Three essays in empirical labor and political economics

Tora Kjærnes Knutsen

Dissertation for the Ph.D. degree

Department of Economics

University of Oslo

Series of dissertations submitted to the
Faculty of Social Sciences, University of Oslo
No. 915

ISSN 1504-3991

Cover: UiO.
Print production: Graphics Center, University of Oslo.

Table of Contents

Acknowledgements

Many people have contributed to this thesis. First, I am grateful for the input from my supervisors, Jo Thori Lind and Andreas Kotsadam, co-authors Marte Rønning and Jørgen Modalsli, and everyone else who has helped me with my research. Second, I want to thank Kjersti Misje Østbakken and Marte Strøm at the Institute for Social Research for allowing me to get data for my work on industry specific wage floors through their project. Finally, I am very grateful for the support I have gotten throughout the PhD from fellow PhD-students, friends, family and maybe most importantly José.

Tora Knutsen, April 2022

Introduction

This thesis consists of three articles focusing on the effects of different policy measures on voting behavior, educational choices and wages. All papers are in applied microeconomics and use empirical methods aiming to identify causal effects. The findings are examples of how policy interventions may not ultimately have the intended effect.

In the first chapter, we study how foreign aid affects the support for the incumbent leader. While foreign aid has received criticism in recent decades for allowing politicians to distribute patronage and influence voting, the evidence that aid can increase incumbent support is so far scarce. Comparing areas before vs. after receiving aid, we find that aid projects do, on average, lead to increased incumbent support. However, this effect is only found for World Bank projects and seems to act through trust in politicians.

In the second chapter, we find that a large geographical expansion of higher education in Norway involving the establishment of 9 nursing colleges, 11 new engineering colleges and 13 regional colleges, did not increase the education level for those living in commuting distance to the new colleges. While the aim of the expansion was to increase access to education, the new colleges instead affected the choice of field of study. The establishment of nursing colleges was, for example, followed by a decrease in the propensity for local youth to undertake a degree in teaching.

In the third and final chapter, I study the extension of wage floors introduced in industries with a rising share of labor immigrants. I find that this policy most likely did not increase wages even for workers who were low-waged prior to the intervention. However, there is some evidence that workers new to the labor market or the industry did experience wage gains and that some groups of immigrant workers may have experienced higher wage growth compared to other immigrant workers.

These findings highlight the importance of studying the effects of policy interventions.

In all chapters effects are estimated using variations of a difference-in-difference research design. In recent decades, empirical microeconomics has undergone what has been dubbed a "credibility revolution", driven by the focus on the quality of empirical research designs (Angrist and Pischke, 2010). In earlier research, econometricians had attempted to statistically model the world to make causal inferences, for example by controlling for confounders and making functional form assumptions. In essence, the "Credibility Revolution" moved the focus from statistical modeling to research design and causal inference. With more focus on research design, experimental and quasi experimental methods received more attention. In particular, empirical economics has drawn from methods used in clinical medical trials, where a group of participants is randomized into a treatment group receiving the medication and a control group receiving a placebo medication. As the treatment is randomized, any difference between the two after the intervention can be attributed to the treatment received. In economics, randomized controlled trials (RCTs) have increased rapidly in popularity and is seen by many as the "gold standard" in studies of policy interventions. In a RCT, like in a clinical medical trial, who receives the treatment is randomly selected from the eligible population. While using random assignment is a clear advantage, experiments are usually expensive and not possible in practice if one wants to study the effect on an intervention ex post. Therefore, quasi experimental research designs have become a popular tool to study policy interventions. In this context, the difference-in-difference (DiD) estimation method has been widely used and in fact almost one fourth of all NBER working papers in applied micro today make references to to this method (Currie et al., 2020).

In the "canonical" Difference-in-Difference set-up with two groups and two time periods, a researcher compares the change in outcomes in a (non-random) treatment group before vs. after treatment to the change in outcomes in a comparison group over the same time period. Under certain assumptions, the DiD approach removes selection bias by differencing out pre-treatment differences between the treatment group and the comparison group, and the time trends by differing out pre-period vs. post-period changes in outcomes within the comparison group. The DiD approach then leaves us with a credible quasi-experimental estimate of the treatment effect of interest. However, often one wants to study reforms or interventions that affect different units at different times. In these cases researchers often add time fixed-effects and unit-fixed-effects. This is often referred to as "two-way-fixed-effects" and forms the basis for the methods used in all three papers in this thesis. In recent years, the two-way-fixed-effects has been under scrutiny by econometricians suggesting that the estimates stemming from such analysis are poorly understood (Callaway and SantAnna, 2020; De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021). In particular, the ordinary least square (OLS) estimate coming from this procedure cannot necessarily be interpreted as the average effect as the two-way-fixed-effects model weighs units depending on when they are treated and the relative sizes of the treated and control groups.

In the following, I first summarize the findings of my three thesis chapters. Next, I turn to discuss the identification model used in the three papers in light of new developments in econometrics suggesting that our understanding of the two-way fixed effect model with variation in treatment timing has been limited. I argue that the two final chapters, to some extent, are robust to this critique, and I discuss how new identification models can be used for the analysis of aid and incumbent support. Finally, I discuss whether the research in this thesis comply with new standards when it comes to transparency in research. While using register data enables us to do analysis on the whole population, restricted data access makes reproducability more challenging than in studies using publicly available survey data.

Chapter 1: Foreign aid and support for the incumbent leader

In the first paper, co-authored with Andreas Kotsadam, we study how foreign aid affects voting for the incumbent leader. The manifested goal of foreign aid is to strengthen economic and social development. However, the political economy of foreign aid has been widely debated for decades and remains one of the most controversial debates in development economics (Qian, 2015). Critics point to the potential adverse effects of aid on institutions and corruption. One line of argument suggests that politicians use aid to distribute patronage or to influence voting (Jablonski, 2014). Yet, there is little evidence on whether the voters actually respond by increasing their support for the incumbent leader. In this paper, we find that foreign aid may increase the support for the incumbent leader, but that this effect is not universal across donors; we find a positive effect from World Bank projects and no effect from Chinese aid project. Our results suggest that this effect goes through increased trust in politicians.

Theoretically, the effects of aid on the incumbency advantage are ambiguous. Voters may reward incumbent leaders who strategically locate aid to their area, either if the leaders do it in order to please core supporters or to attract swing-voters. Incumbency support may also increase due to economic voting motives, if aid is creating favorable economic conditions in the areas. Aid may also lead to lower support for the incumbent leader if it undermines the capacity and legitimacy of recipient governments (Knack, 2001) or if it reduces the political accountability (Ahmed, 2012).

Most previous studies investigating the political effects of aid use cross-sectional data at the country level. The question of strategic allocation of aid projects for political gains is however, inherently local. We therefore follow the recent and growing trend towards analyzing the effects at the local level made possible by the availability of geocoded data on both aid and outcomes (Briggs, 2017; Dreher et al., 2019; Hodler and Raschky, 2014; Jablonski, 2014; Masaki, 2018; Nunnenkamp et al., 2017; Ohler and Nunnenkamp, 2014). In this study, we move beyond testing for biased aid allocation to an investigation of the political effects of targeting. Up until now, such studies have mostly been conducted using single countries. A notable exception is Briggs (2018), who finds that aid lowers the support for incumbent presidents in a local level analysis of aid projects in three countries (Nigeria, Senegal, and Uganda). We add to this literature by conducting a study in many countries, using project level fixed effects, and by analyzing the effects of different donors.

By matching geo-coded data on aid projects to 101 792 respondents in five waves of the Afrobarometer, we investigate the effects of aid on incumbency support using project fixed effects. In the regression model we use fixed effects for areas surrounding an aid project and investigate how the support changes when an aid project starts. We estimate the effects for World Bank aid and Chinese aid separately and find positive effects for the former and no robust effects for the latter. We find little indication of aid locations being selected based on previous incumbency support and we find no effects on overall electoral turnout.

Furthermore, the effects of aid on incumbency support are not more positive when power is more contested, i.e. in countries and periods with higher political competition. A possible explanation for the differential impact of aid on incumbency support is that the effects are mediated by the effects on local living conditions. While this explanation is compatible with previous research finding that economic conditions become better with World Bank aid, while corruption (Isaksson and Kotsadam, 2018) as well as favoritism (Dreher et al., 2019) increase with Chinese aid, we fail to corroborate this mechanism as we do not find that self-reported living conditions are differentially affected or that they seem to mediate the effects. Following Briggs (2018), we investigate the effects of aid on trust in government and find that the positive effects for the World Bank aid projects seem to be mediated by trust in politicians while we find no such effect of Chinese aid.

Chapter 2: The location of higher education institutions, education level and field of study

Currently, the location of higher education institutions is subject to political controversy. Local college access is put forward as an important way of securing access to education and increasing the supply of skills demanded by the local labor market. However, so far evidence that local college access in fact has the potential to affect the education level of the local population is scarce. In this paper, we find that distance to college might matter for the choice of educational field, but not educational level.

The location of university colleges in Norway has been subject to political controversy for decades. In the middle of the twentieth century, the level of higher education in Norway was low with about 2 percent of the 1940 birth cohort and 7 percent for the 1950 cohort having a college degree. Geographical background was important in predicting education level, but the educational advantage by geographical centrality was reduced considerably in the period after 1960 (Lindbekk, 1998). In this period policies aimed at equalizing social and economic differences across the country gained wide popular support. One such intervention was the establishment of higher education institutions across the country, and throughout the 1960s and 1970s Norway experienced a rapid increase in university colleges whose location was decided politically, and was not necessarily referring to local demands or resources. Instead, the location of the colleges was chosen with the purpose of providing each of Norway's nineteen counties colleges offering degrees in nursing, engineering, teaching and business administration (Ottoesen, 1969).

The exact location of the colleges was decided by the parliament, where it was subject to extensive debate. Even though recruitment of students and professional environment were supposed to be considered, debates in the parliament suggest that the equalization of educational opportunities across regions was more important. The minister of education, Kjell Bondevik, later regretted that regional political considerations had been decisive for the location of colleges (Johnsen, 1999). As a result of this targeted district policy measure, the number of colleges outside of the biggest cities increased substantially: 9 new nursing colleges, 11 new engineering colleges and 13 business administration colleges were spread around the country. While the number of students enrolled in the universities remained the same, there was a substantial increase in the number of students enrolled in the colleges (Johnsen, 1999).

In this paper, we study how this geographical expansion of higher education institutions in Norway affected outcomes - in terms of choice of field of study, education length and income as adult - for individuals growing up close to the new college establishments. The increase in education institutions across the whole country increased the take-up of the degrees provided at the newly established colleges. As percentage of the population, the number of nurses and college engineers doubled between 1960 and 1980.

We use administrative data on education and earnings as well as censuses going back to 1960 to track field of study, earnings and the municipality of residence in a given year for the entire Norwegian population. Using a difference-in-difference approach with flexible age cut-offs and fixed effects at the cohort and municipality level, we compare individuals within a municipality getting college access at different ages to individuals living without local college access. Our findings indicate that young adults residing within distance to a new college show a significant increased take-up of the new education opportunities. However, men do not change their behavior in response to nursing degrees being offered nearby, and women do not respond to engineering degrees. This fits in with the general pattern seen in Norway and elsewhere; women often choose health related fields and work in the public sector, while men choose science and engineering related fields (Card and Payne, 2020). For business and administration, the take up rates do not differ across men and women.

We find no increase in the overall attainment of education, implying that the increase in the degrees offered came at the expense of field of study not offered locally. The establishment of a nursing college is associated with a decrease in the propensity to do a teaching degree for women. For men, establishment of new engineering colleges is negatively correlated with take-ups of degrees in other technological fields and business administration. Furthermore, when the new business administration degree was introduced, men chose it at the expense of engineering, while for women there are indications that the increase in this degree came at the expense of teaching and social sciences.

Although the college premium is important in explaining wage inequality, it is by now well-documented that field of study also matters for labor market outcomes including the gender wage gap (Altonji et al., 2012; Hastings et al., 2013; Kirkeboen et al., 2016). While higher education is associated with higher earnings, the effect of a shift in field is not clear. If people choose field according to their comparative advantage as found in Kirkeboen et al. (2016), decreasing the cost of particular fields may induce people to not choose optimally. We find few pronounced changes in mid-life income and labor market participation, suggesting that the changes in field of study for the treated population did not result in higher earnings. However business administration is one exception: While we find zero or small positive changes in male wages, we do observe a decrease in the labor market participation for women living in areas with a new established higher education institution offering a college degree in business administration. We ask two questions in pursuit of plausible explanations for this result. First, to what extent did the establishment of a college induce women to remain in their home municipality? If they did, they could have missed out on better labor market opportunities elsewhere. Second, did the business administration degree offer particularly bad labor market opportunities for women? We find some support for both explanations: The college establishments had an effect in retaining women in their home municipality, while we find a smaller effect on men. Also, the mapping from degree to occupation differs across gender: A business administration degree is associated with manager positions for men, while women concentrate in occupations that tend to pay less, such as general office workers.

To ensure that our results are driven by the college openings and not other confounding factors, we control for municipality level time trends in our main specification and show that our results are robust to a range of different geographical definitions of college access. In addition, we follow a previous study by Bhuller et al. (2017) showing that municipality level characteristics in 1950 and 1960 cannot predict reform year.

Chapter 3: Wages and extensions of collective bargaining

Similar to the other Scandinavian countries, Norway has never had a national minimum wage. Instead, 70% of workers in the private industry work for employers covered by a collective agreement. Since the enlargement of the European union in 2004, with new member countries included in the European labor market, legally enforced wage floors have been introduced in several industries in Norway to avoid too large wage differentials between domestic and immigrant workers. These wage floor introductions provide a unique setting for studying the effect of legally enforced wage floors. While voluntary and legally enforced collective bargaining coverage have been studied extensively, less is known about the effects of transitioning from a voluntary to a legally enforced wage floors.

In this chapter, I focus on the first three industries where wage floors were introduced: construction, shipyards, and cleaning. In these, minimum wages were introduced in 2006/2007, 2009 and 2012, respectively. The study of the different industries offers insight into mandatory wage floor introductions in very different parts of the labor market: While cleaning is an industry with a large share of low-wage workers, this share is negligible in shipyards and low in construction. The minimum wage to median wage ratio is between 0.5 and 0.65, highest in cleaning, and lowest in shipyards. These ratios are fairly high in an international context.

Using Norwegian administrative registers on all employees in Norwegian companies, I can trace workers wages and labor market trajectories over a period of 15 years. This enables me to control for all time-invariant unobserved individual characteristics. To study the effect of introducing a minimum wage, I first use a difference-in-difference strategy with individual fixed effects to compare workers' wage growth in the targeted industry to other workers. Surprisingly, workers in construction and cleaning experience less annual wage growth after the minimum wage introduction compared to other workers.

While an individual fixed-effects model does control for the correlation between the outcome and any background variable that is constant over time, it does not pick up whether segments of the population experienced specific trends over time not linked to changes in sectoral wage policy. Norwegian register data includes a rich set of background variables on educational field and length, immigrant background and work experience. I use post-estimation LASSO to select a set control variables, finding that the the results do not change substantially when adding controls. Moreover, these industries may form part of a subsection of the economy that has had less wage growth. To find similar sectors, I construct synthetic control sectors with similar wages prior to the minimum wage introduction.

The results suggest that introducing a wage floor has mixed effects: wages in cleaning and construction decreased compared to the synthetic control industry, while for shipyards wages seem to have increased, but changes in wage growth seem not to be timed according to the wage floor introduction.

Most studies of the minimum wage use either aggregated or repeated cross-sectional data, offering no possibility of tracking individual trajectories. Tracking individuals over time enables me to understand who is driving the effect. I compare workers in the targeted industries to workers in the synthetic control industry to explore which groups experienced less wage growth after the minimum wage introduction. The results at the individual level are subject to larger uncertainty due to the selection process of the sample. Yet, the results do indicate that the workers experiencing less wage growth in cleaning and construction are full-time workers. In cleaning, the negative effect is driven by low-wage workers while this is not the case in construction. Decline in wage growth is also associated with union membership and experience. Thus, the zero or negative change in wage growth for construction and cleaning workers could come from a lower relative growth (or decline) in the return to experience. Moreover, I find no evidence that employment is affected by the wage floor policy. At the individual level, the probability of exiting the labor market or changing sector seem not to change.

How much can we learn from two-way-fixed-effects regressions?

A difference-in-difference analysis with time- and unit-fixed-effects is often referred to as ”two-way-fixed-effects”. This means that there may be many groups and that they are observed over several time periods. Time-fixed effects capture time trends common to all units, while unit-fixed-effects allows every unit to have a separate intercept capturing any effect, not changing over time, specific to the unit. When different units receive treatment at different points in time it is often called a ”staggered” difference-in-difference. The three chapters of this thesis all include a version of this set-up. Traditionally one of the main worries in such set-ups has been the assumption of common trends and similar potential outcomes of the treatment and control group in the absence of treatment. However, in recent years there has been developments questioning the use of two-way-fixed effects itself. While the estimator found in a typical regression set-up was previously thought of as an estimate of the average treatment effect on the treated, recent papers suggest that this estimator is, at best, a weighted average that is difficult to interpret (see for example Callaway and SantAnna, 2020; De Chaisemartin and d'Haultfoeuille, 2020; Goodman-Bacon, 2021). In particular, several papers have pointed out that early adopters serve as controls for later adopters. This may pose a serious threat to the common trend assumption.

Goodman-Bacon (2021) shows that the estimated coefficient in a two-way fixed effects estimate with variation in treatment timing is a weighted average of (1) comparisons between early adopters and later adopters over the periods when the later adopters are not yet treated; (2) comparisons between early adopters and later adopters over the periods when the early adopters are treated; and (3) comparisons between different timing groups and the never-treated group, if there is one. According to Goodman-Bacon (2021) the most troubling is (2), that already treated units serve as controls. When treatment effects are not constant over time, using already treated units as controls necessarily biases estimates of the treatment effect. The weights each of the different treatment groups get depend also on how big the treatment and control group is compared to each other. Moreover, OLS over-weights units with more variance in treatment status in order to achieve a more precise estimate of the treatment effect. Hence, units that are treated near the middle of the evaluation window receive relatively more weight. Thus, some re-weighting is required to estimate the average treatment effect on the treated.

Callaway and SantAnna (2020) propose a new way of estimating a staggered difference-in-difference including also the case when the control group consists of the not-yet-treated. However, they recommend discarding already treated units from the control group as they get treated. In the paper on foreign aid this criticism is very relevant as we use the variation between not-yet-treated and treated. This means that it is unclear as to where exactly the treatment effect comes from and whether it is biased. However, using the method suggested by Callaway and SantAnna (2020), it is unclear how to estimate the treatment effect for units treated at the end of the time-period as there are no more non-treated units left. This is illustrated in Figure 1 that shows separate regression by cohort, that is the first year a unit is treated using only not-yet-treated as the control group. This means that, for example, 2004 refers to a regression with units treated in 2004 as the treated group and only units receiving aid projects later as the control group. The figure shows that the confidence interval increases as the control group gets smaller. Moreover, comparing units treated in 2012 to those treated in 2013 and 2014 yield a negative estimated effect. However this estimate is based on few units not necessarily representative even for the sample. Figuring out under which circumstances foreign aid increases incumbent support and when it possibly does not seems like an interesting avenue for future research.

A way for this to work could be to include areas without foreign aid projects in the control group. However, areas not receiving aid projects are potentially different from areas receiving aid projects on important unobserved factors. It is therefore not clear whether

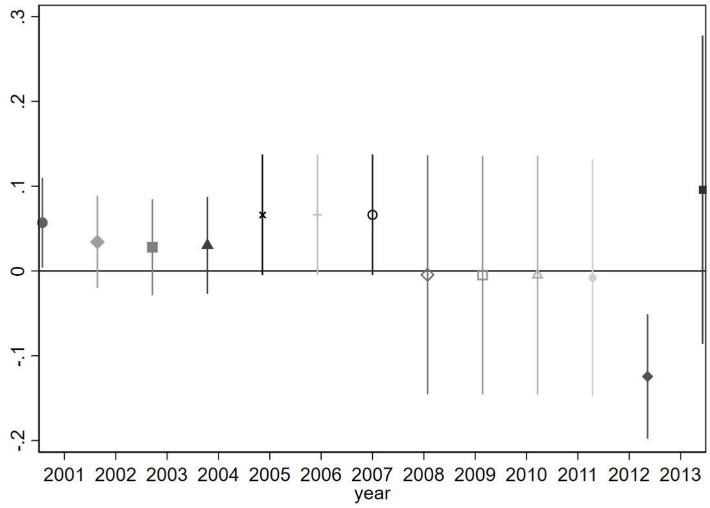

Figure 1: Effect of aid on incumbent support. Estimates using only not yet treated as control

Note: Coefficient plot shows the coefficient from the regression on the effect of aid on incumbent support including time- and project fixed effect (as in chapter 1 of the thesis). Each estimate x year combination refers to separate regressions including only units treated in year t and units that have not yet received treatment.

they would be suitable as a control group. Thus, there is a clear trade-off between the traditional worry in a difference-in-difference regression, namely developments in potential outcomes, and the new insights suggesting that a research design with staggered treatment timing where all units are eventually treated has problems. Potentially one could find ways to predict aid project location and compare areas with the same likelihood of receiving an aid project. Other approaches include the synthetic control method and prediction using Machine Learning methods such as LASSO such as those used in the third chapter of this thesis.

In the second chapter on educational expansion the control group of never-treated is large in comparison with the treated group. Therefore, the already treated units are less weighted when they serve as the control group. However, the estimates in the paper might still be biased with some treated units serving as control. Moreover, the municipalities that are treated towards the middle of the time-period studied may be attributed more weight in the OLS regression than those getting college access earlier or later. This is problematic if there is heterogeneity in treatment effects across municipalities or across time. Thus, the estimated effects may not be representative for all municipalities getting new colleges. If this data becomes available to me again, it could be helpful to explore heterogeneity in treatment effects by cohorts and understand which exact college establishments are really driving the treatment effects.

In the third paper, instead of doing a staggered DiD, I do separate regressions for each minimum wage introduction. The reason for doing this initially was that I suspected a wage floor would not have the same effect across sectors. Moreover, in this case the control group mostly consist of never-treated individuals and therefore the already treated are given lower weights. However, using the whole population as a control group potentially threatens other assumptions in a DiD analysis as explained above. I propose two solutions to this problem. First, I use a post-estimation LASSO procedure to choose a set of control variables. Next, given the richness of the data that includes the universe of industries in Norway, I can use the synthetic control method to find a weighted average of industries matching the pre-trend of the treated industry. Thus, the different papers in this thesis with staggered DiD designs illustrate how complex two way difference-in-difference estimation is and that having a lot of data and many time periods does not necessarily solve the fundamental issues such analyses are subject to.

