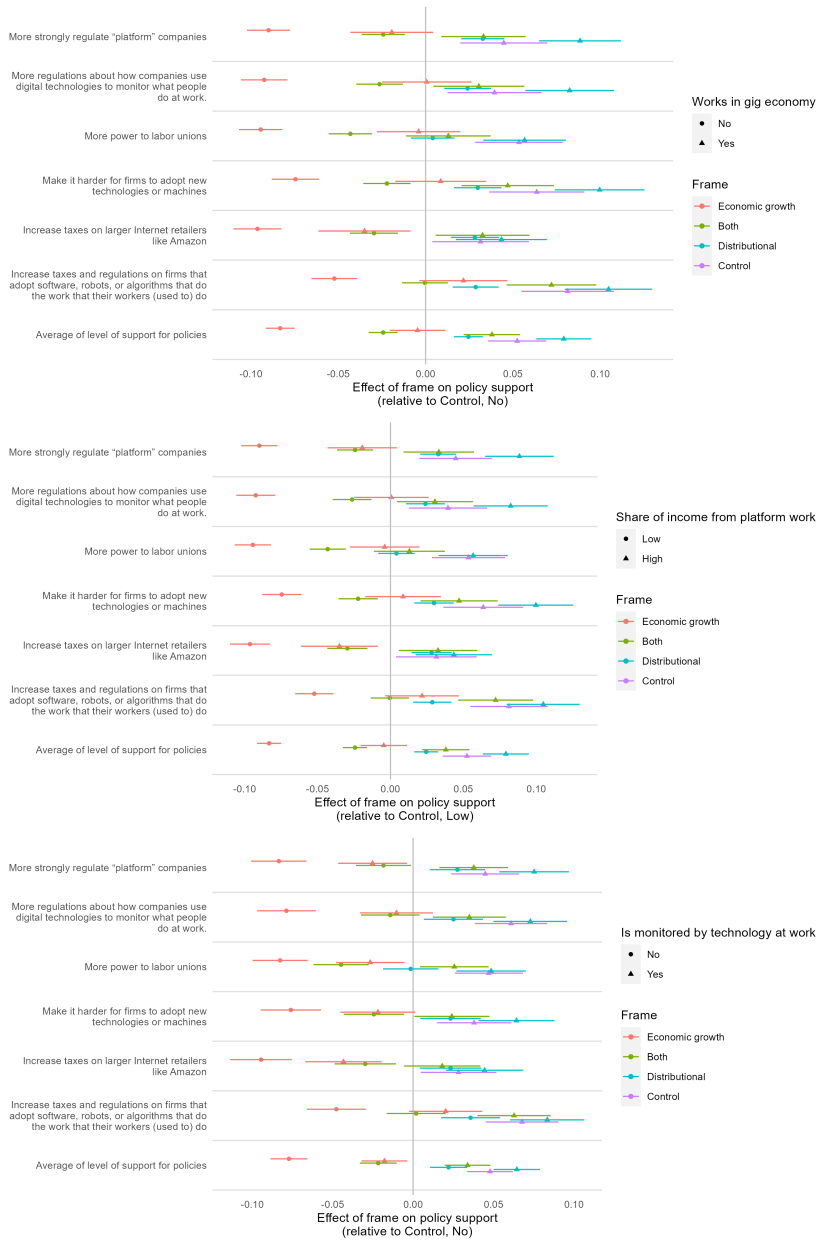

Note: The figure reports the treatment effects by work in the gig economy, percentage of income from the gig economy and experience of being monitored at work.