Transparency and publication bias: Register data versus survey data

Even though published research may have been rightfully executed and is perfectly reproducible, the results may still suffer from biases. Moreover, studies with null findings may not survive the publication process due to reviewers and editorial decisions, while studies with large effects are more likely to get published (Franco et al., 2014). Brodeur et al. (2016) find a hump around the z-statistics related to a p-value of 0.05 when looking at the z-statistics of the results published in three of the most important journals in economics between 2005-2011. This pattern is an indication that researchers in economics search for specifications giving marginally significant findings as it implies that published findings more often have a z-value close to 1,96 than one would expect in a world without any censoring of p-values.

In recent years, the need for transparency in social science has received large attention. Those advocating for more transparency call for reproducibility, replicability and pre-registered research plans (Christensen and Miguel, 2018). An advantage of survey data such as the Afro Barometer, is that it is publicly available online (although a special application is needed in order to get the geographical locations of respondents). Thus, the question on incumbent support can in theory be studied again by anyone. They might use new data or employ a different empirical strategy and find different results.

Traditionally, a concern in science has been that research finding are less likely to be true when studies conducted have a small number of observations (Ioannidis, 2005). The Nordic countries have the advantage of a long history of public data collection covering the whole population. Having data for the whole population mostly removes any small-sample bias as it provides large scale evidence, at least at the country level. However, there are ethical concerns linked to the use of this data as the combination of different registers may reveal information that makes it possible to identify individuals. Privacy is a legal concern and the use of register data is therefore strictly regulated. As a consequence, data openness and reproducibility in registry-based research is challenging compared to research using anonymous survey data including only samples of the population.

Luckily, initiatives such as microdata.no, allow register data to be available for researchers and students connected to Norwegian research institutions.[1] Moreover, a newly started

[1] To access microdata.no you have to be a researcher, PhD student or master student at a research institution or employed in a ministry or directorate. The institution needs to sign an agreement with Microdata in order to participate.

journal named "The Journal of Comments and Replications in Economics", will as the name suggests, publish replication and comments of published articles in economics. In terms of my research, I hope that the extension of wage floors will be studied again as new wage floors have been introduced. If this policy does not have the desired effect we need to think of other policies or measures to increase compliance.

None of these studies were pre-registered. It is not obvious how to pre-register research on observational data as the researcher has the chance to look into the data ahead of the study. However, pre-analysis plans can also be used in non-experimental settings as demonstrated by Clemens and Strain (2021) in a study on the effects of minimum wage changes. They used early data to estimate short-run effects and committed themselves to analyze later data with a common set of estimation framework. As there have been heated debates over research designs for studying minimum wages, pre-registering the design may be a fruitful way forward. It would for example be possible to pre-register a research design to study the effects of new extensions of wage floors. In the case of a relatively inexperienced PhD student however, it does seem less realistic to commit to a pre-plan as the project one does seldom ends up as planned. Moreover, surprising results may change the overall research question. This is the case for the wage floor project, where the original idea was to check out secondary outcomes implied from wage increases. As wage increases were impossible to detect, the overall research question changed fundamentally too.

References

Ahmed, Faisal Z (2012). "The perils of unearned foreign income: Aid, remittances, and government survival." In: *American Political Science Review* 106.1, pp. 146–165.

Altonji, Joseph G, Erica Blom, and Costas Meghir (2012). "Heterogeneity in human capital investments: High school curriculum, college major, and careers." In: *Annu. Rev. Econ.* 4.1, pp. 185–223.

Angrist, Joshua D and Jörn-Steffen Pischke (2010). "The credibility revolution in empirical economics: How better research design is taking the con out of econometrics." In: *Journal of economic perspectives* 24.2, pp. 3–30.

Bhuller, Manudeep, Magne Mogstad, and Kjell G Salvanes (2017). "Life-cycle earnings, education premiums, and internal rates of return." In: *Journal of Labor Economics* 35.4, pp. 993–1030.

Briggs, Ryan C. (2017). "Does Foreign Aid Target the Poorest?" In: *International Organization* 71.01, pp. 187–206. ISSN: 0020-8183.

— (2018). "Receiving foreign aid can reduce support for incumbent presidents." In: *Political Research Quarterly*.

Brodeur, Abel, Mathias Lé, Marc Sangnier, and Yanos Zylberberg (2016). "Star wars: The empirics strike back." In: *American Economic Journal: Applied Economics* 8.1, pp. 1–32.

Callaway, Brantly and Pedro HC SantAnna (2020). "Difference-in-differences with multiple time periods." In: *Journal of Econometrics*.

Card, David and A Abigail Payne (2020). "High school choices and the gender gap in STEM." In: *Economic Inquiry*.

Christensen, Garret and Edward Miguel (2018). "Transparency, reproducibility, and the credibility of economics research." In: *Journal of Economic Literature* 56.3, pp. 920–80.

Clemens, Jeffrey and Michael R Strain (2021). *The Heterogeneous Effects of Large and Small Minimum Wage Changes: Evidence over the Short and Medium Run Using a Pre-analysis Plan*. Tech. rep. National Bureau of Economic Research.

Currie, Janet, Henrik Kleven, and Esmée Zwiers (2020). "Technology and big data are changing economics: Mining text to track methods." In: *AEA Papers and Proceedings*. Vol. 110, pp. 42–48.

De Chaisemartin, Clément and Xavier d'Haultfoeuille (2020). "Two-way fixed effects estimators with heterogeneous treatment effects." In: *American Economic Review* 110.9, pp. 2964–96.

Dreher, Axel, Andreas Fuchs, Roland Hodler, Bradley C Parks, Paul A Raschky, and Michael J Tierney (2019). "African leaders and the geography of China's foreign assistance." In: *Journal of Development Economics* 140, pp. 44–71.

Franco, Annie, Neil Malhotra, and Gabor Simonovits (2014). "Publication bias in the social sciences: Unlocking the file drawer." In: *Science* 345.6203, pp. 1502–1505.

Goodman-Bacon, Andrew (2021). "Difference-in-differences with variation in treatment timing." In: *Journal of Econometrics*.

Hastings, Justine S, Christopher A Neilson, and Seth D Zimmerman (2013). *Are some degrees worth more than others? Evidence from college admission cutoffs in Chile*. Tech. rep. National Bureau of Economic Research.

Hodler, Roland and Paul A Raschky (2014). "Regional favoritism." In: *The Quarterly Journal of Economics* 129.2, pp. 995–1033.

Ioannidis, John PA (2005). "Why most published research findings are false." In: *PLoS medicine* 2.8, e124.

Isaksson, Ann-Sofie and Andreas Kotsadam (2018). "Chinese aid and local corruption." In: *Journal of Public Economics* 159, pp. 146–159. ISSN: 0047-2727.

Jablonski, Ryan S. (2014). "How Aid Targets Votes: The Impact of Electoral Incentives on Foreign Aid Distribution." In: *World Politics* 66.02, pp. 293–330. ISSN: 0043-8871.

Johnsen, Bodil Wold (1999). "Fra universitetsvisjon til hyskoleintegrasjon." PhD thesis. Kristiansand: Hyskoleforlaget.

Kirkeboen, Lars J, Edwin Leuven, and Magne Mogstad (2016). "Field of study, earnings, and self-selection." In: *The Quarterly Journal of Economics* 131.3, pp. 1057–1111.

Knack, Stephen (2001). "Aid Dependence and the Quality of Governance: Cross-Country Empirical Tests." In: *Southern Economic Journal* 68.682, pp. 310–329.

Lindbekk, Tore (1998). "The education backlash hypothesis: The Norwegian experience 1960-92." In: *Acta Sociologica* 41.2-3, pp. 151–162.

Masaki, Takaaki (2018). "The political economy of aid allocation in africa: Evidence from zambia." In: *African Studies Review* 61.1, pp. 55–82.

Nunnenkamp, Peter, Hannes Öhler, and Maximiliano Sosa Andrés (2017). "Need, Merit and Politics in Multilateral Aid Allocation: A District-level Analysis of World Bank Projects in India." In: *Review of Development Economics* 21.1, pp. 126–156.

Ohler, Hannes and Peter Nunnenkamp (2014). "Needs-Based Targeting or Favoritism? The Regional Allocation of Multilateral Aid within Recipient Countries." In: *Kyklos* 67.3, pp. 420–446.

Ottoesen, Kristian (1969). *St.prp. nr. 136 Instilling om videreutdanning for artianere [Recommendation on further education for high school graduates].*

Qian, Nancy (2015). "Making progress on foreign aid." In: *Annu. Rev. Econ.* 7.1, pp. 277–308.

CHAPTER 1

The political economy of aid allocation: Aid and incumbency at the local level in Sub Saharan Africa

The political economy of aid allocation: Aid and incumbency at the local level in Sub Saharan Africa*

Tora Knutsen[1] and Andreas Kotsadam[2]

[1]University of Oslo. t.k.knutsen@econ.uio.no
[2]Ragnar Frisch Centre for Economic Research. andreas.kotsadam@frisch.uio.no

Abstract

Aid allocation within countries is often thought of as a strategic action by the incumbent leaders to further their own goals. Theoretically, however, the effects of aid may be either positive or negative and the empirical evidence is limited. By matching geocoded data on aid projects to 101 792 respondents in five waves of the Afrobarometer, we investigate the effects of aid on incumbency support using project fixed effects. We estimate the effects for World Bank aid and Chinese aid separately and find positive effects for the former and no robust effect for the latter. For neither project donor do we find effects on turnout and that aid is not targeting areas with previously higher incumbency support. We find little support for the notion that economic voting is driving the result as individuals self-perceived economic conditions are not affected. The positive effects for the World Bank aid projects seem to be mediated by trust in the politicians, whereas we find no effects of Chinese aid on trust.

Keywords: Aid; Africa; Politics; Development; Incumbency

*We would like to thank Ryan Briggs for very helpful and generous comments. In addition we thank two anonymous referees. Work on this study is supported by a grant from the Research Council of Norway (grant number. 250301)

1 Introduction

The debate around whether foreign aid has overall positive or overall negative consequences has been labeled one of the most controversial in development economics (Qian, 2015). Critics point to the potential adverse effects of aid on institutions and corruption (e.g. Easterly 2006; Deaton 2013; Knack 2001). One line of argument suggests that politicians use aid to distribute patronage or to influence voting (Jablonski, 2014). Yet, there is little evidence on whether the voters actually respond by increasing their support for the incumbent leader. We test this question by merging survey data on political preferences with geocoded data on aid projects for a large number of African countries.

Most previous studies investigate the political effects of aid using cross-sectional data at the country level. This literature provides mixed evidence, which is perhaps not surprising given the myriad of factors that are likely to affect country level institutions over time and the difficulty of identifying causal effects in cross-country regressions. The question of strategic allocation of aid projects for political gains is, furthermore, an inherently local one. We therefore follow the recent and growing trend towards analyzing the effects at the local level that has been spurred by the availability of geocoded data on both aid and outcomes.[1]

Several studies investigate the distribution of aid at the local level (Briggs, 2017; Dreher et al., 2019; Hodler and Raschky, 2014; Jablonski, 2014; Masaki, 2018; Nunnenkamp et al., 2017; Ohler and Nunnenkamp, 2014). In this study, we move beyond testing for biased aid allocation to an investigation of the political effects of targeting. Up until now, such studies have mostly been conducted using single countries (Briggs, 2012, 2014; Jablonski, 2014). A notable exception is Briggs (2018), who finds that aid lowers the support for incumbent

[1] See e.g. Kotsadam et al. (2018) on aid and infant mortality, Findley et al. (2011) on aid and conflict; Francken et al. (2012) on relief aid allocation in Madagascar; Powell and Findley (2012) on donor coordination; Dionne et al. (2013) on aid allocation in Malawi; Dreher and Lohmann (2015) on aid and growth at the regional level; Brazys et al. (2017) and Isaksson and Kotsadam (2018a) on the effects of aid on corruption; Isaksson and Kotsadam (2018b) on the effects of aid on unionization; and Berlin et al. (2017) on the effects of aid on gender outcomes in Malawi and Uganda.

presidents in a local level analysis of aid projects in three countries (Nigeria, Senegal, and Uganda). We add to this literature by conducting a study in many countries, using project level fixed effects, and by analyzing the effects of different donors.

The theoretical effects of aid on incumbency advantage are ex-ante ambiguous. Voters may reward incumbent leaders who strategically locate aid to their area, either if the leaders do it in order to please core supporters or to attract swing-voters. Incumbency support may also increase due to economic voting motives if aid is creating favorable economic conditions in the areas. Aid may also lead to lower support for the incumbent leader if it undermines the capacity and legitimacy of recipient governments (Knack, 2001) or reduces the political accountability (Ahmed, 2012). Furthermore, if aid has negative effects at the local level, or less positive effects than expected, it may also lower incumbency support.

To test for the effects of aid on incumbency support we use fixed effects for areas surrounding an aid project and investigate how the support changes when an aid project starts. We estimate the effects for World Bank aid and Chinese aid separately and find positive effects for the former and no robust effects for the latter. We find little indication of aid locations being selected based on previous incumbency support and we find no effects on overall electoral turnout. Furthermore, the effects of aid on incumbency support are not more positive when power is more contested, i.e. in countries and periods with higher political competition. A possible explanation for the differential impact of aid on incumbency support is that the effects are mediated by the effects on local living conditions. While such an explanation is compatible with previous research finding that economic conditions become better with World Bank aid (as measured by nighttime light) , while corruption (Isaksson and Kotsadam, 2018a) as well as favoritism (Dreher et al., 2019) increase with Chinese aid, we fail to corroborate this mechanism as we do not find that self-reported living conditions are differentially affected or that they seem to mediate the effects. Following Briggs (2018), we investigate the effects of aid on trust in Government and find that the positive effects for the World Bank aid projects seem to be mediated by trust in the politicians while we find no effect of Chinese aid on trust.

2 Aid and incumbency advantage

A positive relationship between incumbent support and aid location may arise both in cases where incumbent leaders locate aid to their own supporters, and when voters reward the incumbent after receiving aid. In the context of sub-Saharan Africa, patronage and favoritism are often considered important factors in shaping electoral outcomes. Some scholars claim that aid provides leaders with additional financial resources to distribute patronage, buy off political support, and ultimately to consolidate power (see e.g. Briggs (2012), Jablonski (2014), and Morrison (2009). Briggs (2015) finds a correlation between changes in aid before elections, and incumbent advantage in Africa suggesting that leaders potentially use aid to remain in power.

The empirical evidence on whether leaders allocate aid to please core supporters or to attract swing-voters is nevertheless still inconclusive. Hodler and Raschky (2014) find that higher aid inflows are associated with more regional favoritism. Ohler and Nunnenkamp (2014), find no evidence for the influence of favoritism in the location of World Bank projects in Africa. Likewise, Dreher et al. (2019) find no indications of favoritism in the location of World Bank projects in Africa, but they find evidence that regions where leaders were born receive more Chinese funded aid. Jablonski (2014) observes a strong bias toward constituencies with high vote shares for the incumbent in the aid allocation in Kenya. In contrast to the Kenyan case, in a study of the distribution of aid projects in Zambia, Masaki (2018) finds that fewer aid projects are allocated to districts where the ruling party enjoys greater popularity and suggest that aid in Zambia is used to attract swing-voters.

One potential mechanism for aid positively affecting voting is retrospective economic voting, i.e. the idea that voters reward incumbent leaders when the economic conditions are good (see Lewis-Beck and Stegmaier, 2007 for a review). The correlation between economic conditions and incumbency advantage is robust and has been found also in studies taking advantage of exogenous shocks such as the Spanish Christmas lottery (Bagues and Esteve-Volart, 2016), and variation in oil prices (Snowberg et al., 2007). Thus, those living close to an aid project may increase their support for the incumbent leader because their economic conditions are generally better, and they therefore prefer the status quo. The presence of economic voting hinges on aid having positive effects on economic well-being at the local level. The literature on this is, however, inconclusive. Using exogenous variation in Chinese aid supply following changes in Chinese steel production, Dreher et al. (2021) find that Chinese development assistance boosts economic growth in recipient sub-national regions. However, positive effects have not always been detected at the local level in studies using geocoded locations of aid projects. Isaksson and Kotsadam (2018) find that Chinese aid leads to increased corruption and has no effect on economic activity at the local level while aid from the World Bank increases economic activity.

There are both empirical and theoretical reasons to believe that aid can also lower the support for incumbent politicians. The first one is that aid may have negative consequences at the local level and that the incumbent politicians are punished for this. Another possible mechanism is that aid may alienate citizens from national politics and undermine the capacity and legitimacy of recipient governments (Knack, 2001). Even though this is plausible, there is very little empirical evidence pointing in this direction. For example, Blair and Roessler (2021) find no indications that Chinese aid has diminished state legitimacy. Similarly, Dietrich and Winters (2015) do not find an effect of individuals learning that a health project is foreign-funded in a survey experiment in India. Briggs (2018) is the only study to date that has investigated local level effects on incumbency support in several countries. He follows a similar strategy as in this paper but without project fixed effects and he only investigates the effects in a sample of three countries: Nigeria, Senegal and Uganda. Interestingly he finds that aid lowers the support for incumbent presidents. He further finds that aid leads to mistrust and favors a mechanism whereby aid fails to meet the expectations of the citizens.

3 Data

In order to analyze the effects of aid on incumbency support we spatially merge geocoded data on aid with geocoded data from the Afrobarometer, geocoded by BenYishay et al. (2017). The locations of the local areas are shown in Figure 1.

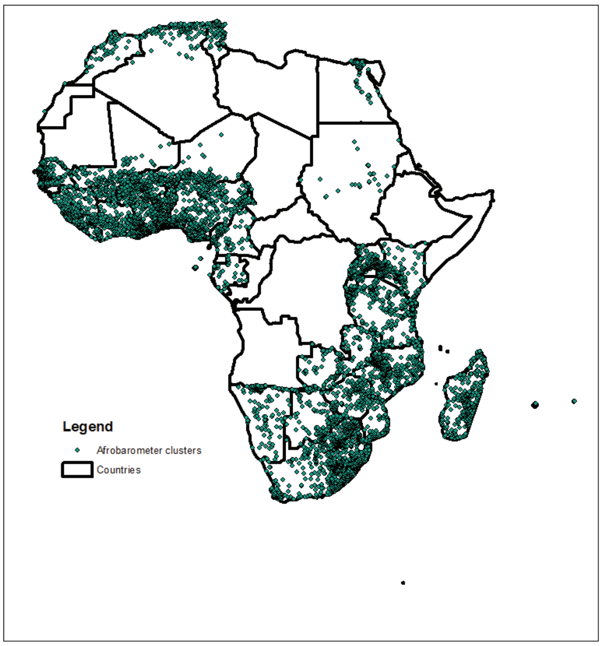

Figure 1: Location of Afrobarometer clusters

We present descriptive statistics for our main variables of interest in Table 1. The samples in this table are based on the baseline regressions which most importantly implies that they only include observations within 50 kilometers from an aid project.

Our measure of incumbency comes from the Afrobarometer and is available in rounds 2-6, albeit with a slightly different coding in round 2. In rounds 3-6 of the Afrobarometer survey the question asked is "If an election was held tomorrow which president (party) would you

Table 1: Descriptive statistics.

	World Bank		China	
	Mean	SD	Mean	SD
Dependent variables: Variants of incumbency advantage				
Incumbent	0.526	(0.499)	0.467	(0.499)
Incumbent 2	0.404	(0.491)	0.348	(0.476)
Incumbent 3	0.548	(0.498)	0.490	(0.500)
Aid variables				
Distance (km)	14.769	(12.683)	16.128	(15.237)
Active 50 km	0.908	(0.290)	0.662	(0.473)
Future aid 50 km	0.092	(0.290)	0.338	(0.473)
Active 25 km	0.702	(0.457)	0.507	(0.500)
Future aid 25 km	0.088	(0.283)	0.197	(0.398)
Control variables				
Urban	0.446	(0.497)	0.668	(0.471)
Age	36.189	(14.405)	35.665	(14.313)
Female	0.486	(0.500)	0.488	(0.500)
Other variables				
Electoral competition	40.867	(9.953)	41.123	(9.933)
Electoral democracy index	0.580	(0.140)	0.617	(0.153)
Turnout	0.806	(0.396)	0.773	(0.419)
Living conditions	−0.098	(1.003)	−0.021	(1.002)
N	40621		14983	

Notes: The samples are based on the baseline regressions (columns 1 in Tables 2 and 3) except for the variables Incumbent 2 and Incumbent 3 which are based on their corresponding baseline regressions.

vote for?". In addition to the different candidates, the respondents can also answer that they would not vote, that they do not know or that they would have voted for another candidate. In round 2 the question was "Do you feel close to any particular political party or political organization? If so, which party or organization is that?".

We always code all instances where the respondent answers the name of the incumbent president or party as one.[2] We define our dependent variable in three different ways to ensure that the results are not sensitive to how we define the missing and zero category in the dependent variable. In the baseline regressions we put zero on those that would vote for other parties than the incumbent and those that respond that they would not have voted in an election. We see that around 53 percent vote for the incumbent in the World Bank aid sample and 47 percent in the Chinese aid sample. In addition, we try two alternative definitions, measuring the incumbent's support at the extensive and the intensive margin. First, we include also those that respond that they do not know or do not want to respond as zeroes (Incumbent 2). Second, we look at the support at the intensive margin, by including those that would not vote in the missing category (Incumbent 3). This naturally increases the share of incumbency voters. The former is to be interpreted as the incumbent's support among those that answered, while the latter is to be interpreted as the support for the incumbent leader among those stating what they would have voted.

The data on Chinese aid projects is obtained from georeferenced project-level data of version 1.1 of AidData's Chinese Official Finance to Africa dataset, introduced by Strange et al. (2017) and geocoded by Dreher et al. (2016) (see Strange et al. (2013 and 2017) for a detailed description of the data collection methodology). Since this paper focuses on local effects of aid projects we focus on projects with recorded locations coded as corresponding to an exact location or as 'near', in the 'area' of, or up to 25 km away from an exact location

[2]We manually coded the names of the incumbents in all countries in the Afrobarometer and used the Afrobarometer codebooks for the various countries to match the answers.

(precision categories 1 and 2 in Strandow et al. 2011). We follow Isaksson and Kotsadam (2018a, 2018b) and limit our analysis to the Chinese aid projects that have been classified as official development assistance. The location of the Chinese aid projects are shown in Figure 2a.

The World Bank data is obtained from AidData (World Bank IBRD-IDA, Level 1, Version 1.4.1). We again limit the sample to projects with precise geocodes and information about start year, giving us 4,245 project locations. The location of the World Bank aid projects are shown in Figure 2b.

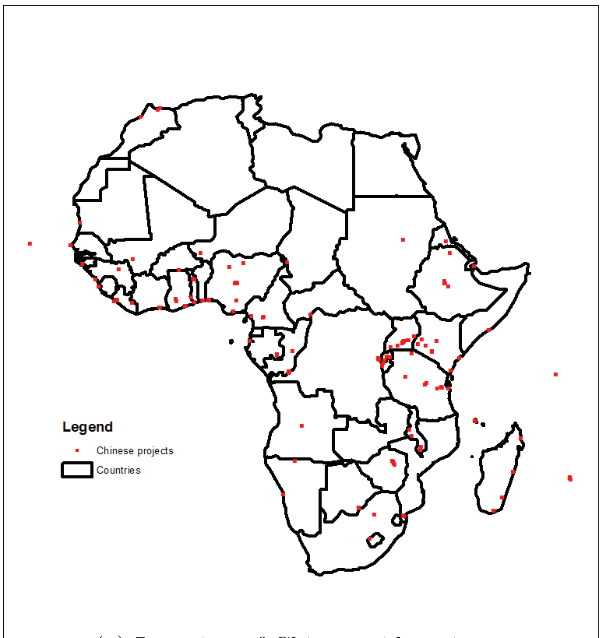

(a) Location of Chinese aid projects

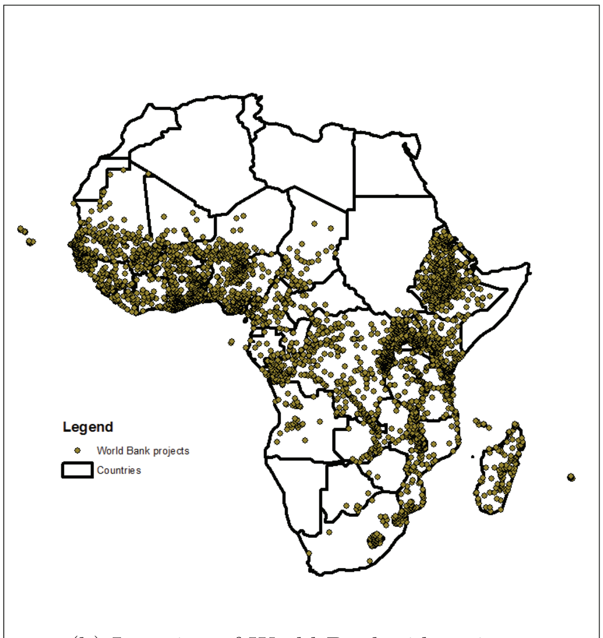

(b) Location of World Bank aid projects

Figure 2: Location of aid projects

In Table 1 we see that the average distance to a World Bank aid project is 15 kilometers and 16 kilometers to a Chinese aid project. Note again that the sample is restricted to individuals already having or that will get an aid project within 50 km.[3] For the World Bank sample we see that around 91 percent of the observations have an active project within

[3]We discuss in Section 7 how this sample restriction affects the external validity of our results.

50 kilometers while this number is 66 percent for Chinese aid. In the table we also present summary statistics for the control variables. These are a dummy variable for living in an urban area, a continuous variable for age, and a dummy variable for being female. We create a dummy variable for turnout (from the question ”Did you vote in the last election?”) and a variable for self-reported living conditions from the question ”In general, how do you rate your living conditions compared to those of other [people from the same country]” (ranging from Much worse (-2) via Same (0) to Much better (2)).

We also include country/year level measures of two other variables in order to explore mechanisms. We use the percent of votes received by the non-incumbent parties (Vanhanen, 2016), as a proxy for political competition. As a proxy for democratic elections we use the Electoral democracy index, from the V-dem dataset. This index is a weighted average of indices measuring freedom of expression, freedom of association and clean elections.

4 Empirical Strategy

To estimate the effect of aid projects on support for the incumbent leader, we match the GPS coordinates from the five Afrobarometer waves and the location of aid projects from AidData. We follow a similar strategy as the one in Knutsen et al. (2017) and Isaksson and Kotsadam (2018) by distinguishing between sites where an aid project is under implementation and sites where the project has yet to be implemented at the time of the Afrobarometer survey. We exclude respondents who live in areas where the implementation of a project has already been completed. Even though the Afrobarometer does not have a panel structure, it happens to revisit some places both before and after aid projects start in the area, and we are therefore able to use a fixed effects estimation strategy. We construct the geographical units by matching the geo-coordinates of aid projects to Afrobarometer clusters so that a respondent gets an indicator variable Active if an individual lives within 25 or 50 kilometers from an aid project.

The baseline regression connects the Afrobarometer respondents and location and timing of aid in the following way:

$$Incumbent_{it} = \beta_1 Active_{it} + \gamma_t + \alpha_i + \delta_1 X_{it} + \epsilon_{it} \tag{1}$$

The dependent variable is a dummy equal to 1 for individuals who report that they would vote for the incumbent leader at the time of the survey. Active indicates whether the individual lives within 50 or 25 kilometers of an active aid project at the time of the survey. If there are several projects nearby we code an area as soon as the first project starts.

To control for the average support for the incumbent across areas and across time, we include year,γ, and project area fixed effects α, in addition we also control for country fixed effects since the areas sometimes cross borders. We include a vector (X) of individual level controls from the Afrobarometer, these are gender, age, age-squared and a dummy for urban residence. These variables are included to control for compositional changes in the areas. We only include variables that are unlikely to be affected by aid. However, aid projects may change the composition of areas so we also present results without including these variables and the results are very similar (see Appendix Tables A2 and A3). The standard errors are clustered at the area level.

Without the area fixed effects, in order for β to have a causal interpretation we would need to assume that the introduction of an aid project is not correlated with the previous level of support for the incumbent leader or any other omitted variable that is correlated with both aid and voter preferences. This is a very strong assumption given that one of the mechanisms that we are testing for is that of strategic aid location, i.e. that the incumbent leader gives aid projects to his supporters in return for votes. With area fixed effects, however, we are able to control for previous levels of support as well as all other factors that are stable over time. The area fixed effects imply that identification comes solely from areas that are observed both before and after projects. As we discuss in Section 7, this has implications for the external validity of the results. A remaining identifying assumption is that there are no time varying omitted variables. One particular worry would be that there are trends in e.g. aid distribution and other political variables such as reforms or investments. To some extent we control for this with the inclusion of country and year fixed effects. Any such time varying variable would then have to evolve differently within countries and years. By comparing two different donors we would also be more worried if the results are very similar for them. If they are not, it would imply that the potential time varying factors differ across donors. Furthermore, we are able to test for selection of aid projects to areas where the support for the incumbent leader is high by excluding the area fixed effects and regressing Incumbent on Future, a dummy equal to one if the respondent lives in an area where there will be an aid project in the future. In addition, we also estimate regressions where we interact the year dummies with the country fixed effects and investigate if the effects are different.

To supplement the baseline regression, we explore theoretical mechanisms. First, we investigate whether the effect is different in countries where political competition is higher to understand whether the observed effect is due to aid being used deliberately as a tool to increase the support for the incumbent leader. The idea is that the incumbent leader is more interested in swing-voters when there exists some uncertainty regarding the electoral outcome. To explore an additional dimension of the electoral channel we examine whether the effect is more prevalent in countries with more democratic compared to less democratic elections.

5 Effects of aid on support for incumbents

We start by analyzing the effects of World Bank aid and in Table 2 we show the baseline results. Since we include area fixed effects, the coefficient for active can be interpreted directly as the increase in the support for the incumbent leader when an aid project starts.

We estimate this for individuals living close to aid projects and start by following the previous literature and use a 50 km cutoff in column 1. We see that the probability of supporting the incumbent politician increases by 5.5 percentage points in areas with active aid projects as compared to the same areas before the aid project had started. This effect is large and implies an increase of over 10 % from the mean. In Appendix Figure A1 we show that the effect seems to positive across the distances within 50 kilometers. In column 2 of Table 2 we present regression results with a 25 kilometers cutoff and the results are very similar (note that the sample is smaller as we are now restricting it to only contain individuals living within 25 km of an aid project).

A possible mechanism for the observed effect could be that the incumbent leaders give aid projects to areas that support them, and as a result they increase their support for them even more. In column 3 we test whether there is a selection in aid location by regressing support for incumbent on a dummy equal to 1 if the respondent lives in an area that will receive an aid project in the future. These regressions are run without the area fixed effects so that we can compare the support with a control group that is further away. This implies that the sample is larger and now includes respondents living further away. We also include a dummy variable for Active 50 km so that the Future 50 km coefficient only compares incumbency support in areas that will eventually get aid projects with areas that are further away than 50 kilometers to an aid project. We find no support for selection based on incumbency support as the coefficient on Future 50 km is imprecisely estimated and not significant.

Investigating the role of electoral competition shows that the coefficient on the interaction term Active*Electoral competition is negative and significant at the 10 percent level, implying that the effect of aid on support for the incumbent is weaker in countries with more political competition. This indicates that the incumbent effect does not seem to be driven by politicians locating aid projects to attract swing-voters.[4] Adding the Electoral democracy

[4]It is important to note that the sample is halved as the electoral competition variable is missing for some

index to the regression in (5) indicates that the effect is not larger in democratic countries as the coefficient on Active*Electoral democracy is not statistically significant. The coefficients on electoral competition and democracy are standardized with mean zero and standard deviation equal to 1 so that their size can be interpreted as the change in the dependent variable from a standard deviation increase in the regressors. Finally, we show that the effects of aid are not running via increased turnout in general (column 6), rather it seems to shift the support of the local electorate.

The results seem robust to various choices and definitions. We perform the same regressions using alternative missing categories, that is including those that do not know as 0 and excluding those that do not know or would not have voted (see Appendix Tables A6 and A7). The results are qualitatively similar to those in the baseline specification. We also test if the results are different in countries where the incumbent is not changing over the period and countries where they do.[5] Creating a dummy variable for being in a country without incumbency change and interacting active with this variable show that there is no statistically significant difference in the effects (Appendix Table A5). The samples are unfortunately too small to conduct meaningful analysis in single countries. We also test to add country times year fixed effects in Appendix Table A8 and note that the results point in the same direction. It seems, however, as if this is demanding too much of our data and the standard errors become larger. We cannot reject that the results are the same with and without these controls but neither is the effect statistically significant.

The results of the baseline regressions using Chinese aid projects as the independent variable are shown in Table 3. The main samples are now restricted to individuals living close to Chinese projects and Active in regressions 1-5 indicates that the respondent lives

countries throughout the period and some countries in particular years. However, we show in columns 1 and 2 of Table A4 that the baseline effects are larger in these samples.

[5]The countries without changes in incumbency are: Algeria, Botswana, Cameroon, Ethiopia, Gabon, Mozambique, Namibia, South Africa, Sudan, Swaziland, Tanzania, Uganda and Zimbabwe.

Table 2: World Bank aid and incumbency.

	(1) Incumbent	(2) Incumbent	(3) Incumbent	(4) Incumbent	(5) Incumbent	(6) Turnout
Active 50 km	0.055**			0.099*	0.11***	-0.010
	(0.028)			(0.053)	(0.031)	(0.019)
Future aid 50 km			0.0097			
			(0.015)			
Active 25 km		0.048				
		(0.031)				
Electoral competition				0.045		
				(0.052)		
Active*Electoral competition				-0.095*		
				(0.051)		
Electoral democracy index					0.23***	
					(0.058)	
Active*Electoral democracy					0.020	
					(0.031)	
Mean dep. var	0.53	0.50	0.54	0.51	0.52	0.79
No. of observations	40621	29306	64373	19052	30504	47301
R-squared	0.18	0.20	0.10	0.23	0.20	0.16
Project FE	Yes	Yes	No	Yes	Yes	Yes

Notes: All regressions control for country and year fixed effects, age, age squared, gender, and urban. Column three includes active 50 km as a control. Robust SE clustered at the project level in parentheses.

within 50 (1) or 25 (2) kilometers of an active aid project financed by China. The coefficient on Active now has the opposite sign, but the effect is only statistically significant in the 25 km sample. It should be noted that the effects of Chinese aid are identified based on only 46 projects and as such the precision is lower. Despite this we are, however, able to reject that the effects are more positive than 0.02 using an equivalence testing approach with two one-sided t-tests (TOST). We therefore conclude that the effects of Chinese aid seem different than the effects of World Bank aid, but they are less precisely estimated. In Appendix Figure A2 we further show that there is no clear relationship between the effects of Chinese aid and distance to the project. In addition Chinese aid seems to be located in areas where the incumbent leader already had less support before the initiation of an aid project (column 3). Adding proxies for political competition and democracy do not seem to alter the results substantially. The coefficient on Active drops when controlling for political competition, but it is too imprecisely estimated to draw any conclusions. Note that we are

Table 3: Chinese aid and incumbency.

	(1) Incumbent	(2) Incumbent	(3) Incumbent	(4) Incumbent	(5) Incumbent	(6) Turnout
Active 50 km	-0.074			-0.023	-0.12**	-0.018
	(0.057)			(0.063)	(0.062)	(0.027)
Future aid 50 km			-0.037***			
			(0.014)			
Active 25 km		-0.096***				
		(0.034)				
Electoral competition				-0.037		
				(0.046)		
Active*Electoral competition				-0.022		
				(0.059)		
Electoral democracy index					0.30**	
					(0.13)	
Active*Electoral democracy					0.013	
					(0.055)	
Mean dep. var	0.47	0.44	0.54	0.45	0.46	0.76
No. of observations	14983	10496	64823	7161	11626	16707
R-squared	0.13	0.11	0.11	0.20	0.14	0.18
Project FE	Yes	Yes	No	Yes	Yes	Yes

Notes: All regressions control for country, year, and project fixed effects, age, age squared, gender, and urban. Column three includes active 50 km as a control. Robust SE clustered at the project level in parentheses.

left with only around 7,000 observations in this regression and as seen in column 3 of Table A4, the baseline regression is not statistically significant in this sample either.[6]

6 Potential mechanisms for the different results across donors

One mechanism for the different results may be that Chinese and World Bank aid have differential impacts at the local level. Previous studies have indeed found differences whereby local economic conditions are improved more in areas receiving World Bank aid, and increased corruption as well as trade union memberships going down in areas with Chinese aid (Isaksson and Kotsadam, 2018a,b). One way to test this mechanism is to test whether there is a differential impact on peoples self-perceived economic conditions. In Table 4 we show that neither World Bank (column 1) nor Chinese aid (column 3) seem to affect how

[6] Appendix Tables A3, A9, A10,and column 2 of Table A5 show the same robustness checks for Chinese aid as those discussed for World Bank aid.

Table 4: Aid, self-reported living conditions, and incumbency.

	(1) Living conditions	(2) Incumbent	(3) Living conditions	(4) Incumbent
Active 50 km	0.072	0.058**	-0.019	-0.085
	(0.068)	(0.029)	(0.11)	(0.065)
Living conditions		0.034***		0.035***
		(0.0034)		(0.0060)
Mean dep. var	-0.10	0.52	-0.02	0.47
No. of observations	32524	32524	12217	12217
R-squared	0.11	0.20	0.05	0.13
Project donor	World Bank	World Bank	China	China

Notes: All regressions control for project, country and year fixed effects, age, age squared, gender, and urban. Robust SE clustered at the project level in parentheses.

voters perceive their economic conditions. We also see that the results remain more or less the same if we control for such perceptions and that perceptions themselves are positively correlated with incumbency support.

Another mechanism may be that aid affects trust in politicians and this may increase incumbent's support. We see in Table 5 that World Bank aid increases trust in the President, the parliament, and the ruling party, but not the trust in the opposition party. Chinese aid has no statistically significant effect on trust in politicians. We further see in Table 6 that trust in the president may actually mediate the positive World Bank results.[7] The mediation results should be interpreted with care, however, as we do not control for all factors that are correlated with the mediator and aid onset.

[7] In Appendix Tables A11 and A12 we show the same type of mediation analysis for the variables Trust Parliament and Trust ruling Party. We note that these variables also seem to mediate the effects of World Bank aid, but not completely.

Table 5: Aid and trust in politicians.

	(1) president	(2) president	(3) parliament	(4) parliament	(5) ruling	(6) ruling	(7) opposition	(8) opposition
Active 50 km	0.26***	0.052	0.12**	0.059	0.16**	0.047	-0.053	0.066
	(0.072)	(0.13)	(0.054)	(0.095)	(0.061)	(0.10)	(0.048)	(0.087)
Mean dep. var	1.83	1.68	1.61	1.49	1.58	1.43	1.18	1.17
No. of observations	39608	14624	38656	14467	39182	14606	38733	14431
R-squared	0.22	0.14	0.19	0.12	0.19	0.13	0.12	0.08
Project donor	World Bank	China	World Bank	China	World Bank	China	World Bank	China

Notes: All regressions control for project, country and year fixed effects, age, age squared, gender, and urban. Robust SE clustered at the project level in parentheses.

Table 6: Aid, trust in president, and incumbency.

	(1) Trust president	(2) Incumbent	(3) Trust president	(4) Incumbent
Active 50 km	0.26***	0.020	0.052	-0.074
	(0.072)	(0.024)	(0.13)	(0.046)
Trust president		0.15***		0.17***
		(0.0056)		(0.011)
Mean dep. var	1.83	0.53	1.68	0.47
No. of observations	39608	39608	14624	14624
R-squared	0.22	0.27	0.14	0.26
Project donor	World Bank	World Bank	China	China

Notes: All regressions control for project, country and year fixed effects, age, age squared, gender, and urban. Robust SE clustered at the project level in parentheses.

7 External validity

Our sample restrictions imply that we are only capturing a subset of the Chinese and World Bank aid projects. In particular, we restrict the projects to having precise geocodes and to have observations in the Afrobarometer both before and after project start. While the restrictions are described in the Data section, in this section we describe the effects of these restrictions on our final sample and discuss the implications this has for the external validity of our results.

Starting with the Chinese aid projects, we follow Isaksson and Kotsadam (2018a, 2018b) and limit the analysis to only ODA-like projects. This reduces the number of projects in the data from 1,955 to 1,272. Limiting the precision of the recorded project location to an exact location or as 'near', in the 'area' of, or up to 25 km away from an exact location (precision categories 1 and 2 in Strandow et al. 2011) reduces the number of projects to 491. As noted in Dreher and Lohmann (2015) and Isaksson and Kotsadam (2018a), the geographical coding precision is related to the sectoral composition of aid. We also restrict the projects to the ones having an actual start data, which reduces the number of projects to 227. Finally, restricting the projects to the ones that are driving identification in our design, i.e. the first project that opens in an area where we have observations both before and after project start, we are down to 46 projects.

We show the sectoral composition of Chinese aid in Table 7, for the full sample (Panel A) and our effective sample (Panel B) separately. We see that the three largest specified sectors for Chinese aid are "Government and Civil Society", "Education", and "Health". For the projects included in our effective sample the composition is different. While health is still a large sector, it's importance is actually even larger, we see that "Transportation and storage" accounts for around 30 percent of the project. Transportation projects account for less than 5 percent of the total Chinese aid.

The restrictions seem to have similar consequences for the representativeness of the World Bank projects. There are 1,702 projects in Africa without any restrictions and this reduces to 688 after restrictions on having precise geocodes. When restricting the projects to the ones in our effective sample we are left with 168 projects. We show the sectoral composition of World Bank aid in Table 7 and we note that we again clearly loose projects that relate to "Public Administration, Law, and Justice". As for the Chinese projects, the sector with most projects in the effective sample is "Transportation". In all, we note that the aid projects are not representative for aid projects overall.

In addition to examining the representativeness of the aid projects we can also investigate how representative our sample of *individuals* in the Afrobarometer are across samples. In investigating this we use the whole sample for which there are observations on incumbency and we create a dummy variable for being in our effective sample. In Table 9 we show how the individuals in the effective World Bank sample differ from the rest of the individuals. We see that there are few differences, and once we control for country and year fixed effects there is only a difference whereby individuals in the effective sample are less likely to live in an urban area. For Chinese aid, however, we see in Table 10 that there are more differences. Individuals in the effective Chinese aid sample show less support for the incumbent, are younger, and more likely to live in urban areas than individuals in the total sample when country and year fixed effects are controlled for. Hence, we conclude that individuals in the

Table 7: Sectoral division of Chinese aid projects

Panel A: All Chinese ODA-like aid projects		
Sector	Freq.	Percent
Action relating to debt	59	4.64
Agriculture, Forestry and Fishing	82	6.45
Banking and financial services	4	0.31
Business and other services	2	0.16
Communications	45	3.54
Developmental food aid	20	1.57
Education	146	11.48
Emergency response	72	5.66
Energy generation and supply	27	2.12
General budget support	3	0.24
Government and civil society	187	14.70
Health	186	14.62
Industry, Mining, Construction	11	0.86
Non-food commodity assistance	1	0.08
Other multisector	45	3.54
Other social infrastructure	58	4.56
Population policies programmes	11	0.86
Support to NGO	2	0.16
Trade and tourism	7	0.55
Transport and storage	63	4.95
Unspecified	211	16.59
Water supply and sanitation	24	1.89
Women in development	6	0.47
Total	1,272	100.00
Panel B: Projects in the effective sample		
Sector	Freq.	Percent
Agriculture, Forestry and Fishing	2	4.35
Education	4	8.70
Government and Civil Society	1	2.17
Health	14	30.43
Other multisector	2	4.35
Other social infrastructure	6	13.04
Transport and storage	14	30.43
Water supply and sanitation	2	4.35
Women in development	1	2.17
Total	46	100.00

Table 8: Sectoral division of World Bank aid projects

Panel A: All World Bank aid projects		
Sector	Freq.	Percent
Agriculture, Forestry and Fishing	159	9.34
Education	141	8.28
Energy and mining	131	7.70
Finance	54	3.17
Health and other social services	284	16.69
Industry and trade	70	4.11
Information and communications	26	1.53
Public Administration, Law, and Justice	548	32.20
Transportation	167	9.81
Water, sanitation and flood protection	122	7.17
Total	1,702	100.00
Panel B: Projects in the effective sample		
Sector	Freq.	Percent
Agriculture, Forestry and Fishing	21	12.50
Education	1	0.60
Energy and mining	52	30.95
Health and other social services	8	4.76
Industry and trade	7	4.17
Information and communications	4	2.38
Public Administration, Law, and Justice	8	4.76
Transportation	54	32.14
Water, sanitation and flood protection	13	7.74
Total	168	100.00

Chinese effective sample are different from the overall population. If these differences imply that there are effects in other areas or whether part of the differences are actually driven by aid is, however, difficult to know. In any case, it is important to note that the external validity of our results is limited.

Table 9: Comparing individuals in the effective sample to individuals in the total sample. World Bank aid.

	(1) Incumbent	(2) Incumbent	(3) Age	(4) Age	(5) Female	(6) Female	(7) Urban	(8) Urban
Effective sample	-0.0053	0.010	-0.44**	-0.027	0.0015	-0.0015	-0.0078	-0.041***
	(0.0065)	(0.0067)	(0.19)	(0.20)	(0.0065)	(0.0071)	(0.0065)	(0.0067)
No. of observations	68663	68663	68663	68663	68663	68663	68663	68663
R-squared	0.00	0.09	0.00	0.03	0.00	0.00	0.00	0.11
Country and Year FE	No	Yes	No	Yes	No	Yes	No	Yes

Table 10: Comparing individuals in the effective sample to individuals in the total sample. Chinese aid.