Our overall takeaway from the predicted probabilites displayed in the figures is that many of the most straightforward demographic and objective risk indicators mentioned in the treatment text do not strongly moderate (if at all) treatment effects. Neither the stronger negative effects of the socio-tropic frame, nor modestly positive effects of the distributive frame, appear to be strongly driven by the most obvious theoretically relevant demographic and occupational indicators. As the figures show, strikingly, individuals who are clearly indicated to be beneficiaries of various regulatory or taxation policies are not more responsive to treatments that highlight this. For example, less educated and older individuals are not more likely to favor the three labor-protecting policies more even when such policies are framed as such. Notably as well, our most precise measure of AI risk by occupation does not clearly moderate treatment effects for the distributional treatment (or any other). We overall find limited support for our hypotheses predicting moderating effects. [12]

5 Conclusions

We first summarize our results from the observational and experimental data with respect to our hypotheses, and then turn to extensions and implications. First, we find considerable support for all six technology regulation policies within the control group. Second, we find fairly limited to no evidence objective or subjective risks, across a wide variety of precise measures, correlate highly with regulatory support. A coarse measure such as older working-age correlates more highly than detailed occupation-based measures.

Third, we find that frames that connect to the core narrative that putting limits to technology carries costs for economic growth or for consumers in general are much more persuasive than frames that mention winners and losers of policies. This is generally consistent across countries and policies.

Fourth, we find that the aforementioned direct objective indicators do not moderate the effects of winner-loser frames, not even for people with the characteristics mentioned in those frames. For example, in frames that a regulatory policy can help older workers, older respondents in our dataset do not become more likely to support the policy. We do not find that our main treatment results are particularly driven by hypothesized moderators of age, education, and various forms of technological risk, either objective nor subjective;

[12] Due to space constraints we do not show the graphs for our battery of measures of subjective risk. For most such measures, the distributional prime has no conditional effect.

within the control group, there are not strong correlations with these demographic features and support for policies.

We conclude with caveats and thoughts about how these may inform further research. The frames we provide in this article are of course not exhaustive. While we think our experiments are sensible operationalization of theoretical concepts, there are alternative arguments in favor of regulation that span non-economic dimensions. Implicitly, our design has focused on testing the key "techno-optimist" position, but other arguments can be considered. More research needs to be done to identify the types of arguments that are more frequently used and those that are more effective, specially regarding the regulation of artificial intelligence.

An additional caveat is that we are studying attitudes towards the regulation of technology in a context in which this issue has a low political salience. Capturing preferences about an important but low salience topic is relevant as it provides a benchmark to track public debate and opinion in the future, something akin to what has happened in relation to climate change.

Given the importance of AI, we anticipate a growth in the salience of these topics in the future. There are some indications that the types of policies we are studying have the potential to become relevant for several reasons. First, there are abundant historical examples of cases when various political actors, including workers and citizens, have preferred not to adopt new technologies (Rosenberg and Curci, 2023; Johnson and Acemoglu, 2023). Second, the examples discussed in the introduction point in the direction of rising public debate about the regulation and taxation of AI. Third, even though survey measures of support for steering policies so far have been sparse and indirect, there is indication in current survey data that these policies may be palatable to workers affected by technological change (Gallego et al., 2022). Fourth, if technological change continues to increase income inequality, some policy-makers may arrive at the conclusion that preventing some changes is more feasible than dealing with the consequences of accelerated automation. Colantone and Stanig (2018) note that trade policies to preserve the status quo have been perceived as a more viable or politically appealing proposal as welfare states have become strained and the limits and costs of redistribution have increased. Given the potential of new technologies to increase income inequality, a similar logic may apply to technology.

For these reasons, it is entirely plausible that the regulation of technology, specially AI, becomes more salient as workers or communities concerned about potential ego-tropic or socio-tropic negative implications of technological change demand protection against change. In this context, we hope that our results can be considered as a benchmark for future studies on this topic.

References

Acemoglu, D. (2021). Harms of ai. Technical report, National Bureau of Economic Research.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022a). Artificial intelligence and jobs: evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D., Lelarge, C., and Restrepo, P. (2020). Competing with robots: Firm-level evidence from france. In *AEA papers and proceedings*, volume 110, pages 383–388. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Acemoglu, D. and Lensman, T. (2023). Regulating transformative technologies. Technical report, National Bureau of Economic Research.