	(1) Incumbent	(2) Incumbent	(3) Age	(4) Age	(5) Female	(6) Female	(7) Urban	(8) Urban
Effective sample	-0.089***	-0.058***	-0.85***	-0.70***	0.0017	-0.0016	0.19***	0.19***
	(0.0051)	(0.0054)	(0.15)	(0.16)	(0.0051)	(0.0057)	(0.0051)	(0.0053)
No. of observations	68663	68663	68663	68663	68663	68663	68663	68663
R-squared	0.00	0.10	0.00	0.03	0.00	0.00	0.02	0.13
Country and Year FE	No	Yes	No	Yes	No	Yes	No	Yes

8 Conclusion

Aid allocation within countries is often thought of as a strategic action by the incumbent leaders to further their own goals. Theoretically, however, the effects of aid on incumbency support may be either positive or negative and the empirical evidence is limited.

In this paper we merge survey data on political preferences with geocoded data on aid projects for a large number of African countries. Using project fixed effects, we can compare incumbency support in the same communities close to projects before and after aid projects start and find that World Bank aid increases incumbent's support, whereas Chinese aid has no robust effect on this. In addition, we compare communities close to aid projects before they actually start with communities further away to test for selection of aid into incumbent-supporting areas. According to our results, aid is not targeting areas that had higher incumbency support before the aid project is initiated. There is no relationship between World Bank aid and initial incumbency support and Chinese aid is more likely to appear in areas with lower incumbency support.

Furthermore, we explore potential mechanisms mediating the positive effect on incumbent's support for World Bank aid. We find little support for the notion that economic voting is driving the result as individuals self-perceived economic conditions are not affected. For neither project donor do we find effects on turnout. The positive effects for the World Bank aid projects seem to be mediated by trust in the politicians, but we find no effects of Chinese aid on trust.

While we believe that our results have strong internal validity, this comes at the cost of potentially lower external validity. First of all, we do not know whether the findings of World Bank and Chinese projects generalize to other donors. Furthermore, using geocoded data for these two donors implies that not all of their aid is analyzed. To the extent that the effects of aid differ across sectors and across projects with and without precise geocodes we may not be able to speak to the effects of aid in general. Furthermore, by using project fixed effects we are leveraging variation from areas which we observe both before and after aid projects are implemented. We hope that future studies use similar methods as we do when more data becomes available. Future research could also use African Development Bank data to compare the effects of yet another multilateral donor in addition to the World Bank. For instance, following the work of Briggs (2019) who analyses aid project success across donors, researchers could analyze several outcomes so that we would reach a more generalized knowledge about the effects of aid.

References

Ahmed, Faisal Z (2012). “The perils of unearned foreign income: Aid, remittances, and government survival.” In: *American Political Science Review* 106.1, pp. 146–165.

Bagues, Manuel and Berta Esteve-Volart (2016). “Politicians’ Luck of the Draw: Evidence from the Spanish Christmas Lottery.” In: *Journal of Political Economy* 124.5, pp. 1269–1294. ISSN: 0022-3808.

Berlin, Maria Perrotta, Evelina Bonnier, and Anders Olofsgård (2017). *The donor footprint and gender gaps*. Tech. rep. World Institute for Development Economic Research (UNU-WIDER).

Blair, Robert A and Philip Roessler (2021). “The Effects of Chinese Aid on State Legitimacy in Africa: Cross-National and Sub-National Evidence from Surveys, Survey Experiments, and Behavioral Games.” In: *World Politics* 73.2, pp. 315–357.

Brazys, Samuel, Johan A Elkink, and Gina Kelly (2017). “Bad neighbors? How co-located Chinese and World Bank development projects impact local corruption in Tanzania.” In: *The Review of International Organizations* 12.2, pp. 227–253.

Briggs, Ryan C. (2012). “Electrifying the base? Aid and incumbent advantage in Ghana.” In: *The Journal of Modern African Studies* 50.4, pp. 603–624.

— (2014). “Aiding and abetting: project aid and ethnic politics in Kenya.” In: *World Development* 64, pp. 194–205.

— (2015). “The influence of aid changes on African election outcomes.” In: *International Interactions* 41.2, pp. 201–225.

— (2017). “Does Foreign Aid Target the Poorest?” In: *International Organization* 71.01, pp. 187–206. ISSN: 0020-8183.

— (2018). “Receiving foreign aid can reduce support for incumbent presidents.” In: *Political Research Quarterly*.

Briggs, Ryan C. (2019). "Results from single-donor analyses of project aid success seem to generalize pretty well across donors." In: *The Review of International Organizations*, pp. 1–17.

Dietrich, Simone and Matthew S Winters (2015). "Foreign aid and government legitimacy." In: *Journal of Experimental Political Science* 2.2, pp. 164–171.

Dreher, Axel, Andreas Fuchs, Roland Hodler, Bradley C Parks, Paul A Raschky, and Michael J Tierney (2019). "African leaders and the geography of China's foreign assistance." In: *Journal of Development Economics* 140, pp. 44–71.

Dreher, Axel, Andreas Fuchs, Bradley Parks, Austin Strange, and Michael J Tierney (2021). "Aid, China, and growth: Evidence from a new global development finance dataset." In: *American Economic Journal: Economic Policy* 13.2, pp. 135–74.

Dreher, Axel and Steffen Lohmann (2015). "Aid and growth at the regional level." In: *Oxford Review of Economic Policy* 31.3-4, pp. 420–446.

Findley, Michael G, Josh Powell, Daniel Strandow, and Jeff Tanner (2011). "The localized geography of foreign aid: A new dataset and application to violent armed conflict." In: *World Development* 39.11, pp. 1995–2009.

Francken, Nathalie, Bart Minten, and Johan FM Swinnen (2012). "The political economy of relief aid allocation: evidence from Madagascar." In: *World Development* 40.3, pp. 486–500.

Hodler, Roland and Paul A Raschky (2014). "Regional favoritism." In: *The Quarterly Journal of Economics* 129.2, pp. 995–1033.

Isaksson, Ann-Sofie and Andreas Kotsadam (2018a). "Chinese aid and local corruption." In: *Journal of Public Economics* 159, pp. 146–159. ISSN: 0047-2727.

— (2018b). "Racing to the bottom? Chinese development projects and trade union involvement in Africa." In: *World Development* Volume 106, pp. 284–298.

Jablonski, Ryan S. (2014). "How Aid Targets Votes: The Impact of Electoral Incentives on Foreign Aid Distribution." In: *World Politics* 66.02, pp. 293–330. ISSN: 0043-8871.

Knack, Stephen (2001). "Aid Dependence and the Quality of Governance: Cross-Country Empirical Tests." In: *Southern Economic Journal* 68.682, pp. 310–329.

Kotsadam, Andreas, Gudrun Østby, Siri Aas Rustad, Andreas Forø Tollefsen, and Henrik Urdal (2018). "Development aid and infant mortality. Micro-level evidence from Nigeria." In: *World Development* 105, pp. 59–69.

Lewis-Beck, Michael S and Mary Stegmaier (2007). "Economic models of voting." In: *The Oxford handbook of political behavior*.

Masaki, Takaaki (2018). "The political economy of aid allocation in africa: Evidence from zambia." In: *African Studies Review* 61.1, pp. 55–82.

Morrison, Kevin M (2009). "Oil, nontax revenue, and the redistributional foundations of regime stability." In: *International Organization* 63.1, pp. 107–138.

Nunnenkamp, Peter, Hannes Öhler, and Maximiliano Sosa Andrés (2017). "Need, Merit and Politics in Multilateral Aid Allocation: A District-level Analysis of World Bank Projects in India." In: *Review of Development Economics* 21.1, pp. 126–156.

Ohler, Hannes and Peter Nunnenkamp (2014). "Needs-Based Targeting or Favoritism? The Regional Allocation of Multilateral Aid within Recipient Countries." In: *Kyklos* 67.3, pp. 420–446.

Powell, Joshua and Michael G Findley (2012). "The Swarm Principle? A Sub-national Spatial Analysis of Donor Coordination in Sub-Saharan Africa." In.

Qian, Nancy (2015). "Making progress on foreign aid." In: *Annu. Rev. Econ.* 7.1, pp. 277–308.

Snowberg, Erik, Justin Wolfers, and Eric Zitzewitz (2007). "Partisan impacts on the economy: evidence from prediction markets and close elections." In: *The Quarterly Journal of Economics* 122.2, pp. 807–829.

Strange, Austin, Bradley Park, Michael J Tierney, Andreas Fuchs, Axel Dreher, and Vijaya Ramachandran (2013). "China's development finance to Africa: A media-based approach to data collection." In: *Center for Global Development Working Paper* 323.

Strange, Austin M, Axel Dreher, Andreas Fuchs, Bradley Parks, and Michael J Tierney (2017). "Tracking underreported financial flows: Chinas development finance and the aid–conflict nexus revisited." In: *Journal of Conflict Resolution* 61.5, pp. 935–963.

Vanhanen, Tatu (2016). "The polyarchy dataset: Vanhanen's index of democracy." In: *URL: http://www. prio. no/page/Project detail//9649/42472. html*.

Appendices

Appendix Figures

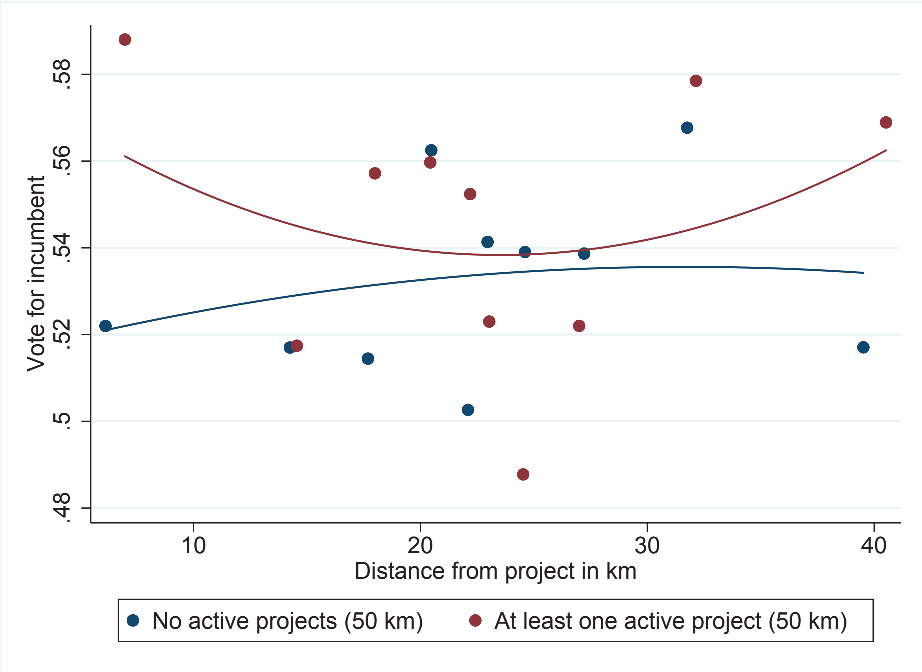

Figure A1: Distance to project and voting for incumbent. World Bank aid.

Notes: The figure shows a binned scatterplot by active status where the dots represent equal sized bins. The estimates are residualised for country, year, and project fixed effects, age, age squared, gender, and urban. The sample only includes areas that are already having or that will get an aid project within 50 km and excludes areas where aid has been suspended.

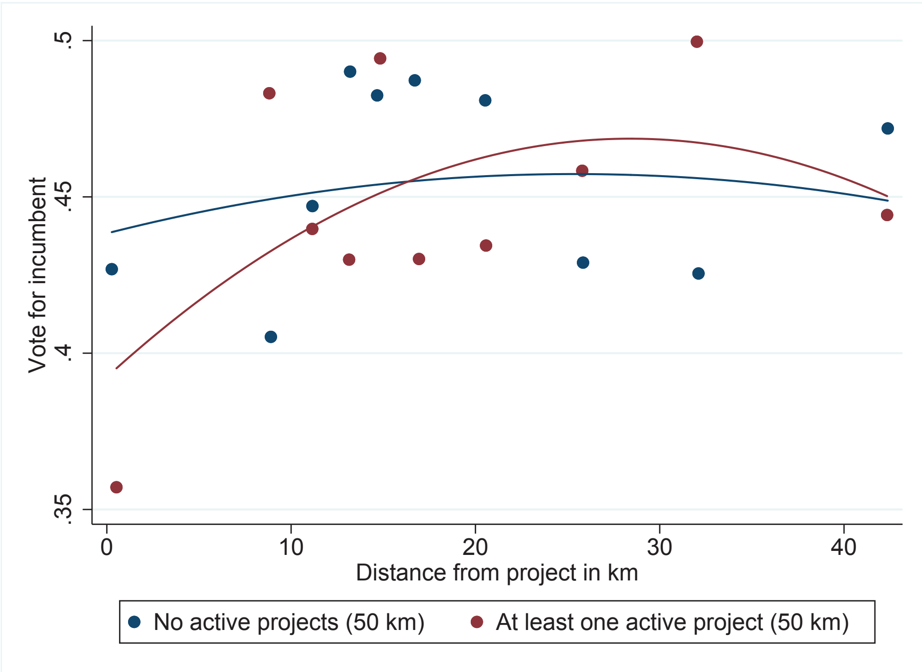

Figure A2: Distance to project and voting for incumbent. Chinese aid.
Notes: The figure shows a binned scatterplot by active status where the dots represent equal sized bins. The estimates are residualised for country, year, and project fixed effects, age, age squared, gender, and urban. The sample only includes areas that are already having or that will get an aid project within 50 km and excludes areas where aid has been suspended.

Appendix Tables

	Table A1: Variable descriptions	
Variable	**Description**	**Data Source**
$Incumbent_i$	The main outcome variable: A dummy equal to 1 for individuals who report that they would vote for the incumbent leader at the time of the survey and 0 for those that reported support to other candidates or did not want to vote. Otherwise missing	Afrobarometer

$Incumbent2_i$	Voting on the extensive margin: Different from *incumbent* in that also those refusing to answer the question and those that do not know are given the value 0.	Afrobarometer
$Incumbent3_i$	Voting on the intensive margin: Different from *incumbent* in that those not voting, refusing to answer or who do not know are given missing.	Afrobarometer
$Active_i$	A dummy equal to 1 if the respondent lives within 50 or 25 kilometers of an active aid project at the time of the Afrobarometer survey. Observations that have a completed project within the distance are excluded.	AidData
$Future_i$	Dummy equal to one if respondent lives in area that will receive an aid project in the future. Observations that have a completed project within the distance are excluded.	AidData
α_i	Project area fixed effect.	AidData
$Rural_i$	Dummy equal to one if respondent lives in a urban area as defined by the Afrobarometer	Afrobarometer
Electoral competition	A measure of electoral success of smaller parties. The variable is calculated by subtracting from 100 the percentage of votes won by the largest party (the party which wins most votes) in parliamentary elections or by the party of the successful candidate in presidential elections.	Varhanen(2016)
Electoral Democracy	An index which is a weighted average of indices measuring freedom of expression, freedom of association and clean elections.	V-dem dataset

Turnout	An indicator variable taking the value 1 if the respondent reports to have voted in the last election.	Afrobaromenter
Economic conditions	Self-reported living conditions as reported in the response to the question: "In general, how do you rate your living conditions compared to those of other [people from the same country]" (ranging from Much worse(-2) via Same(0) to Much better(2)).	Afrobarometer
Trust	Answer to the question: "How much do you trust each of the following [institution]?" (ranging from Not at all(0) via Just a little(1) , Somewhat(0) to A lot(3))	Afrobarometer

Table A2: Robustness without controls.

	(1) Incumbent	(2) Incumbent
Active 50 km	-0.079	
	(0.058)	
Active 25 km		-0.097
		(0.084)
Mean dep. var	0.47	0.44
No. of observations	15083	10545
R-squared	0.12	0.10
Project FE	Yes	Yes

Notes: All regressions control for country, year, and mine fixed effects. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A3: Robustness without controls: Chinese aid.

	(1) Incumbent	(2) Incumbent
Active 50 km	-0.079	
	(0.058)	
Active 25 km		-0.097***
		(0.033)
Mean dep. var	0.47	0.44
No. of observations	15083	10545
R-squared	0.12	0.10
Project FE	Yes	Yes

Notes: All regressions control for country, year, and project fixed effects. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A4: Different samples for the interaction analysis.

	(1) Incumbent WB	(2) Incumbent WB	(3) Incumbent China	(4) Incumbent China
Active 50 km	0.14***	0.12***	-0.033	-0.089
	(0.049)	(0.031)	(0.055)	(0.063)
Mean dep. var	0.51	0.52	0.45	0.46
No. of observations	19052	30504	7161	11626
R-squared	0.23	0.19	0.20	0.13
Project FE	Yes	Yes	Yes	Yes

Notes: All regressions control for country, year, and project fixed effects, age, age squared, gender, and urban. Robust SE clustered at the Afrobarometer cluster level in parentheses. Columns 1 and two show results for world bank aid whereas columns 3 and 4 show results for Chinese aid.

Table A5: Different effects on countries with stable incumbents.

	(1) Incumbent	(2) Incumbent
Active 50 km	0.073*	-0.086
	(0.041)	(0.067)
Countries without change in incumbency	0.16	0.24***
	(0.21)	(0.063)
Active*No change in incumbency	-0.034	0.040
	(0.043)	(0.069)
Mean dep. var	0.53	0.47
No. of observations	40621	14983
R-squared	0.18	0.13
Project donor	World Bank	China

Notes: All regressions control for country, year, and project fixed effects, age, age squared, gender, and urban. Robust SE clustered at the Afrobarometer cluster level in parentheses. Column 1 and two show results for world bank aid whereas columns 2 shows results for Chinese aid.

Table A6: Aid and incumbency 2: World Bank aid.

	(1) incumbent2	(2) incumbent2	(3) incumbent2	(4) incumbent2	(5) incumbent2
Active 50 km	0.070***			0.099*	0.11***
	(0.025)			(0.053)	(0.029)
Future aid 50 km			-0.010		
			(0.014)		
Active 25 km		0.060**			
		(0.027)			
Electoral competition				0.045	
				(0.052)	
Active*Electoral competition				-0.095*	
				(0.051)	
Electoral democracy index					0.20***
					(0.047)
Active*Electoral democracy					0.048
					(0.030)
Mean dep. var	0.40	0.38	0.42	0.51	0.52
No. of observations	52917	38308	83864	19052	38884
R-squared	0.14	0.15	0.08	0.23	0.15
Project FE	Yes	Yes	No	Yes	Yes

Notes: All regressions control for country and year fixed effects, age, age squared, gender, and urban. Column three includes active 50 km as a control. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A7: Aid and incumbency 3: World Bank aid.

	(1) incumbent3	(2) incumbent3	(3) incumbent3	(4) incumbent3	(5) incumbent3
Active 50 km	0.066**			0.099*	0.13***
	(0.029)			(0.053)	(0.032)
Future aid 50 km			0.0098		
			(0.016)		
Active 25 km		0.062*			
		(0.032)			
Electoral competition				0.045	
				(0.052)	
Active*Electoral competition				-0.095*	
				(0.051)	
Electoral democracy index					0.22***
					(0.059)
Active*Electoral democracy					0.047
					(0.031)
Mean dep. var	0.55	0.52	0.57	0.53	0.55
No. of observations	38988	28029	61580	19052	29056
R-squared	0.19	0.21	0.10	0.23	0.21
Project FE	Yes	Yes	No	Yes	Yes

Notes: All regressions control for country, year, and project fixed effects, age, age squared, gender, and urban.Column three includes active 50 km as a control. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A8: Aid and incumbency with country specific year fixed effects: World Bank aid.

	(1)
	Incumbent
Active 50 km	0.014
	(0.031)
Mean dep. var	0.53
No. of observations	40621
R-squared	0.22
Project FE	Yes
Country*Year FE	Yes

Notes: The regression controls for country times year fixed effects in addition to project fixed effects, age, age squared, gender, and urban. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A9: Aid and incumbency 2: Chinese aid.

	(1)	(2)	(3)	(4)	(5)
	incumbent2	incumbent2	incumbent2	incumbent2	incumbent2
Active 50 km	-0.020			-0.012	-0.065
	(0.049)			(0.059)	(0.057)
Future aid 50 km			-0.037***		
			(0.013)		
Active 25 km		-0.031			
		(0.031)			
Electoral competition				-0.029	
				(0.038)	
Active*Electoral competition				-0.013	
				(0.055)	
Electoral democracy index					0.25**
					(0.11)
Active*Electoral democracy					0.0068
					(0.047)
Mean dep. var	0.35	0.32	0.41	0.45	0.46
No. of observations	20094	14504	85069	8588	14983
R-squared	0.09	0.07	0.09	0.17	0.11
Project FE	Yes	Yes	No	Yes	Yes

Notes: All regressions control for country, year, and project fixed effects, age, age squared, gender, and urban.Column three includes active 50 km as a control. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A10: Aid and incumbency 3: Chinese aid.

	(1) incumbent3	(2) incumbent3	(3) incumbent3	(4) incumbent3	(5) incumbent3
Active 50 km	-0.060			-0.023	-0.11
	(0.065)			(0.063)	(0.071)
Future aid 50 km			-0.038***		
			(0.014)		
Active 25 km		-0.084**			
		(0.036)			
Electoral competition				-0.037	
				(0.046)	
Active*Electoral competition				-0.022	
				(0.059)	
Electoral democracy index					0.33**
					(0.13)
Active*Electoral democracy					0.014
					(0.063)
Mean dep. var	0.49	0.47	0.56	0.47	0.49
No. of observations	14300	9978	62104	7161	10979
R-squared	0.13	0.11	0.10	0.20	0.14
Project FE	Yes	Yes	No	Yes	Yes

Notes: All regressions control for country, year, and project fixed effects, age, age squared, gender, and urban.Column three includes active 50 km as a control. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A11: Aid, trust in parliament, and incumbency.

	(1) Trust president	(2) Incumbent	(3) Trust president	(4) Incumbent
Active 50 km	0.12**	0.044*	0.059	-0.077
	(0.054)	(0.026)	(0.095)	(0.057)
Trust parliament		0.098***		0.10***
		(0.0041)		(0.011)
Mean dep. var	1.61	0.53	1.49	0.47
No. of observations	38656	38656	14467	14467
R-squared	0.19	0.22	0.12	0.16
Project donor	World Bank	World Bank	China	China

Notes: All regressions control for project, country and year fixed effects, age, age squared, gender, and urban. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Table A12: Aid, trust in ruling party, and incumbency.