Acemoglu, D., Makhdoumi, A., Malekian, A., and Ozdaglar, A. (2022b). Too much data: Prices and inefficiencies in data markets. *American Economic Journal: Microeconomics*, 14(4):218–256.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6):1488–1542.

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. (2022). Tasks, automation, and the rise in us wage inequality. *Econometrica*, 90(5):1973–2016.

Achen, C. and Bartels, L. (2017). *Democracy for realists: Why elections do not produce responsive government*. Princeton University Press.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2):645–709.

Autor, D. H. (2019). Work of the past, work of the future. In *AEA Papers and Proceedings*, volume 109, pages 1–32. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

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., 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.

Ayyagari, R., Grover, V., and Purvis, R. (2011). Technostress: Technological antecedents and implications. *MIS quarterly*, pages 831–858.

Baldwin, R. (2019). *The globotics upheaval: Globalization, robotics, and the future of work*. Oxford University Press.

Benabou, R. and Tirole, J. (2006). Belief in a just world and redistributive politics. *The Quarterly journal of economics*, 121(2):699–746.

Boix, C. (2021). Ai and the economic and informational foundations of democracy. In Bullock, J., Chen, Y.-C., Himmelreich, J., Hudson, V., Korinek, A., Young, M., and Zhang, B., editors, *The Oxford Handbook of AI Governance*. Oxford University Press.

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.

Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What can machines learn and what does it mean for occupations and the economy? In *AEA papers and proceedings*, volume 108, pages 43–47. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Burgisser, R. (2023). Policy responses to technological change in the workplace. Technical report, European Joint Research Centre (Seville site).

Busemeyer, M. R., Gandenberger, M., Knotz, C., and Tober, T. (2023). Preferred policy responses to technological change: Survey evidence from oecd countries. *Socio-Economic Review*, 21(1):593–615.

Busemeyer, M. R., Kemmerling, A., Van Kersbergen, K., and Marx, P. (2022). *Digitalization and the welfare state*. Oxford University Press.

Busemeyer, M. R. and Sahm, A. H. (2022). Social investment, redistribution or basic income? exploring the association between automation risk and welfare state attitudes in europe. *Journal of Social Policy*, 51(4):751–770.

Busemeyer, M. R. and Tober, T. (2023). Dealing with technological change: Social policy preferences and institutional context. *Comparative Political Studies*, 56(7):968–999.

Chong, D. and Druckman, J. N. (2007). Framing theory. *Annu. Rev. Polit. Sci.*, 10:103–126.

Colantone, I. and Stanig, P. (2018). The trade origins of economic nationalism: Import competition and voting behavior in western europe. *American Journal of Political Science*, 62(4):936–953.

Converse, P. E. (1964). The nature of belief systems in mass publics. *Critical review*, 18(1-3):1–74.

Dauth, W., Findeisen, S., Suedekum, J., and Woessner, N. (2021). The adjustment of labor markets to robots. *Journal of the European Economic Association*, 19(6):3104–3153.

Djourelova, M. (2023). Persuasion through slanted language: Evidence from the media coverage of immigration. *American Economic Review*, 113(3):800–835.

Felten, E., Raj, M., and Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12):2195–2217.

Gallego, A., Kuo, A., Manzano, D., and Fernández-Albertos, J. (2022). Technological risk and policy preferences. *Comparative Political Studies*, 55(1):60–92.

Gallego, A. and Kurer, T. (2022). Automation, digitalization, and artificial intelligence in the workplace: implications for political behavior. *Annual Review of Political Science*, 25:463–484.

Goos, M., Manning, A., and Salomons, A. (2009). Job polarization in europe. *American economic review*, 99(2):58–63.

Goos, M., Manning, A., and Salomons, A. (2014). Explaining job polarization: Routine-biased technological change and offshoring. *American economic review*, 104(8):2509–2526.

Huber, L. M. (2022). Beyond policy: the use of social group appeals in party communication. *Political Communication*, 39(3):293–310.

Jeffrey, K. (2021). Automation and the future of work: How rhetoric shapes the response in policy preferences. *Journal of Economic Behavior & Organization*, 192:417–433.