	(1) Trust president	(2) Incumbent	(3) Trust president	(4) Incumbent
Active 50 km	0.16**	0.036	0.047	-0.072
	(0.061)	(0.023)	(0.10)	(0.048)
Trust ruling party		0.14***		0.16***
		(0.0053)		(0.015)
Mean dep. var	1.58	0.53	1.43	0.47
No. of observations	39182	39182	14606	14606
R-squared	0.19	0.26	0.13	0.24
Project donor	World Bank	World Bank	China	China

Notes: All regressions control for project, country and year fixed effects, age, age squared, gender, and urban. Robust SE clustered at the Afrobarometer cluster level in parentheses.

Coding of incumbency

Round	Country	Q type[8]	Year of survey	Incumbent
5	Algeria	2	2013	National Liberation Front (FLN)
6	Algeria	2	2015	FLN
3	Benin	2	2005	Kerekou
4	Benin	2	2008	Independent
5	Benin	2	2011	Boni
6	Benin	2	2014	Boni
2	Botswana	4	2003	Botswana Democratic Party (BDP)
3	Botswana	2	2005	BDP
4	Botswana	2	2008	BDP
5	Botswana	2	2012	BDP
6	Botswana	2	2014	BDP
4	Burkina Faso	2	2008	Congress for Democracy and Progress (CDP)
5	Burkina Faso	2	2012	CDP
6	Burkina Faso	2	2015	
5	Burundi	2	2012	National Council for the Defense of Democracy (CNDD-FDD)

[8] Question asked in survey: 1. What did you vote? 2. What would you vote in presidential election? 3. What would you vote in a national election (president and prime minister)? [In Lesotho only prime minister]. 4. Do you feel close to any political part, if so which?

6	Burundi	2	2014	CNDD-FDD
5	Cameroon	2	2013	Cameroon People's Democratic Movement (RDPC)
6	Cameroon	2	2015	RDPC
2	Cape Verde	4	2002	African Party for the Independence of Cape Verde (PAICV)
3	Cape Verde	2	2005	PAICV
4	Capo Verde	3	2008	PAICV
5	Capo Verde	2	2011	Movement for Democracy (MpD)
6	Capo Verde	2	2014	MpD
5	Côte d'Ivoire	2	2013	Rally of the Republicans (RDR)
6	Côte d'Ivoire	2	2014	RDR
5	Egypt	2	2013	Freedom and Justice Party (FJP)
6	Egypt	2	2015	
5	Ethiopia	3	2013	Ethiopian People's Revolutionary Democratic Front (EPRDF)
6	Gabon	2	2015	Gabonese Democratic Party (PDG)
2	Ghana	4	2002	New Patriotic Party (NPP)
3	Ghana	2	2005	NPP

4	Ghana	2	2008	NPP
5	Ghana	2	2012	National Democratic Congress (NDC)
6	Ghana	2	2014	NDC
5	Guinea	2	2013	Guinean People's Assembly (RPG)
6	Guinea	2	2015	RPG
2	Kenya	4	2003	National Rainbow Coalition/ Liberal Demoratic Party (NARC/LDP)
3	Kenya	2	2005	NARC
4	Kenya	2	2008	Party of National Unity (PNU)
5	Kenya	2	2011	PNU
6	Kenya	2	2014	The National Alliance (TNA)
2	Lesotho	4	2003	Lesotho Congress for Democracy (LCD)
3	Lesotho	3	2005	LCD
4	Lesotho	3	2008	LCD
5	Lesotho	2	2012	LCD
6	Lesotho	2	2014	All Basotho Convention (ABC)
4	Liberia	2	2008	Unity Party (UP)
5	Liberia	2	2012	UP

6	Liberia	2	2015	UP
3	Madagascar	2	2005	Tiako I Madagasikara (TIM)
4	Madagascar	2	2008	TIM
5	Madagascar	2	2013	TIM
6	Madagascar	2	2014/2015	New Forces for Madagascar (HVM)
2	Malawi	4	2002	United Democratic Front (UDF)
3	Malawi	2	2005	UDF
4	Malawi	2	2008	Democratic Progressive Party (DPP)
5	Malawi	2	2012	People's Party (PP)
6	Malawi	2	2014	PP
2	Mali	4	2002	Touré/Mouvement
3	Mali	2	2005	Touré/Mouvement
4	Mali	2	2008	
5	Mali		2012	
6	Mali	2	2014	Rally for Mali (RPM)
5	Mauritius	3	2012	Labour
6	Mauritius	2	2014	Labour
5	Morocco	3	2013	Justice and Development Party (PJD)
6	Morocco	2	2015	PJD

2	Mozambique	4	2002	The Mozambique Liberation Front (FRELIMO)
3	Mozambique	2	2005	FRELIMO
4	Mozambique	3	2008	FRELIMO
5	Mozambique	2	2012	FRELIMO
6	Mozambique	2	2015	FRELIMO
2	Namibia	4	2003	South West Africa People's Organization (SWAPO)
3	Namibia	2	2006	SWAPO
4	Namibia	2	2008	SWAPO
5	Namibia	2	2012	SWAPO
6	Namibia	2	2014	SWAPO
5	Niger	2	2013	Nigerien Party for Democracy and Socialism (PNDS)
6	Niger	2	2014	PNDS
2	Nigeria	4	2003	PDP
3	Nigeria	2	2005	PDP
4	Nigeria	2	2008	PDP
6	Nigeria	2	2014	PDP
5	Nigeria	2	2012	PDP
6	Sao Tome and Principe	2	2015	Independent Democratic Action (ADI)
2	Senegal	4	2002	Senegalese Democratic Party PDS
4	Senegal	2	2008	PDS

5	Senegal	2	2013	Alliance for the Republic (APR)
6	Senegal	2	2014	APR
3	Senegal	2	2005	PDS
5	Sierra Leone	2	2012	All People's Congress (APC)
6	Sierra Leone	2	2015	APC
2	South africa	4	2002	African National Congress (ANC)
3	South Africa	3	2006	ANC
4	South Africa	2	2008	ANC
5	South Africa	3	2011	ANC
6	South Africa	2	2015	ANC
5	Sudan	2	2013	National Congress (NC)
6	Sudan	2	2015	NC
5	Swaziland	2	2013	
6	Swaziland	2	2015	
2	Tanzania	4	2003	Chama Cha Mapinduzi (CCM)
3	Tanzania	2	2005	CCM
4	Tanzania	2	2008	CCM
5	Tanzania	2	2012	CCM
6	Tanzania	2	2015	CCM
5	Togo	2	2012	Union for the Republic (UNIR)

6	Togo	2	2014	UNIR
5	Tunisia	2	2013	Congress for the Republic (CPR)
6	Tunisia	2	2015	Nidaa Tounes
2	Uganda	4	2002	National Resistance Movement (NRM)
3	Uganda	2	2005	NRM
4	Uganda	2	2008	NRM
5	Uganda	2	2011/2012	NRM
6	Uganda	2	2015	NRM
2	Zambia	4	2003	Movement for Multi-party Democracy (MMD)
3	Zambia	2	2005	MMD
4	Zambia	2	2009	MMD
5	Zambia	2	2013	Patriotic Front (PF)
6	Zambia	2	2014	PF
2	Zimbabwe	4	2004	The Zimbabwe African National Union Patriotic Front (ZANU-PF)
3	Zimbabwe	2	2005	ZANU-PF
4	Zimbabwe	2	2009	ZANU-PF
5	Zimbabwe	2	2012	ZANU-PF
6	Zimbabwe	2	2014	ZANU-PF

Notes by Country

Below are country specific comments on choices made when coding *Incumbent*.

Benin: The president between 2006 and 2011, Boni, was an independent candidate. In the 2007 parliamentary elections, Boni supported a coalition called FCBE, which got the largest share of votes. Therefore, we consider FCBE as the incumbent party.

Burkina Faso: President Blaise Compaoré faced a coup in October 2014. The Afrobarometer carried out a survey in the following year. In 2015, Burkina Faso was ruled by a transitional president, Michel Kafando. He was not from a political party. Therefore, respondents from Burkina Faso in the sixth round of the survey have missing on all the incumbent variables.

Cape Verde: Is a semi-presidential representative democracy where the president appoints the prime minister. From 2011 to 2015, the prime minister and the president were from different parties. Until round four, the Afrobarometer asked about prime minister in Cape Verde, while in round five and six, they asked about the president. Therefore, we consider the president as the incumbent in round five and six, and the party of the prime minster as the incumbent in the prior rounds. Cape Verde had an election in August 2011, the same year as the fifth round was carried out. In this election, Movement for Democracy's (MpD) candidate won against the candidate from the incumbent party. As the Afrobarometer carried out the survey in December 2011, we consider MpD the incumbent party in round five.

Egypt: Muhammed Mursi was removed in a coup in July 2013. Afrobarometer carried out the fifth survey round in March 2013. Therefore, Mursi was still incumbent at the time of the survey. In 2015, the year of the sixth round, Abdel Fattah al-Sisi was the sitting president. He had no party affiliation, and was not a response option in the Afrobarometer. Therefore, we give respondents from Egypt the value missing on the incumbent variables in this round

Ethiopia:Afrobarometer carried out a survey in Ethiopia only in 2013. However, Afrobarometer raises questions on the comparability of Ethiopian results with those form other surveyed countries, in particular with regard to attitudes toward democracy. Therefore, Ethiopia is not included in in the multi-country dataset. We also do not include Ethiopia.

Kenya: Kenya African National Union (KANU) had held the presidency uninterrupted since Kenyan independence. In 2002, The National Rainbow Coalition (NARC), a coalition consisting of Liberal Democratic Party (LDP) and National Alliance Party of Kenya (NAK) among others, won the presidential election. Their presidential candidate was Mwai Kibaki, from Democratic Party which was now a part of NAK. The coalition dissolved in 2005 due to a disagreement regarding the division of power between the president and the prime minister. The disagreement led to a vote on a proposed change to the constitution. While, Kibaki and NAK led the yes-campaign, LDP led the no-campaign. The yes- campaign later established the new party Party of National Unity (PNU), while the no-campaigners created Orange Democratic Movement (ODM). Thus, we consider NARC and LDP to be the incumbent up until 2005. As LDP broke out of the coalition in 2005, we do not consider LDP to be incumbent in 2005. The other parties that joined NARC in 2002 are not separate response alternatives. In the election in 2007, Kibaki, now in PNU and this time supported by KANU, run against Odinga and ODM. Kibaki claimed victory, but the result was contested and violence broke out in Kenya. Kibaki and Odinga achieved a diplomatic solution in 2008, forming a coalition government with Kibaki as president and Odinga as the prime minister. As the Afrobarometer in Kenya does ask about presidential candidate, we consider PNU and Kibaki as the incumbent president in this time period. The coalition between PNU and ODM held power until 2013 when Uhuru Kenyatta won the election, this time as the leader of The National Alliance (and later the Jubilee alliance).

Malawi: President Mutharika (Democratic Progressive Party) died suddenly in 2012.

Joyce Banda (People's Party) succeeded him and took office the in April. The fifth round of the Afrobarometer survey was carried out in June/July 2012, so Banda was incumbent at the time of the survey. In May 2014, the year of round six, Malawi had an election. Due to the election result, People's Party (PP) handed power to Democratic Progressive Party (DPP). Nevertheless, PP was still in power when Afrobarometer surveyed Malawi before in April/March.

Mali: In June 2002, the year of the second round of the Afrobarometer, Mali had a change in president. As Mali was surveyed in November and December, the new president, Touré, was the incumbent leader at the time of the survey. He ran as an independent candidate, but was supported by a coalition called Movement citoyen. In the spring of 2012, the year of round five, Mali had two different presidents due to a coup in March where the incumbent president agreed to resign. In the Afrobarometer, the overthrown president is not a response option (Touré). Therefore, we give respondents in Mali the value missing for incumbent in this round.

Swaziland: Is not a part of the analysis in this paper, as the survey has been conducted multiple rounds without asking respondents about voting.

CHAPTER 2

Distance and choice of field: Evidence from a Norwegian college expansion

Distance and choice of field: Evidence from a Norwegian college expansion *

Tora K. Knutsen† Jørgen Modalsli‡ Marte Rønning§

Abstract

How can geographical proximity to college explain field of study choices? We empirically address this question using the major expansion of university colleges in Norway in the second half of the twentieth century, when 33 new higher education institutions were established in areas that did not previously have access to higher education. Our findings indicate that take-up of the relevant educations (nursing, engineering and business administration) increased substantially with the establishment of new colleges. However, we do not find evidence of an increase in education or earnings capacity overall, suggesting that the new colleges shifted individuals on the intensive rather than extensive margin, between education tracks of similar length. We discuss challenges related to the estimation of education choices in a population that often started higher education late, well into their twenties, and also document that traditional gender differences in educational fields persisted.

Keywords: University access, Gender wage gap; Field of study; Family background; Geospatial variation.

JELcodes: D31; I23; J62

*We wish to thank Jo Thori Lind and Stefan Leknes as well as seminar participants at 30th conference of the European Association of the Labour Economists for comments and suggestions. Support from the Norwegian Research Council (grant no. 237840) is acknowledged.

†Department of Economics, University of Oslo, `t.k.knutsen@econ.uio.no`

‡Oslo Business School at Oslo Metropolitan University and Statistics Norway, `jorgenmo@oslomet.no`

§Statistics Norway, `marte.ronning@ssb.no`

1 Introduction

Geraphy and place are important in shaping opportunity (Chetty et al., 2014; Markussen and Roed, 2018), and living near a higher education institution is associated with higher educational attainment (Card, 1995; Helland and Heggen 2018). However, less is known about how proximity to education institution shapes the choice of field of study. In this paper, we study how geographical expansion of higher education institutions in Norway during the second half of the twentieth century affected outcomes - in terms of choice of field of study, education length and participation in the labor market as adult - for individuals growing up close to these new college establishments. Between 1955 and 1989 a total of 33 new university colleges were established in Norway, with wide geographical dispersion, predominantly offering higher degrees in nursing, engineering, teaching and business administration. Pre 1950, higher education institutions in Norway were mainly located in the vicinity of the largest cites.

As in other Western countries, higher education in Norway was subject to a very rapid expansion beginning in the 1960s, and developing through the 1970s and 1980s. The increase in education institutions across the whole country increased the take-up of the degrees provided at the newly established colleges. Figure 1 shows that the overall share of each cohort that obtained a degree in nursing and engineering increased substantially for cohorts born between 1950 and 1960.[1] The increase in business continued for cohorts born between 1960 and 1970. These increases coincide with the regionally staggered expansion of university colleges, on which we base our identification strategy.

Using rich administrative data on education and earnings as well as censuses going back to 1960, we are able to track field of study, earnings and the municipality of residence in a given year for the entire Norwegian population. Our findings indicate that young adults residing within commuting distance to a new college show a significant rate of take-up of the new degrees being offered. However, men do not respond to the opportunity to take a nursing degree, and women do not respond to engin-

[1] Other historical sources tell similar stories. The documentation of censuses 1960, 1970 and 1980 (Vassenden, 1987) contains statistics on education. In the 1960 census education from type of institution was registered. In the 1970 and 1980-census, type of education was registered. As percentage of the population, the share of nurses increased from 0.59 % in 1960 to 0.98 % in 1970 and 1.3 % in 1980 . The share of college engineers increased from 0.5% in 1960 to approximately 1 % in 1980. Hence, the number of nurses and college engineers doubled as a percentage of the population between 1960 and 1980.

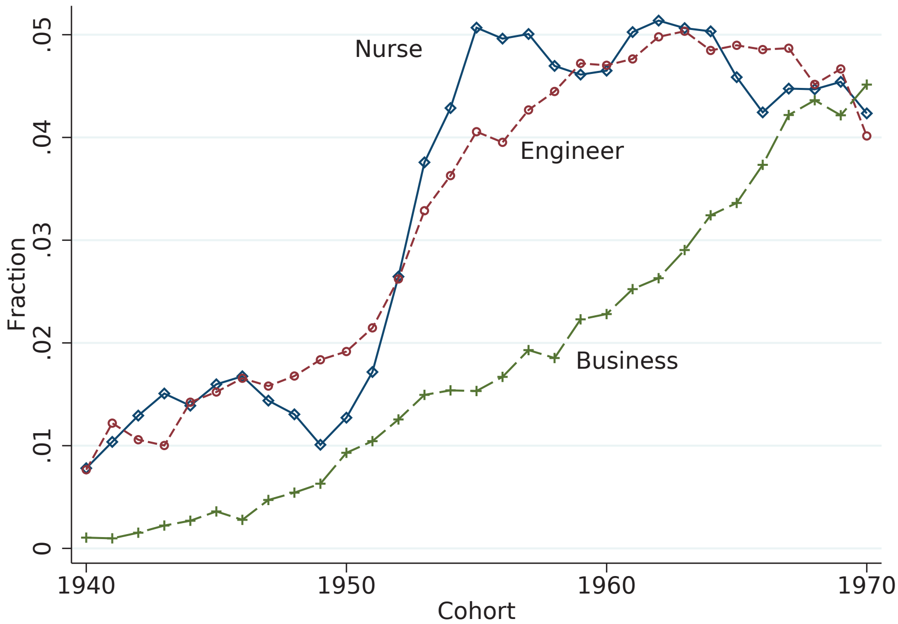

Figure 1. Fraction with a degree in nursing, engineering and business in Norway, by birth cohorts 1940-1973.

eering (STEM) degrees. This fits in with the general pattern seen in Norway and elsewhere, women often choose health related fields and work in the public sector, while men choose STEM related fields (Card and Payne, 2020). For business administration, the take up rates do not differ across men and women.

We find no increase in the overall educational attainment, implying that the local expansion of these new college degrees came at the expense of other (and more established) degrees at the same level not offered locally. The establishment of a nursing college is associated with a decrease in the propensity to do a teaching degree for women. For men, the new engineering colleges are negatively correlated with take-ups of degrees in other technological fields and business administration. Furthermore, when the new business administration degree was introduced, men chose it at the expense of engineering, while for women there are indications that the increase in this degree came at the expense of a degree in teaching and social sciences.

Although the college premium is important in explaining wage inequality, it is by now well-documented that field of study also matters for labour market outcomes including the gender wage gap (Altonji et al., 2012; Hastings et al., 2013; Kirkeboen et al., 2016). While higher education is associated with higher earnings, the effect of a shift in field is not clear. If people choose field according to their comparative advantage as found in Kirkeboen et al. (2016), decreasing the cost of particular fields may induce people to not choose optimally. We find few pronounced changes in mid-life income and labor market participation, suggesting that the changes in field of study for the treated population did not result in higher earnings. However business administration is one exception: While we find zero or small positive changes in male wages, we do observe a decrease in the labor market participation for women living in areas with a newly established college offering degrees in business administration. We ask two questions in pursuit of plausible explanations for this result. First, to what extent did the establishment of a college induce women to remain in their home municipality? If they did, they could have missed out on better labor market opportunities elsewhere. Second, did the business administration degree offer particularly bad labor market opportunities for women? We find some support for both explanations: The college establishments had an effect in retaining women in their home municipality, while we find a smaller effect on men. Also, the mapping from degree to occupation differs across gender: A business administration degree is associated with manager positions for men, while women concentrate in occupations that tend to pay less, such as general office workers. This finding is in line with a a recent study by Andersen et al. (2020), who using discontinuities that randomize applicants near admission cut-offs, find that fields with larger gender gaps causally reduce female earnings in Denmark.

This paper is related to three strands of the literature. First, it contributes to the small literature looking at how changing prices can affect educational choices (Evans, 2017; Stange, 2015; Denning and Turley, 2017; Denning et al., 2019). Even though tuition fees are non-existing for public education at all levels in Norway, the large land surface of Norway combined with scattered settlement involves substantial living and travel cost for individuals who have to leave their home-region to take higher education (which was the case pre 1950). Our findings indicate that changing the costs by reducing distance to college can alter educational choices. However, distance to college does not seem to be important in influencing the decision on whether or not to enroll in higher education.

Second, our paper is also a contribution to the growing literature that seeks to understand differences in field of study choices by groups. This reform did not induce women to become engineers nor men to become nurses. We find that the establishment of engineering colleges increased up-take only among men whose parents had higher education, while nursing appears to have affected all women.

Finally, we also add to the literature looking at the location of higher education institutions. Andersson et al. (2009) and Carneiro et al. (2022) study how geographical expansion of higher education institutions in Scandinavia affect local productivity. Our study shows that regional investment in colleges may not increase the local education level, but may have an effect on individual's propensity to remain in the area.

The location of the colleges was a result of a complex political process aiming to improve access to education across the country. As the locations were determined by the central government, debates in the parliament reveal that local suitability or demand were not prerequisites when politicians decided where to place the new colleges (Johnsen, 1999; Ottoesen, 1969; Knutsen, 2017)[2]

To ensure that our results are not driven confounding factors, we control for municipality level time trends in our main specification and show that our results are robust to a range of different geographical definitions of college access. In addition, we follow a previous study by Bhuller et al. (2017) in showing that municipality level characteristics in 1950 and 1960 cannot predict reform year.

Although one intention of establishing colleges across the country was to increase access to higher education, our results indicate that the people growing up in the affected areas did not change their study length. Such policies may have a larger effect in settings where moving is more costly and education scholarships less generous than in the case of Norway in this period.[3] However, it seems to be room for policy-makers who want to change the composition of the local workforce, to do so through locally offering specific field of study.

The rest of the paper is structured as follows. Section 2 describes the institutional setting and the history of the educational expansion, Section 3 explains our data and empirical strategy and Section 4 presents our main findings on degree take-up and labor market outcomes.

[2]Most regions already had a teaching college, therefore we do not consider teaching as a part of the reform. Teaching education was also not much debated in the parliament in debates regarding new college locations.

[3]Several studies from other countries find a positive relationship between prioximity to newly established colleges and education length and use college distance to instrument educational attainment (Currie and Moretti, 2003; Suhonen and Karhunen, 2019).

2 Institutional settings and background

2.1 Expansion of higher education in Norway

In the middle of the twentieth century, the level of higher education in Norway was low. In the 1950 Census, only 35 percent of men and 26 percent of women had any education above elementary school.[4] In addition to socioeconomic background, geographical background was important in predicting education in Norway, but the educational advantage by geographical centrality was reduced considerably in the period after 1960 (Lindbekk, 1998). After World War 2 (WW2), policies aimed at equalizing social and economic differences across the country gained wide popular support. One such intervention was the establishment of higher education institutions across the country. Throughout the 1950s and 1960s Norway experienced a rapid increase in university colleges whose location was decided politically, and was not necessarily referring to local demands or resources.

Regional university college boards (so called "regionale høgskolestyrer") were established in order to regionally integrate the most common majors at the university college level - such as teaching, engineering, nursing, business administration. In Norway, university colleges, for simplicity also denoted "colleges" in this paper, is the designation of a higher educational institution that traditionally has offered short, career-oriented types of education at or below what today is know as bachelor level. Traditionally, a characteristic difference between colleges and universities has also been that colleges do not conduct academic research. Before WW2 only the three largest cities (Oslo, Bergen and Trondheim) and central municipalites located close to the biggest cities offered higher education at the university- and university college level.[5] Politically it was therefore emphasized that the new colleges should be spread around the country. Thus, the location of colleges was to a large extent regarded as a regional policy measure (Norwegian Ministry of Education, 1975).

The exact location of the regional colleges was decided by the parliament, where it was subject to extensive debate. Even though recruitment of students and professional environment were supposed to

[4] This information is obtained here: https://www.ssb.no/befolkning/artikler-og-publikasjoner/si-meg-har-du-studert

[5] In the period before 1940, there were ten ten nursing colleges and 7 engineering schools in Norway. In addition to being located in Oslo, Bergen and Trondheim, they were located in Stavanger, Skien/Porsgrunn, Tønsberg, Bodø, Lillestrøm, Grimstad and Follo.

be considered, debates in the parliament suggest that equalization of educational opportunities across regions was more important. The minister of education, Kjell Bondevik, later regretted that regional political considerations had been decisive for the location of colleges (Johnsen, 1999). An example of how such regional political considerations looked like in practice is the placement of a college in Bø, a village with 4000 inhabitants in Telemark, which was chosen at the expense of other cities in the same county such as Skien and Porsgrunn with much larger populations.

As a result of this targeted district policy measure, the number of colleges outside of the biggest cities increased substantially in the period after WW2. 9 new nursing colleges, 11 new engineering colleges and 13 business colleges were spread around the country. Figure 3 shows maps of the location of new colleges (established after 1940) offering a degree in nursing, engineering and business administration. The red dot indicates the exact location of the new colleges. The dark blue areas are commuting zones which will be discussed in section 3. Note also that a degree in business administration at the college level did not exist in Norway before 1969. However, approaching 1990 each county/district had at least one college offering this degree. Table A.1 in the appendix gives an exact overview over year of establishment and localization. A thorough description of the establishment of these education institutions is given in Knutsen (2017). Figure A.1 in the Appendix shows a map of colleges established before 1940.

While the number of students enrolled in the universities remained the same in the time period between 1970 and 1985 , there was a substantial increase in the number of students enrolled in college (Johnsen, 1999). The increase is illustrated in Figure 2. Only 2.2 percent of the 1940 birth cohort had a college degree. This share increased to almost 7 percent for the 1950 cohort and 12 percent for the 1970 birth cohort. The first increase in this share, from the late 1940 birth cohort to the 1950 birth cohort, coincides with the increase in nursing and engineering degrees for the same birth cohorts (cf. Figure 1). The second increase coincides with the establishments of the business administration degree from 1969 and onwards.

Tertiary education in Norway relies mainly upon public funding. Public universities and colleges have very low or no tuition fees and funding to cover living expenses is available to everyone, especially after loans and scholarships ceased to be means-tested on parental income in 1968. In 1969,

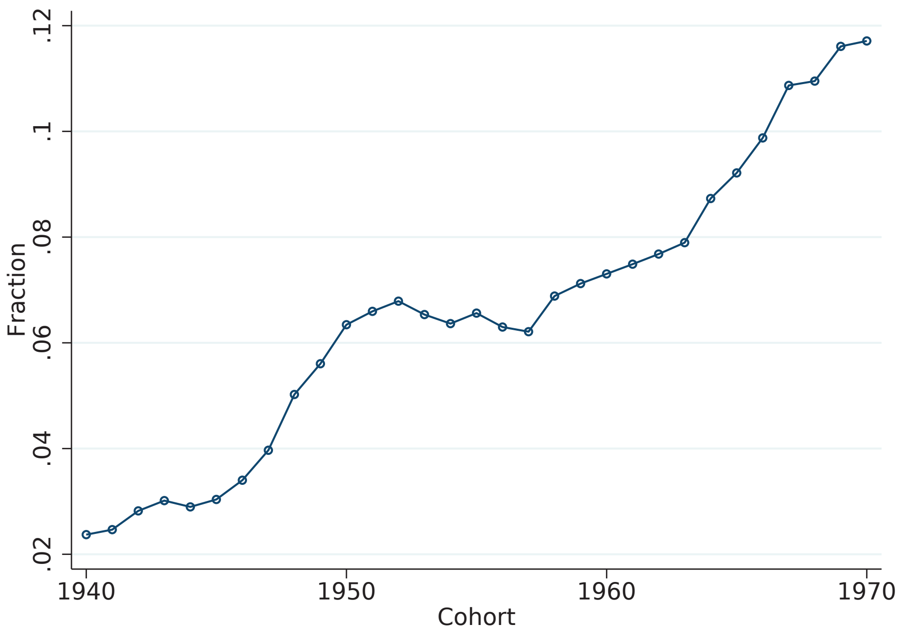

Figure 2. The fraction with a college degree in Norway, by birth cohort

97 percent of those that applied, were granted a scholarship and loan (Ministry of Education, 1969).[6] Thus, the private cost of education was mainly the opportunity cost of not working and potentially moving costs.

2.2 The Norwegian education system

Education in Norway consists of mandatory elementary school and 3 years of high school consisting of academic or vocational tracks. Today a high school diploma from the academic track is required to enter higher education. However, in the time period subject to this study, three years in high school was relatively uncommon, and a high school diploma was not a formal requirement for entering higher education in nursing and engineering. As high school became increasingly common, some types of educations changed the entry criteria. In 1977 for engineering and in 1981 for nursing, the entry criteria changed from a compulsory schooling diploma to a high school diploma. However,

[6] Lånekassen, the public institution that gives loans and scholarships to students was created in 1948, but was strictly means-tested in the early years. Towards the end of the fifties the number of students receiving scholarships increased as the means-testing criteria were relaxed and new education institutions were eligible. In the early 1960's the majority of students were receiving loan and scholarship (Røseth, 2003).

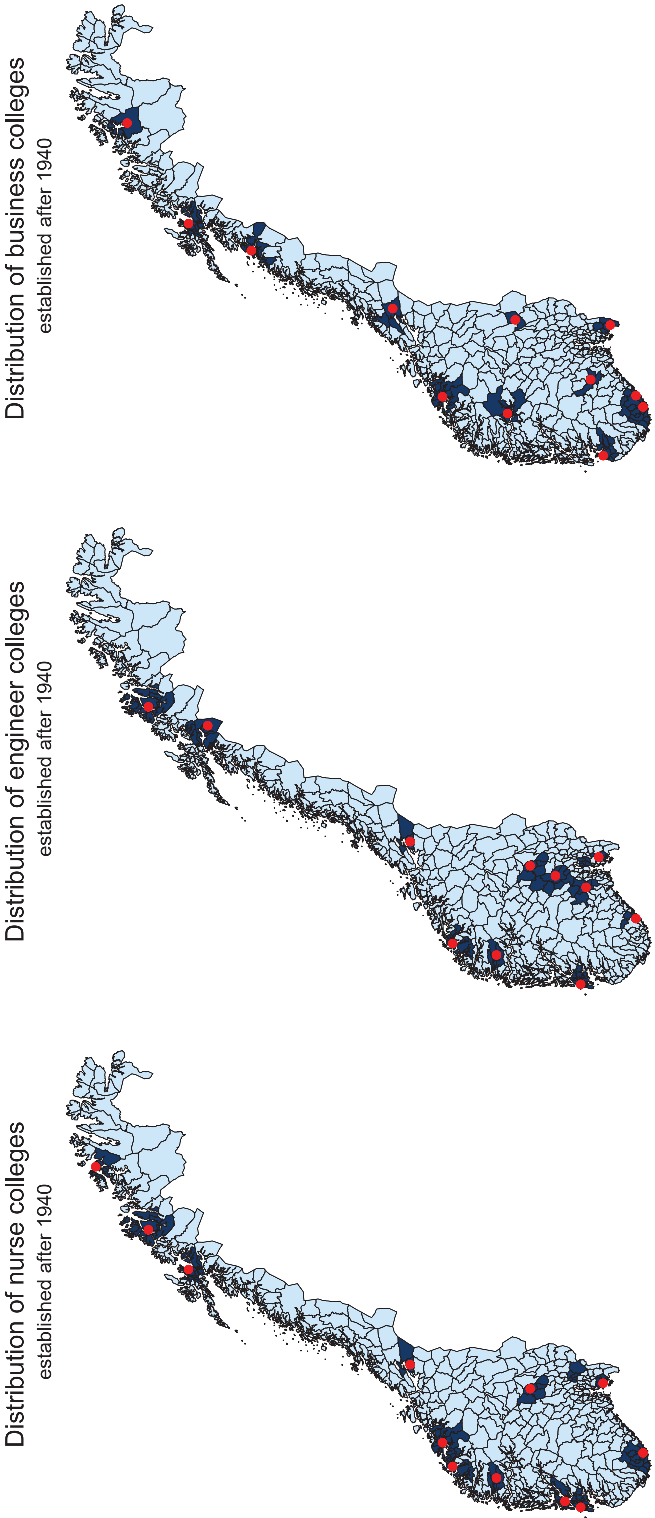

Figure 3. Location of nursing, engineering and regional colleges established after 1940, with commuting zones highlighted

most students entering nursing schools had finished high school also before the entry criteria changed (Aamodt, 1982). In our sample the median age when enrolling in nursing or engineering schools is respectively 22 and 20. For the business administration track, that was first introduced in 1969, a high school diploma has always been required in order to enrol.

Moreover, there was a tendency in Norway in the decades after WW2 that youth entered the work force after elementary school (7-9 years) and came back into the education system later, in particular to colleges and vocational schools (Bostad, 2007). OECD identifies this as a feature that makes the Norwegian tertiary system distinctive even today: students are somewhat older when they commence and graduate than in many other countries (Clark and Sohlman, 2009). In the period we study, there was therefore no clear starting age for higher education. This is the reason why we look at a wide range of age thresholds in our regression specification.

2.3 Business administration degrees

A new 2-year college degree in business administration (''økonomi og administrasjon'') was introduced with the establishment of the regional colleges and its curriculum was determined by the central government: About sixty percent of the curriculum was to consist of methodological subjects such as mathematics, statistics, economics and an introduction to information technology, and the remaining forty percent in applied subjects such as accounting and human resources.[7] The idea was that the degree-holders could enter management and leader positions in local firms. In this way, the new colleges were to support economic growth in the regions in which they were located. The business administration degree turned out to be by far the most popular degree at the new colleges, in terms of number of students (Johnsen, 1999).

2.4 Nursing and engineering schools

Nurses in Norway organized early; in 1912 the Norwegian Nurses Association was established. From the beginning they campaigned for a better and standardized training of nurses. As a result the education of nurses was a 3-year degree starting some places already in the 1920's and national standardized exams were introduced in the 1950's together with a national authorization of nurses

[7] All regional colleges, with an exception of the one in Lillehammer, Oppland, offered this degree from the start.

(Norwegian Nurse Association, 2017). Nursing was at the level of upper secondary school until it was upgraded to university college in 1981. Hence, until 1981, a compulsory schooling diploma was required in order to enrol in a nursing school, while later a diploma from upper secondary school (videregående) was needed. Individuals in our sample are born between 1940 and 1973, hence for a vast majority of our sample, compulsory schooling was sufficient. After the compulsory schooling reform the graduation age from compulsory schooling was 16. However, in practice most students enrolled in the nursing schools had completed high school even though this was not a formal requirement (Aamodt, 1982). Nursing was and continues to be a profession dominated by women. In the data, as much as 89 % of nursing degree holders are women. Average age for initiating a nursing degree for the 1940-1973 cohorts is around 23 years.[8]

Before 1977, the engineering schools were equivalent to technical schools, and a high school diploma was not necessary in order to enter. In 1962, the education in technical schools increased from 1 to 3 years, or 2 years for students who had completed the science track in high school. In 1977, the schools changed name from technical schools to engineering colleges and the criteria for entry changed; now a high school diploma was a prerequisite. With a male percentage of 88, engineering is almost exactly as dominated by men as nursing is by women.

3 Data, sample selection and empirical strategy

The analysis in this paper is based on data from Norwegian administrative registries, which can be linked together using a personal identification number. This unique number was established as part of the National Population Register in 1964, and includes all individuals alive and resident in Norway at some point since that year. The census of 1960, as well as later censuses, can also be linked using the personal identification number.

3.1 Population data, municipal structure and commuting zones

Our sample consists of all individuals born between 1940 and 1973 residing in Norway at the age of 16 and not subsequently emigrating. The data set is constructed from the population register, tax

[8] The individual-level micro data used in the analysis only has information on the year of education for those who started their degree after 1974. When upgraded to a university college, the average age changes slightly, from 22 years between 1974- 1980 to 23,6 between 1981-1990. There is no significant difference in average starting age between treated and untreated regions.

records and the educational register, based on the individual ID numbers. The population registry has recorded the municipality of residence for all individuals each year. We base our analysis on the municipality borders of 1980.[9] We use the municipality in which an individual lives at the age of 16 as recorded in the population registry. For those born before 1949, who are 15 years or older when the population register was established in 1964, we use their municipality of residence recorded in the 1960 census which is the only earlier source available. The reason for using the registered municipality at 16 is to reduce the systematic difference in how we measure municipality of residence, between older and younger cohorts. However, we also do robustness tests using the municipality of residence when individuals are 10 years old, excluding older cohorts for whom we do not have this information. Individuals with unknown municipality (around 3 per cent, mostly immigrants) and individuals registered as living outside Norway are not included in the analysis.

In the 1970 Census, all individuals in Norway reported their municipality of residence as well as municipality of work. Based on this information, we construct commuting zones around the colleges that we base our identification on. A $municipality_i$ is defined as being within commuting distance of an education institution in municipality j if at least 1% of the employed individuals in $municipality_i$ work in $municipality_j$. For the few cases where a municipality falls within the commuting region of two educational programs, starting the same type of education at different times, the earliest date of start up is used. The maps of commuting zones of nursing and engineering colleges are shown in Figure 3 where the commuting zones around a college are represented by a shaded area.

3.2 Education

The primary data used to assess the effect of the college establishments on degree attainment are collected from the Norwegian Education Data Base (NUDB). This data base contains individual-level data on all education completed by 1970 and education undertaken from 1970 and onward (Vangen, 2007). Education is coded at a high level of detail (six digits in the NUS classification) so that we are able to identify particular educational fields. Fields that have changed levels, such as nursing,

[9]A large number of Norwegian municipalities were merged with their neighbouring municipalities in the 1960s, bringing total number down from 747 to 450. We base our analysis on the 1980 municipality borders and use a conversion based on information in the 1960-census in order to find define the 1980-municipality of an individual in 1960 (i.e. all individuals are observed each year, but the granularity of information for each individual does not change over time. None of the policies discussed in this table were set at the municipal level.

are coded with the level they have today. In general, the educational institution is not reported for individuals completing their education before 1999.

NUDB has every education initiated linked to a national identity number. This implies that we can track every individual's educational career. In order to obtain a data set with one observation per individual, the data is collapsed on the highest education level initiated and the respective field/degree we are looking at. This means that we also include individuals with for example a nursing degree who for example continued on a master program at a later stage.

A challenge with the education database is that it lacks education data for around 120 000 individuals, most of these born before 1958. Around half of them are registered in NUDB, but without information on education. This is resolved using the censuses from 1960 and 1970 which also contain information on level and type of education. For these individuals highest achieved education in 1970 is used and the fields subject to this study are re-coded.[10]

3.3 Income

Information of individual-level income is obtained from the tax authorities. These registers are available from 1967 and onward. As a measure of income we use total pre-tax income from work. This includes mostly income from work, but also benefits that enter as a substitute to income from work such as paid sick leave and temporary disability benefits. In order to obtain a meaningful estimate of earning differences we keep work-experience roughly fixed, and measure the mean income for each individual in our sample when they are between ages 35-40. At this age, most individuals have finished their education and have entered the labour market.

We observe many individuals (about 22 000) with missing or zero income at the age of 35 to 40 years, two thirds of these are women. Before 1977 the value missing is more frequent and was used for those that had zero income. Therefore, we insert zero income for all with missing income who are living in Norway at the age of 35-40. Thus individuals not registered as residents in Norway at this age (due to emigration or death) are dropped from the income regressions. Because of missing

[10]Complete educational histories are not available in the census data, only highest achieved education. Moreover, partially finished educations may not be visible in the census data. This could in principle introduce a comparison problem. However, based on the data that we do have, there is little reason to believe that a substantial number of individuals completed a nursing degree and then subsequently a higher, unrelated degree before 1970. An earlier version of this paper conducted analyses using only the education database data and found qualitatively similar results to those we report here.

or zero income, we choose not to use log of income as our main outcome variable. Instead we create two variables capturing whether the individual earns above certain thresholds (low and high). These thresholds are based on the pension base rate ["Grunnpensjon"]. This is a rate adjusted annually and forms the basis for calculating the Norwegian state pensions. The thresholds we look at are incomes 2 and 5 times the pension base rate. In 2017 this corresponds to 187 000 and 468 000 nok.[11] Both income and the pension base rate are adjusted for wage inflation. Earning of at least 187 000 nok (two times the pension base rate) is an indication on whether the individual participates in the formal labour market, whereas earning at least 467 500 nok (five times the pension base rate) is an indication of full-time employment. We make dummies to capture whether an individual earned above or below the pension base rate in the following way: An individual whose earnings in the age-interval 35 to 40, is above the pension base rate for those years, get the value 1 and 0 otherwise. We also report the results using log of income as the outcome variable, excluding those with zero income.

3.4 Sample selection

As mentioned above, our sample consists of individuals born between 1940 and 1973. When studying nursing and engineering we exclude areas that offered degrees in nursing and engineering prior to 1940 since they were not a part of the college establishment taking place in the 1950s- and 60s which had a strong component of regional policy measure. This involves that we mainly drop individuals living in the biggest cities and some nearby areas.[12] Our total sample then consists of 688 939 individuals. The college degree in business was first offered in 1969 as part of the roll-out of university colleges across the country. Hence, we do not drop any observations when looking at business. The data used in this paper are summarized in Table 1.

In the upper left panel we see that nursing is very female oriented and engineering is very male oriented. Having a degree in business on the other hand (lower left panel) is equally distributed across gender. When stratifying on parental education (middle and right panel), we see that all three degrees are more common among people whose parents have higher education than among individuals of

[11]This corresponds to approximately 23 370 USD and 58 500 USD with a currency exchange rate of 8 nok to 1 dollar

[12]The following cities are dropped both for nursing and engineering: Bergen, Bodø, Follo, Grimstad, Lillestrøm, Oslo, Skien, Stavanger, Trondheim, Tønsberg - total 700 835 observations. This implies that we also drop the neighbouring municipalities, belonging to the pre-defined commuting zones. A map of the areas that are dropped is shown in Appendix A.1.

Table 1. Summary statistics: Fraction with a degree in nursing, engineering and business administration and labour market outcomes

	All		Mother low ed		Mother high ed	
	Women	Men	Women	Men	Women	Men
Nursing degree	0.071	0.008	0.051	0.005	0.096	0.012
Engineering degree	0.009	0.062	0.006	0.045	0.014	0.083
Mean earnings age 35-40						
- >2G (187 000 nok)	0.75	0.92	0.71	0.92	0.82	0.94
- >5G (467 500 nok)	0.24	0.69	0.17	0.65	0.32	0.74
- average	168 185	285 223	144 963	257 126	197 914	322 813
Nr of obs	325275	363664	182422	207922	142853	155742
Business degree	0.020	0.022	0.012	0.013	0.028	0.031
Mean earnings age 35-40						
- >2G (187 000 nok)	0.76	0.92	0.71	0.91	0.81	0.93
- >5G (468 000 nok)	0.26	0.71	0.18	0.66	0.34	0.75
- average	173191	296382	146694	262108	202122	335278
N	842975	932393	439367	495027	403608	437366

Note: The individuals in this table are born between 1940 and 1973. For nursing and engineering we only include individuals whose municipality of residence at the age of 16 was not within commuting distance of a college before 1940. For business we include all individuals as the business degree was first introduced in 1969. When looking at average income, we drop observations with missing information (zeros are included). This leaves us with a sample of 362 027 men and 324 002 women in the upper panel and 932 393 men and 842 975 women in the lower panel.

lower educated parents.[13] Turning to labour market outcomes, more men than women earn above both the high and low income threshold. The fraction earning above these tresholds is also higher for men and women with a higher educated mother compared to men and women with a lower educated mother. .

[13]An individual is defined to have lower/higher educated parents if the mother has 10 years of schooling or less/more than 10 years of schooling.

4 Empirical strategy

Does reduced distance to college affect the decision to take higher education and which field to specialize in? Distance to college may not be exogenous to unobserved factors which also affect future labor market outcomes and the relationship between college proximity and educational attainment may be a spurious one. We propose to solve this problem by exploiting the time-variation in the roll-out of colleges across the country, and estimate the reduced form effects of college establishment on the probability of taking the specific degree being offered. This reduced form equation can be modelled in the following way:

$$Outcome_i = \gamma reform_{i,age} + \omega_t + \phi_m + t \times \phi_m + v_i \tag{1}$$

$Outcome_i$ is an indicator variable equal to 1 if an individual has the given outcome - for example a degree in nursing. The probability of obtaining such a degree depends on whether the individual resides in a municipality offering a college degree in nursing or engineering, or is residing in a municipality with commuting distance to another municipality that offered the same degrees. In the remainder of the paper, we will denote this as "access to college". Figure 3 gives an overview over municipalities with commuting distance - affected zones - (in dark blue) to a municipality with a higher education institution offering the degrees of interest (in red).

As explained in Section 2.2, in the period we study, there was substantial variation in the age at which individuals started higher education. Instead of defining a particular age as the first treatment age, we estimate several specifications where we define $reform_{i,age}$ to be a dummy variable taking the value one if the individual gets "access to college" at different ages, ranging from 15 to 26. I.e. $reform_{i,18}$ takes value one if an individual was 18 years or younger when he or she got access to college, $reform_{i,19}$ takes value one if an individual was 19 years or younger when he or she got access to college etc. Thus, the treated population consists of individuals who are younger than a specific age (ranging from 15 to 26) at the time the colleges is established.[14] We compare the treated

[14]We only have information on starting time at the individual level for those undertaking degrees after 1974. The only information we have on this for cohorts entering higher education before 1974 are Statistics Norway reports on the age composition of students in nursing schools. The median age according to these records is similar to what we find in our

population to a comparison population consisting of individuals who were older than the particular age when the college was established. ω_t designates a full set of cohort dummies, while ϕ_m refers to municipality fixed effects. The reduced form effect is then derived by comparing the difference in the outcome variable between the treated and untreated population in the affected zones to the difference in the same outcome variable between the treated and untreated population in the unaffected zones. A positive difference implies that college establishments increase the take-up rate for degrees in nursing, engineering and business administration. We also include municipality specific time trends represented by $t \times \phi_m$. This relaxes the assumption, crucial to a difference-in-difference estimation, that treated and untreated municipalities experience parallel trends before the college establishments. The common trend assumption may also be violated if the roll-out of the reform across municipalities and regions is systematically correlated with characteristics that also affect our outcome variables. We will come back to this issue in the next subsection. v_i is a random error term, and is clustered at the municipality level. Equation (1) is estimated separately for men and women.