Johnson, S. and Acemoglu, D. (2023). *Power and Progress: Our Thousand-Year Struggle Over*

Technology and Prosperity. Hachette UK.

Jost, J. T., Banaji, M. R., and Nosek, B. A. (2004). A decade of system justification theory: Accumulated evidence of conscious and unconscious bolstering of the status quo. *Political psychology*, 25(6):881–919.

Jungherr, A. (2023). Artificial intelligence and democracy: A conceptual framework. *Social Media & Society*, 9(3):20563051231186353.

König, P. D. and Wenzelburger, G. (2019). Why parties take up digitization in their manifestos: An empirical analysis of eight western european economies. *Journal of European Public Policy*, 26(11):1678–1695.

Kurer, T. and Hausermann, S. (2022). Automationrisk, socialpolicypreferences, andpolitical-participation. *Digitalization and the welfare state*, page 139.

Lewis-Beck, M. S. and Paldam, M. (2000). Economic voting: an introduction. *Electoral studies*, 19(2-3):113–121.

Marenco, M. and Seidl, T. (2021). The discursive construction of digitalization: a comparative analysis of national discourses on the digital future of work. *European Political Science Review*, 13(3):391–409.

Mazur, J. and Serafin, M. (2023). Stalling the state: how digital platforms contribute to and profit from delays in the enforcement and adoption of regulations. *Comparative Political Studies*, 56(1):101–130.

Mutz, D. C. (2021). (mis) attributing the causes of american job loss: The consequences of getting it wrong. *Public Opinion Quarterly*, 85(1):101–122.

Rehm, P. (2009). Risks and redistribution: An individual-level analysis. *Comparative political studies*, 42(7):855–881.

Rosenberg, M. and Curci, F. (2023). Factory location: Resistance to technology adoption and local institutions. *Available at SSRN 4428398*.

Saris, W. E. and Gallhofer, I. N. (2014). *Design, evaluation, and analysis of questionnaires for survey research*. John Wiley & Sons.

Thau, M. (2021). The social divisions of politics: How parties' group-based appeals influence social group differences in vote choice. *The Journal of Politics*, 83(2):675–688.

Thewissen, S. and Rueda, D. (2019). Automation and the welfare state: Technological change as a determinant of redistribution preferences. *Comparative Political Studies*, 52(2):171–208.

Walter, S. (2010). Globalization and the welfare state: Testing the microfoundations of the compensation hypothesis. *International Studies Quarterly*, 54(2):403–426.

Webb, M. (2019). The impact of artificial intelligence on the labor market. *Available at SSRN 3482150*.

Weisstanner, D. (2021). Technological change and labour market policy preferences. In Clegg, D. and Durazzi, N., editors, *Handbook of Labour Market Policy in Rich Democracies*. Edward Elgar.

Wu, N. (2022). Misattributed blame? attitudes toward globalization in the age of automation. *Political Science Research and Methods*, 10(3):470–487.

Wu, N. (2023). "restrict foreigners, not robots": Partisan responses to automation threat. *Economics & Politics*, 35(2):505–528.

Zhang, B., Anderljung, M., Kahn, L., Dreksler, N., Horowitz, M. C., and Dafoe, A. (2021). Ethics and governance of artificial intelligence: Evidence from a survey of machine learning researchers. *Journal of Artificial Intelligence Research*, 71:591–666.

NOT FOR PUBLICATION

Appendix

These appendices contain materials, results, and robustness checks that supplement the main text.

A Risk of substitution by AI

We asked respondents about their occupation in the survey. Based on the coding of occupational codes at the four-digit level, and a cross-walk from SOC scores to ISCO cores, we assign each respondent a risk of being substituted by AI, as estimated by Webb (2019). The figure below provides the distribution of AI substitution risk in the data.