Parents may behave strategically in the sense that they choose to move to a municipality with access to college. In order to shed some light on this potential problem, we perform robustness checks where we measure residential municipality at the age of 10 instead of 16. Moreover, we also show that the results are not sensitive to how we create the affected commuting zones and that the results do not change substantially if we drop individuals residing in zones that where never affected. These results are reported in Section 5.5.

Almost at the same time as the college reform took place, compulsory schooling also expanded in Norway, from 7 to 9 years. This reform began in 1960 and was completed in 1975. The compulsory schooling reform also involved a standardization of the curriculum which possibly enabled more students to become eligible for enrolling in nursing and engineering schools. As a large part of our sample, namely those born between 1946 and 1961, were subject to the roll-out of this reform, we also control for being exposed to this reform although this does not alter our results.

In a next step, we also estimate the reduced form effects of being affected by the college reform on labour market outcomes when individuals are 35-40 years old. As outcome variables we use dummy

data for those entering nursing after 1974. However, these records do not report the starting age.

variables taking the value one if the individual earns above a certain threshold, corresponding to part-time work and full-time work, as outlined above. Additionally, we also estimate specifications where we look at the intensive margin, i.e. our outcome variable is log of income involving that individuals with zero and missing values are not considered. The regression model is similar to 1, with the dummy variables for low and high income on the left hand side instead of educational attainment. As the estimated coefficients are reduced form coefficients, this model estimates the effect on labour market participation of being exposed to the college reform. The main reason for not presenting 2SLS estimate is a questionable exclusion restriction; college establishments may have affected wages other than through college degrees. Geographical expansion of higher education has been shown to affect productivity, skilled wages and innovation (Carneiro et al., 2022; Andersson et al., 2009).

4.1 Timing of college establishments and municipality characteristics

If our reduced form estimates are to be interpreted in a causal fashion, we must assume that the timing of college establishments are unrelated to underlying trends at the municipality level. Unfortunately panel data at the municipality level is scarce. However, both the 1950 and 1960 census provide aggregate level municipality characteristics. The characteristics we look at are education level, voting behavior and sectoral compostion of the local labor market. We follow the empirical strategy in (Bhuller et al., 2017) to investigate whether the timing of reform implementation is correlated with municipality characteristics as recorded in the censuses. In other words we want to check whether, given municipality characteristics in 1950 or 1960, we could have predicted in which areas the new colleges were to be established.

We run the following regression

$$T_{mt} = (T_t \times B_{m,1950})'\gamma_t + \varepsilon_m \tag{2}$$

where T_{mt} is equal to 1 if municipality m implemented reform in year t and $B_{m,1950}$ is a vector of municipality level information from the 1950 and 1960 census. In this way γ_t captures whether there is a correlation between the year of college establishment and municipality characteristics, in addition to the difference between municipalities near the new colleges and municipalities that did not get college access throughout the period. The chosen municipality characteristics are based on Bhuller et al. (2017), but for the 1950 census only some of these characteristics are available.

Plots of the coefficients are shown in Appendix B. The coefficient plots show the estimates for γ_t from the regressions with each municipality level characteristic. If γ_t is zero there is no correlation between year of college establishment and municipality census characteristic. There seem to be no systematic negative or positive trend for these coefficients. It would be problematic if for example there seemed to be the case that municipalities where colleges were established early, had a higher education level than those were they were established later or similarly, a different share of the population working in industry or the service sector. In addition, we check whether the voting share for the labor party or the conservative bloc can predict college location as it is possible that the party in power could favor its own voters when deciding where to locate a college. In general there seems to be no systematic differences in municipality characteristics between municipalities that were the first to get college access after 1940 and those that got access later.

5 Results

We now turn to the estimation results obtained by using the empirical approach presented in the previous section. We proceed in four steps. First, we assess the overall effect of the education reform on the choice of field for the affected individuals. Second, we examine whether there is heterogeneity with respect to parental socioeconomic background. Third, we address labor market outcomes, measured as income between between the ages of 35 and 40 years. Finally, we also examine the broader picture of counterfactual outcomes, i.e. what the education choices of the affected individuals had been had the reform not been implemented.

5.1 College proximity and choice of field

The reduced form estimates from estimating Equation (1) for men and women separately are presented in Figure 4. Panel (a) shows the coefficients when the outcome is whether an individual obtains a nursing education; engineering education (panel b) and business administration (panel c). In each panel, coefficients from 24 separate regressions (12 for women and 12 for men) are plotted.

In each regression, for a given age *i*, we define "access to college" as being *i* years or younger at the time a college was established within commuting distance (denoted "affected zone" below). In each regression, municipality and cohort fixed effects are included, in addition to an interaction between these two (municipality-specific time trends) and dummy variables controlling for the compulsory schooling reform.

For nursing, the effect is positive and statistically significant for women who were between 19 and 22 (and younger) when the college was established. For instance, women in affected zones who were 19 year or younger when the college was established had a higher probability of taking a nursing degree (compared to older women in the same zones) than women the same age living in unaffected zones. A point estimate of 0.013 indicates that the share with nursing education increases by 1.3 percentage points (or equivalently, the probability of obtaining a nursing degree increases by 1.3 percentage points) compared to unaffected zones. For men the effect of getting access to a nursing college is close to zero. Looking at engineering, there is a positive and significant effect for men who were between 15 and 21 (and younger) when the engineering college was established. The point estimate is about 0.01 indicating a 1 percent increase in the share with a engineering degree (from a sample mean of 6 percent). The effect for women is zero. Turning to business, men and women living in zones where a business college is established have similar probabilities of taking a degree in business. The point estimates for men are slightly higher than for women, especially for the youngest ages, but they are not significantly different.[15]

Summarized, we find substantial effects of the establishment of colleges on choice of field, in the sense that individuals in affected areas have an increased propensity to take up the educations on offer. However, there are some important exceptions: men do not respond to extension of nursing degrees and women do not respond to engineering degrees. In this way, existing gender patterns are preserved.[16]

[15] All these results are also presented in column (1) in Appendix Table C.1 and C.2.

[16] Although we only report the point estimate for access to nursing (engineering/business) college in the specifications where nursing (engineering/business) degree is the outcome variable, we also control for access to engineering/business (nursing) colleges. This does not affect the the main results.

Figure 4. The reduced form coefficients of the college reform on the probability of taking a degree in nursing, engineering and business administration

(a) Nursing

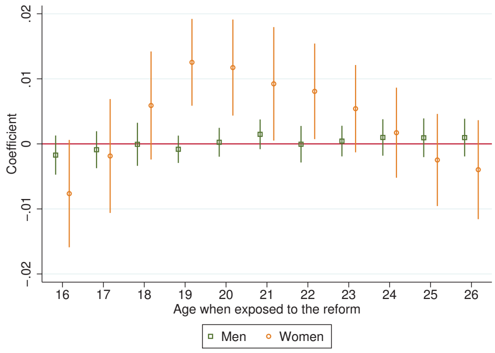

(b) Engineering

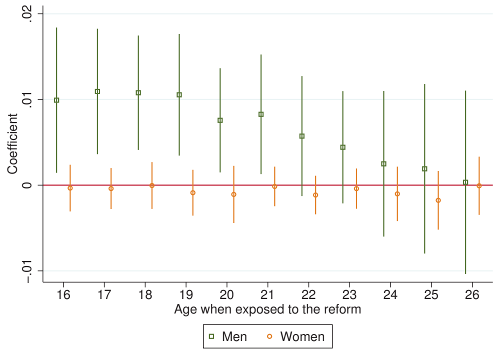

(c) Business Adminstration

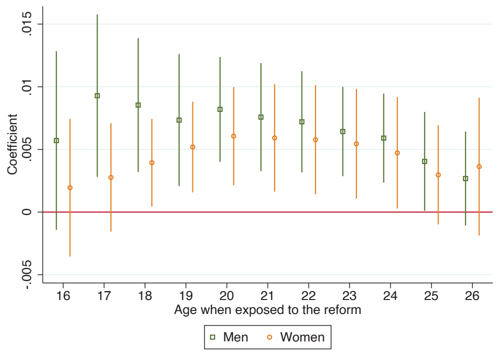

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term. Standard errors are clustered at the municpality level.

5.2 Heterogeneity with respect to parental background

In Figure 5 we report results from estimating Equation (1) when stratifying on mother's education (which is a variable reported at the individual level in the educational database). The results for nursing are presented in the upper panel, the results for engineering are presented in the middle panel and the results for business are presented in the lower panel. For nursing and business administraion the point estimates are of the same magnitude across socio-economic background, although a bit more precise for those with a lower educated mother . For engineering, the results are clearly higher and the estimates more precise for men with a higher educated mother. All the results presented in both Figure 4 and 5 are also reported in columns (2) and (3) in Appendix Figure C. In the sample period we focus on, mothers had lower education than the fathers. The results are unaltered if we stratify on fathers' length of schooling.

5.3 Labor market outcomes

We now turn to a further examination of how the college openings affected the labour market outcomes of the affected individuals when they are between 35 and 40 years old. As our income data is obtained from the public tax and pension system, we do not have direct information on working hours and hourly wages. Rather, we interpret the income data on two different margins. First, by examining the share of individuals above a certain income thresholds, we can assess the effect of the college reform on participation in the formal labor force. In particular we focus on two thresholds, minimum income (187 000 nok)[17] as an indication to what extent the individual participates in the formal labour force, and equivalent full-time salary (467 500 nok)[18] as an indication to what extent the individual works full time. Second, by using log income as outcome in a regression contingent on being part of the labor force, we also investigate the intensive margin.

For both nursing and business, but also to a large extent engineering, the effects on the take up rate of the different degrees seems to be largest for individuals who are between 18 and 22 when treated for the first time (where 19 means 19 years or younger when the college is established, 20 is 20 years or younger when the college is established, etc.). In this analysis we therefore only focus in this age

[17]This figure equals two times the pension rate mentioned in subsection 3.3.
[18]This figure equals five times the pension rate mentioned in subsection 3.3.

Figure 5. The reduced form effects of the college reform on the probability of taking a degree in nursing and engineering, estimated separately for higher (left)- and lower (right) educated mothers

(a) Nursing

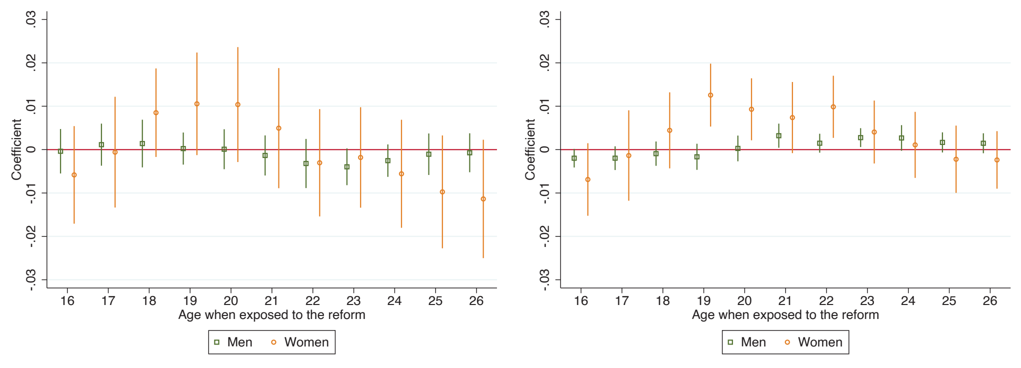

(b) Engineering

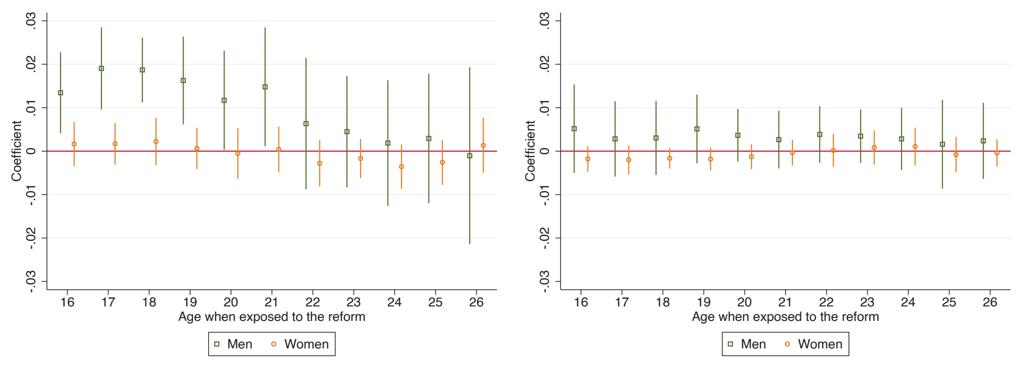

(c) Business Administration

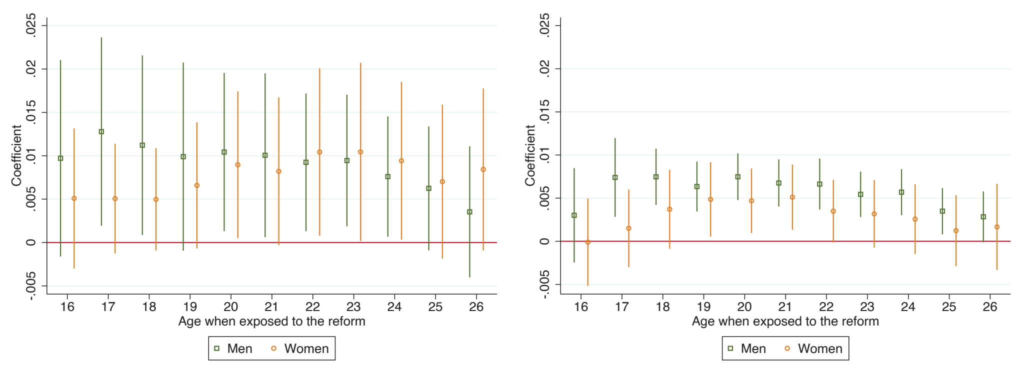

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term.Standard errors are clustered at the municpality level.

interval.

We estimate Equation (1), but change the outcome variable to the different labour market outcomes described above. The results are reported in Figure 6. In the upper panel we report the reduced form effects for women exposed to new nursing colleges and men exposed to new engineering colleges, whereas we in the lower panel report results for men and women exposed to new business colleges.

For nursing and engineering there are no effects of college openings on neither participation in the formal labour force, working full time nor earnings on the intensive margin. For business administration on the other hand, men who get access to a college degree offering this degree seem to perform better in the labour market than men living in regions without access to the same type of education. This is especially true for participation in the formal labour market, but also for earnings on the intensive margin. For women the effect on crossing income thresholds are negative and significant. The labour market behaviour of getting access to nursing and engineering schools does not vary across socio economic background. This is documented in Appendix Figure C.1 and C.2 where we stratify on mothers education. The negative effect of the openings of business administration colleges on women's labour market outcomes is to a large extent driven by those with lower educated mothers. Hence, even though men and women have the same take-up rate for a college degree in business, the labour market outcomes differ across gender. One reason for this could be that, conditional on the same degree, men end up with occupations with higher earnings in the labour market . Unfortunately, there is limited occupation data for the period we study, as high-quality annual occupation data is not available before in 2008.[19]

However, when looking at the available data on occupation we find that, contingent on having a degree in business (10,788 individuals), the most common occupations are administrative and mercantile leaders (occupations classified as management positions), ICT advisors (classified as academic professions), various occupations within the culture and sport sector (classified as college careers) and general office workers. When looking at the variation across gender we see that among those who work as administrative and mercantile leaders, 67 percent are men. As a comparison, 74 percent of

[19]Occupation data for some sectors of the economy (with highest quality for the private sector) are available annually from 2003, and there are full-count censuses every decade until 1980 (and a smaller sample for 1990).

the general office workers are women. For ICT advisors and professions in the culture and sport sector, the division across gender is roughly fifty-fifty. A complementary description of this is reported in Appendix Table C.3.

We also check whether being affected by college openings in nursing, engineering and business affect the probability of moving. In Appendix Figure C.3 and C.4 we present results from regressing Equation (1) when changing outcome variable to a dummy variable taking the value one if the individual still lives in the same municipality (panel a) or region (panel b) as the college at the age of 35.As above, when looking at nursing colleges, we only focus on women. And likewise, when looking at engineering colleges, we only focus on men. For business administration, we focus on both men and women. For nursing and engineering (Figure C.3) the point estimates are in most cases very imprecise. For business administration (Figure C.4) on the other hand, both men and women who get college access seem to have a lower probability of moving out of the region compared to men and women in the control group.

In order to better understand the results in this section, we now turn to investigate the counterfactual outcomes: What would the treated individuals have done in the absence of the establishment of colleges in nursing, engineering and business administration?

5.4 Counterfactual outcomes

In our setting, the reform is associated with changes in the costs of undertaking nursing, engineering and business administration degrees. If we were to fully understand the effect of the reform for those who were induced to undertake these degrees due to a reduction in costs, we would need information on their next-best alternatives (Kirkeboen et al., 2016). Given that there is no information on this in the historical registries, we have to infer the counterfactual outcomes indirectly by investigating how the reform affected other outcomes than the degrees subject to the geographical expansion. Looking at this, we get an indication of what the treated individuals would have done in the absence of the reform. In order to interpret the following results as counterfactual outcomes, we must assume that the reform only affected the uptake of other degrees (i.e. those not part of the reform) by lowering the

Figure 6. The reduced form coefficients of the college reform on labour market outcomes

(a) Nursing and engineering: Formal participation in the labour force (early income above 187 000 nok)

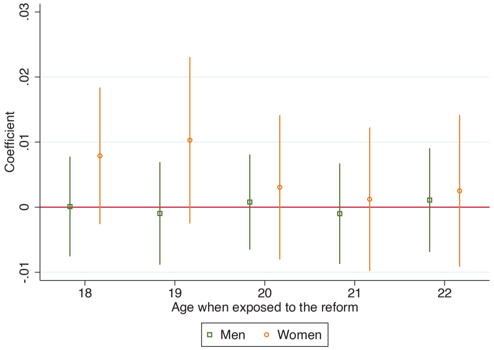

(b) Nursing and engineering: Equivalent full-time salary (yearly income above 467 500 nok)

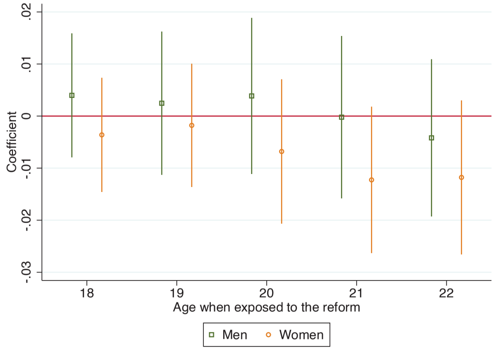

(c) Nursing and engineering: L of income

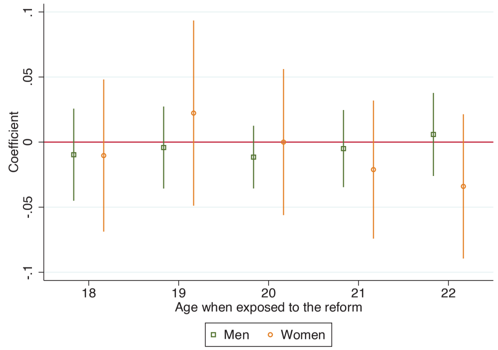

(d) Business: Formal participation in the labour force (early income above187 000 nok)

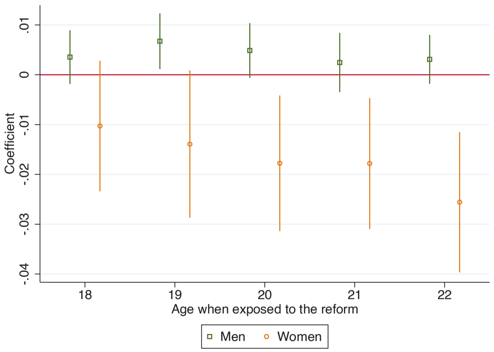

(e) Business: Equivalent full-time salary (yearly income above 467 500 nok)

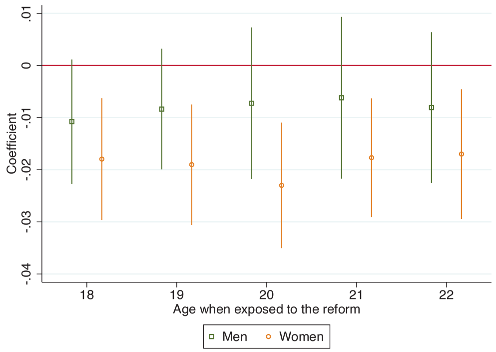

(f) Business: L of income

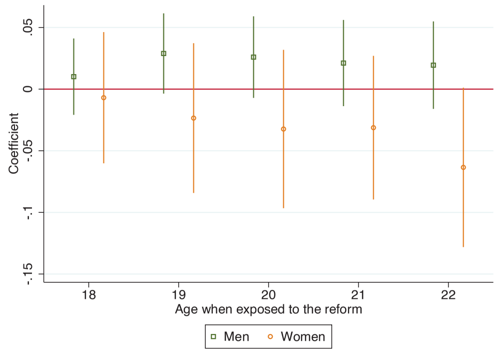

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term. Standard errors are clustered at the municpality level.

Figure 7. The reduced form effects of the college reform on taking 2 or more years of higher education

(a) Nursing and engineering **(b)** Business

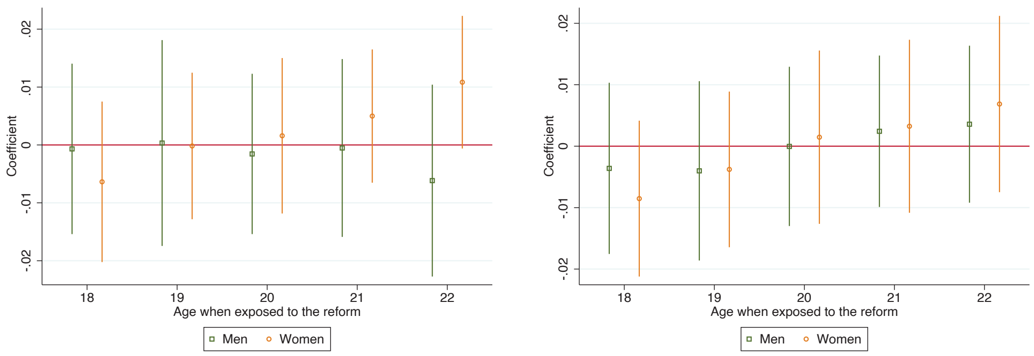

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term.

cost of the degrees that were now available within commuting distance from the treated individual's home municipality.