Figure A-1: *Distribution of risk of substitution by AI*

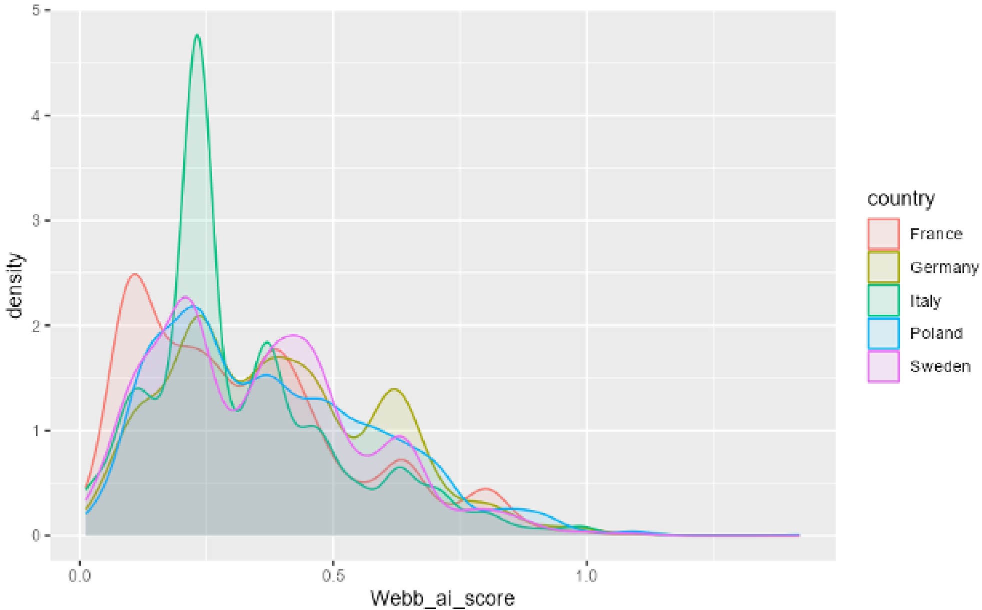

Note: The figure presents the distribution of risk of substitution by AI, using the measure proposed by measured by Webb (2019) in our sample.

B Balance tests

We check balance by regressing various characteristics on assignment to treatment group taking the control group as the baseline.

Figure A-2: *Results of balance tests*

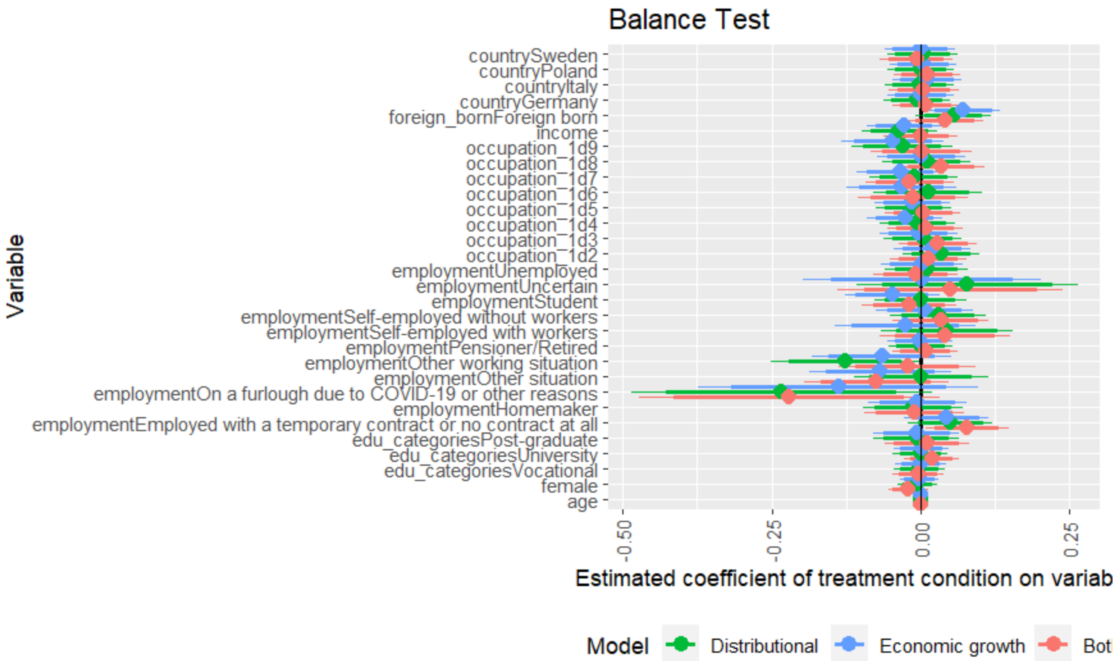

Note: The figure reports the results of regressing treatment assignment on country, origin, income, occupation, employment situation, education, gender, and age.

C Main experimental results with demographic controls

When adding demographic controls, the results remain almost identical as in the case of results presented in the main text, which do not include such controls.

Figure A-3: *Main Experimental results (including controls)*

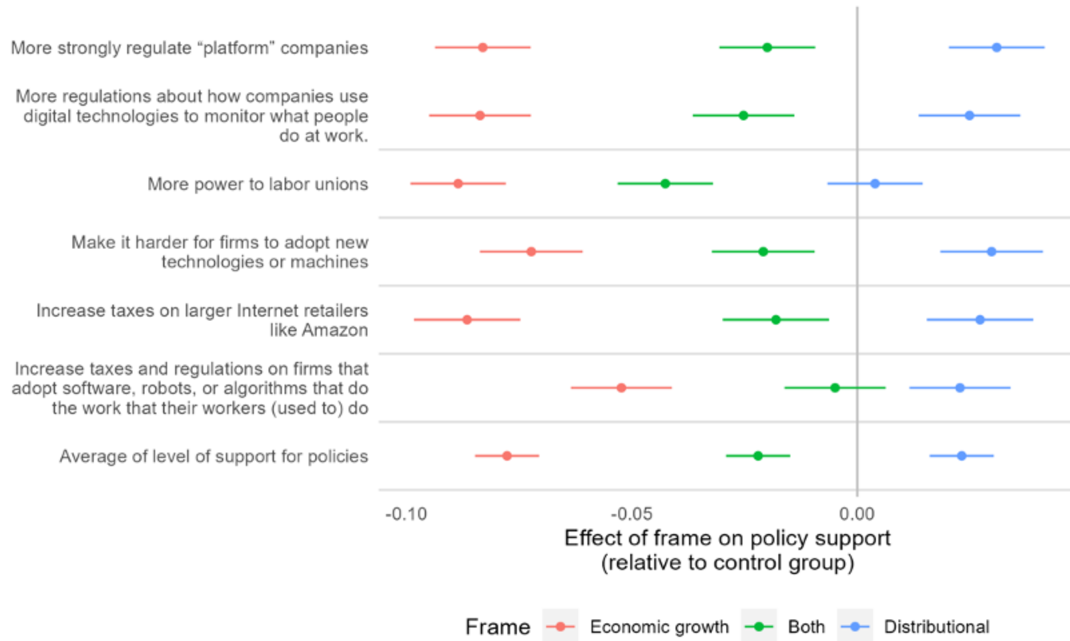

Note: The figure reports the differences in support for the six policy questions about the regulation of technology between each of the three experimental conditions, relative to responses in the control group. The models includes controls for age, education, gender, employment situation, origin, and country.

D Additional moderators

Figure A-4: Subjective risks as moderators

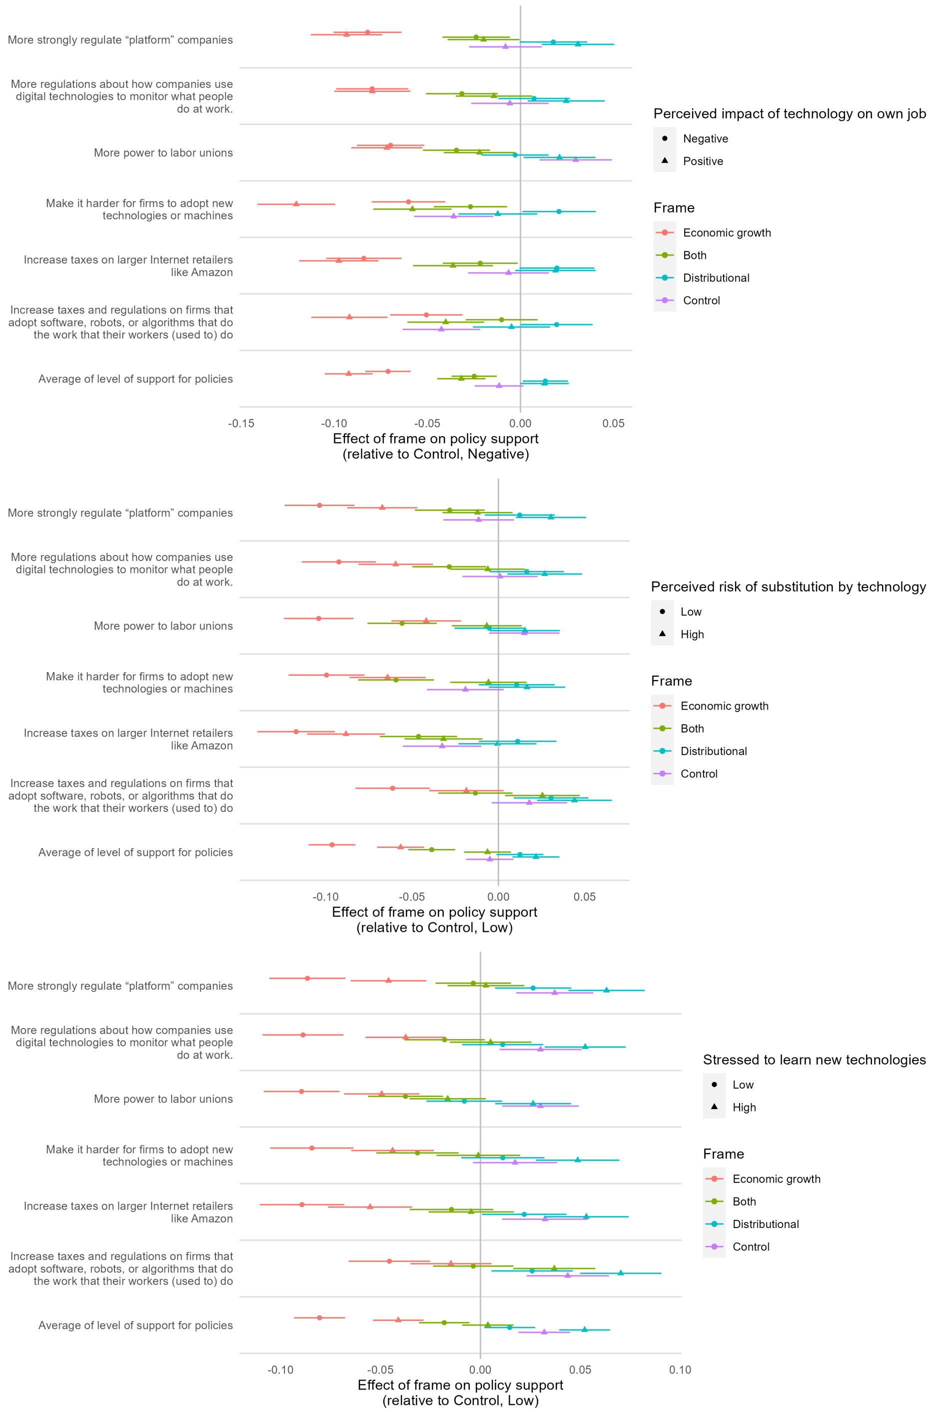

Note: The figure reports the treatment effects for respondents at high and low levels of subjective risk (perceived impact of technology on job, perceived percentage of tasks at risk of substitution, and technostress).

Figure A-5: Other variables as moderators

Note: The figure reports the treatment effects by frequency of online shopping, trust in unions, and vote for left-wing parties.