As a first step, we look at whether the reform affected the probability of taking higher education. Those results are reported in Figure 7. For women we only report the reduced form coefficient from the establishment of nursing colleges, whereas we for men only report the reduced form coefficient from the establishment of engineering colleges. For both men and women we see that the opening of new nursing andengineering colleges as well as business colleges did not affect the probability of taking higher education. This finding implies that the reform altered the field, rather than the level of education.

We pursue this further by looking at how the choice of other educations changed in response to the respective reforms. For each individual, we observe all education spells, and as one individual may obtain several educations, we define a list of priorities where the highest 'ranked' education on the list will be the one assigned to an individual. The exact ranking order does not matter much for the outcomes, with the exception that we rank longer educations (masters degrees) higher than educations with shorter duration; for example, very few individuals in our data obtained degrees in both technical and health related fields. [20]

[20]The variable 'field' is constructed by taking the education with the highest number in this list: 10: Any master or

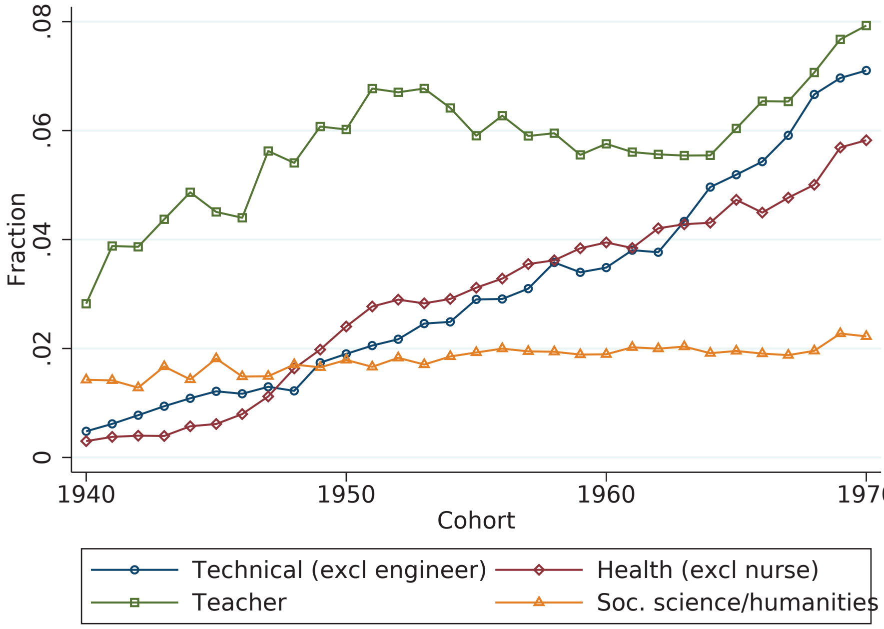

Figure 8. The fraction with a degree teaching, health (excl. nursing), technical subjects (excl. engineering), social science and humanities.

The Norwegian education data base (NUDB) uses the following categories of fields: Teaching, technology, administration, health, social science and humanities.[21] In these registries technology includes STEM subjects, administration includes business, accountancy and management, while social sciences consists of degrees such as law, political science and sociology. Other degrees in health, at the same level as nursing, are for example social worker, physiotherapy and dental hygienists. The level is determined by both the institution offering the degree and its length. Thus, the degrees at the same level as nursing, business administration and engineering will be 2-3 year degrees.

As documented in Figure 8, the proportion who took a degree in one of the educational directions we focus on increased steadily from about 1 per cent among the 1940 cohort to about 5 - 8 percent among the 1970 cohort. The development in the teaching profession differs a bit from the rest. About

Ph.D. degree; 9: Health; 8: Technical fields 7: Teaching; 6: Administration; 5: Social Science /Humanities; 4: Vocational School ("Fagskole"); 3: High school year 3; 2: High school year 1 and 2; and 1: Mandatory schooling. Thus, a person with a degree in teaching at college level, that later pursued a master degree will have a level 10. A person with a teaching and a business/admin degree will end up with a 7 as teaching is ranked higher.

[21] As the number of students in humanities is very small, this category is merged with social sciences in our analysis. There are two more categories, primary industries (such studies in farming and fishing) and transport and securtity (for example police and drivers). These are not used as outcomes because the number of students with degrees in these fields at college level is very small.

3 percent of the 1940 cohort has a degree in teaching (which is higher than for the other fields). The proportion with a teacher education increases to almost 7 percent for the 1950 cohort, then declines to 5.5 percent for the 1960 cohort. It then increases again, and is as high as 8 percent for the 1970 cohort.

The results from estimating Eq 1 when changing outcome variables to obtained degrees in other fields are reported in Figure 9 (nursing and engineering) and 10 (business administration). In all specifications we include the same control variables as in Figures 4 and 5. The only field which is negatively affected by the opening of new nursing colleges, is teaching which suggest that women living in regions with newly opened nursing colleges would have become teachers in the absence of access to a nursing college. As already illustrated in Figure 8, the fraction taking a teacher degree declined from 7 to 5.5 percent from the 1950 to the 1960 cohort. For the same cohorts, the fraction taking a nursing degree increased from 1.2 percent to 4.7 percent (see Figure 1). There is also some evidence that technological educations (excluding engineering) are suppressed by engineering colleges, although those point estimates are not as large as for teachers. For men, the point estimates for business administration is also negative, but not precisely estimated. Turning to the counterfactual outcomes for business, we see that the opening of the new business schools negatively affected the take up rate of engineering degrees among men. For women, there are some indications that women may have chosen a degree in teaching or social science and humanities in absence of a local business college. Summarized, the results in this section indicate that the college reform did not affect the level of education, but altered the choice of field, but not in such a way that it changed the gender composition in a given field.

5.5 Robustness checks

In this subsection we investigate to what extent the point estimates of college access are sensitive to the way we have chosen to define college access. Our measure of college access may be biased if commuting is correlated with factors that affect the outcome through other channels than proximity

Figure 9. The reduced form coefficients of the college reform (nursing and engineering) on the probability of taking degree in other fields

(a) Business

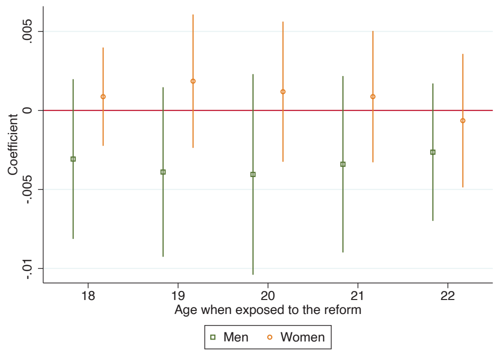

(b) Teaching

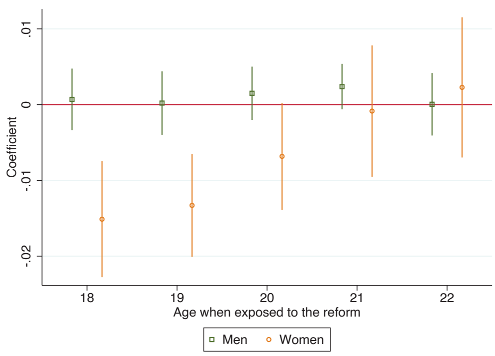

(c) Health education (excluding nurse)

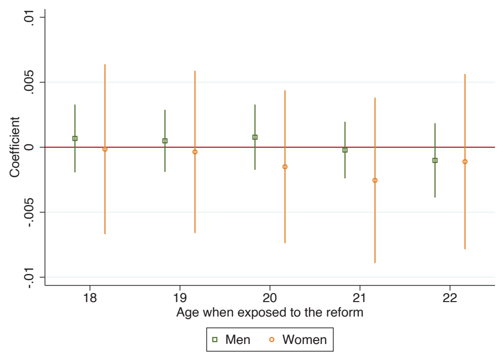

(d) Tech education (excluding engineers)

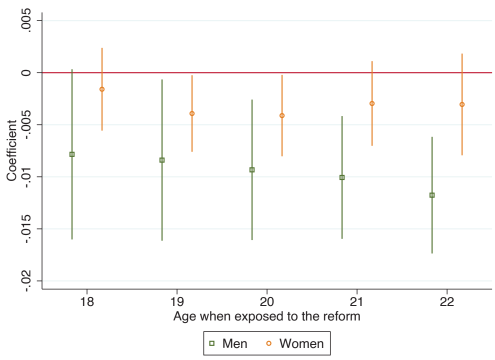

(e) Social science and humanity

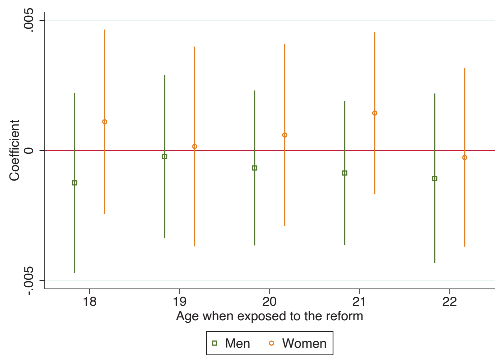

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalities with birth year) and a constant term. Standard errors are clustered at the municpality level.

Figure 10. The reduced form coefficients of the college reform (business) on the probability of taking degree in other fields

(a) Nurse

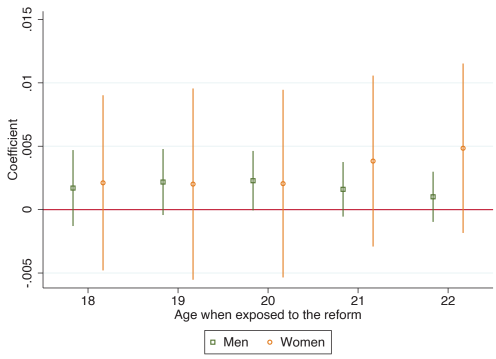

(b) Engineer

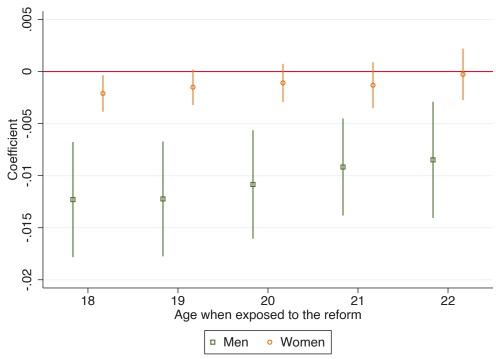

(c) Teaching

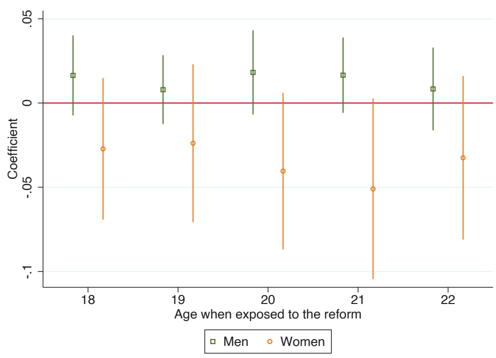

(d) Health education (excluding nurse)

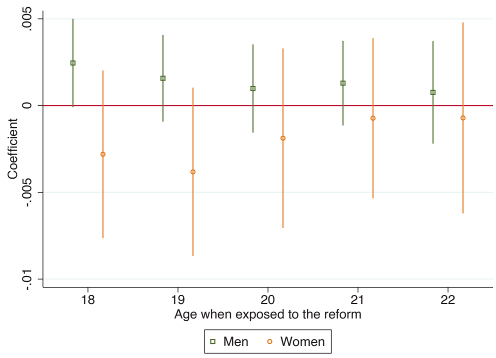

(e) Tech education (excluding engineers)

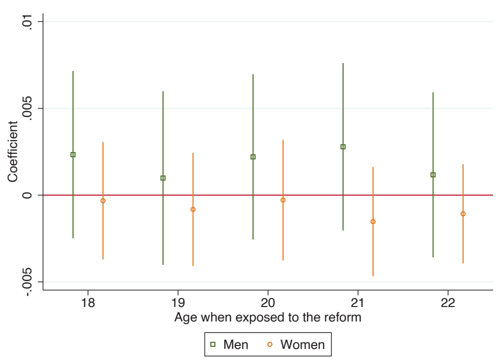

(f) Social science and humanity

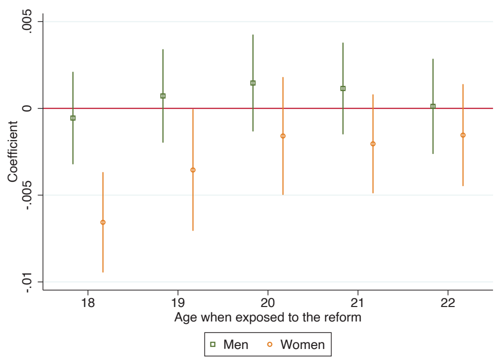

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalities with birth year) and a constant term. Standard errors are clustered at the municpality level.

to college. We therefore create alternative measures of affected zones. As a first and a second robustness measure we define individuals to get access to college if they live within a radius of 30 and 50 kilometres from the centre of the municipality where the college was established.[22] Additionaly,, we check whether our results are robust to using the Classification of Economic Regions, which are based on recent commuting data (Statistics Norway, 2000). As can be seen in Figure 11, this does not seem to affect the results.

Second, to address that parents may move to a municipality with a college before the age of 16, we estimate the baseline equation measuring residential municipality when an individual is 10 instead of 16.[23] Finally, we also check how our estimates change if we only include individuals residing in the commuting zones where the new colleges were established. The reason for doing the latter is that we are afraid that affected and unaffected zones develop differently in ways that are not captured in the municipality specific time trends. The results are presented in Figure 12, and show that the results are qualitatively similar to those in Figure 4 (the baseline results).

[22]The distance is measured using the coordinates of the administrative centre of a municipality

[23]The reason why we choose 16 instead of 10 as the baseline is that the first time we observe municipality of residence is in the 1960 census, and we can therefore only use cohorts born after 1950.

Figure 11. Robustness checks: The reduced form coefficients of the college reform on the probability of taking a degree in nursing and engineering, estimated for alternative definitions of affected zones. 30 km (left), 50 km (middle) and labour market regions (right)

(a) Nurses

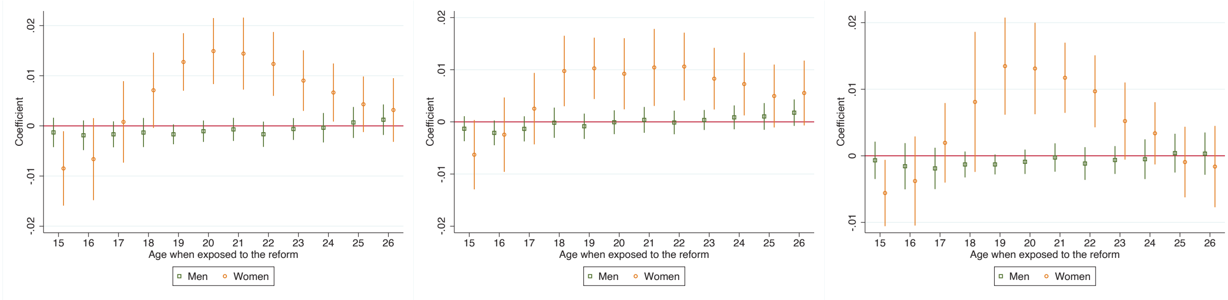

(b) Engineer

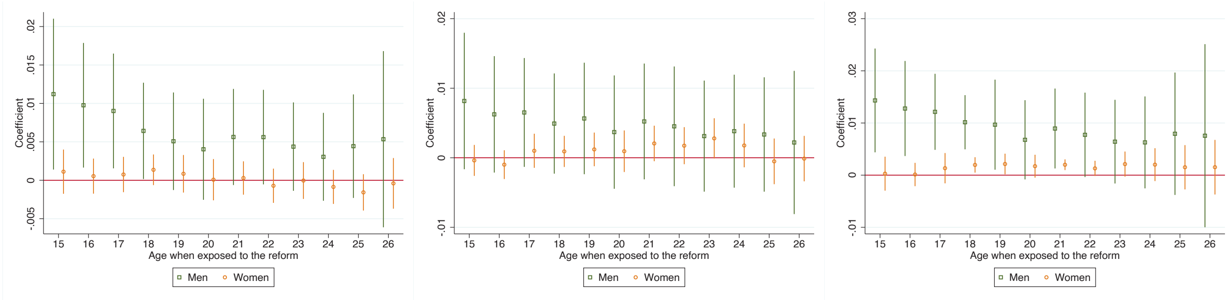

(c) Business administration

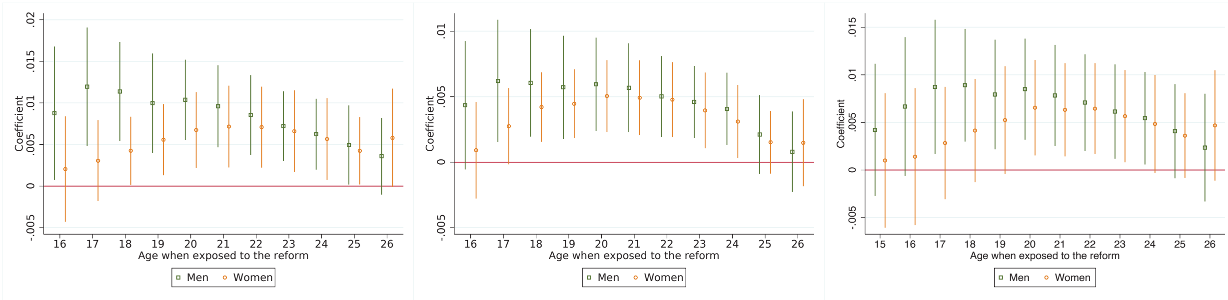

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in specifications in panel a and b are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalities with birth year) and a constant term. Included in all specifications in panel c are region fixed effects, dummy variables for birth year, compulsory schooling reform, and region specific time trends (where we interact regions with birth year) and a constant term. Standard errors are clustered at the municpality level.

Figure 12. Robustness checks: The reduced form coefficients of the college reform on the probability of taking a degree in nursing and engineering. Municipality at age 10 (left) and only affected zones (right)

(a) Nurses

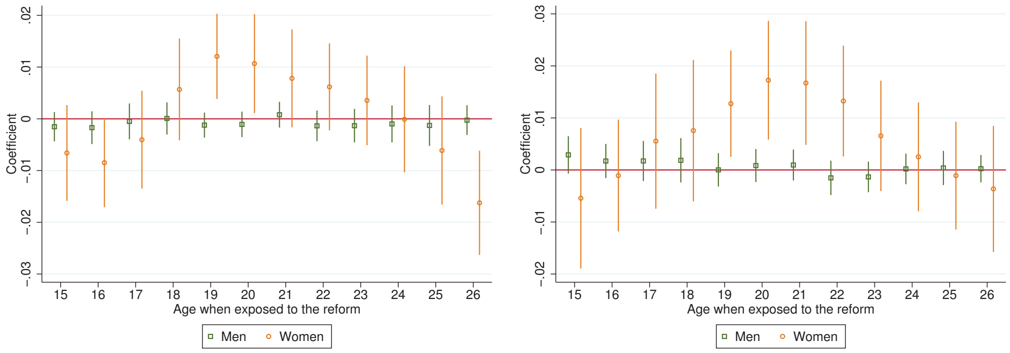

(b) Engineer

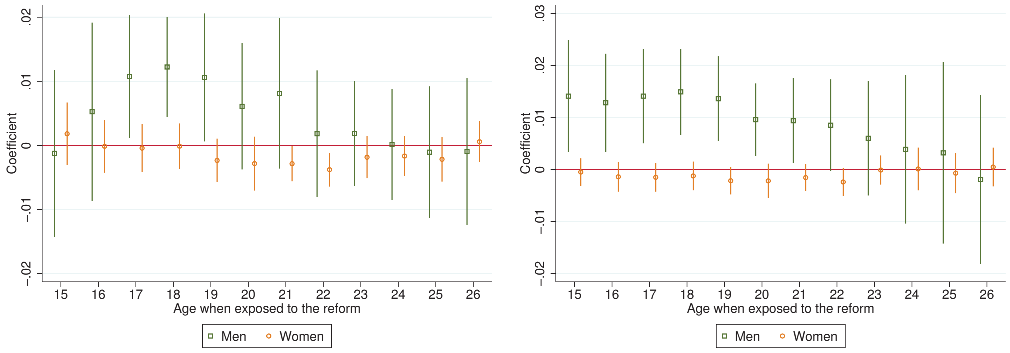

(c) Business administration

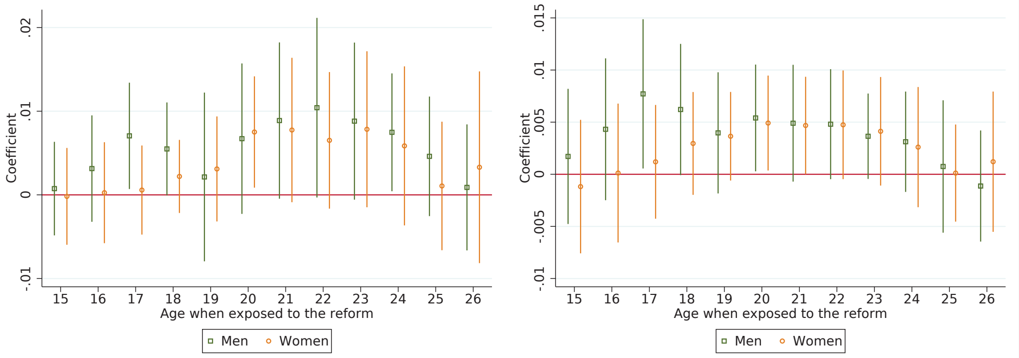

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalities with birth year) and a constant term. Standard errors are clustered at the municpality level.

6 Concluding comments

This paper shows that the 1960-1980 roll-out of higher education institutions in semi-rural Norway substantially impacted the choice of field for affected cohorts. This is found in a regression setup comparing individuals in the affected areas with older cohorts in the same locations, controlling for municipality-specific trends. Although we find large increases in take-up of the college degrees offered at the new colleges, our results suggest that the education level for those whose distance to a higher education institution decreased, did not change. Rather, we find decreases in the propensity to undertake college degrees not offered at the local colleges. Furthermore, there are no effects on incomes neither for men exposed to engineering colleges nor women exposed to nursing colleges. While engineering and nursing both led to clearly defined career paths, a degree in business administration did not. For business degree holders, men dominate the manager positions, while women dominate general office positions. We see this in our analysis as women with access to a college offering a degree in business administration, have lower earnings later in life than other women. The results emphasize the importance of examining field of study - and not only level of education - when major changes to the education structure are considered.

References

Aamodt, P. (1982). Utdanning og sosial bakgrunn [Education and Social Background]. *Samfunnsøkonomiske studier 51, Statistics Norway*.

Altonji, J. G., E. Blom, and C. Meghir (2012). Heterogeneity in human capital investments: High school curriculum, college major, and careers. *Annu. Rev. Econ. 4*(1), 185–223.

Andersen, S., P. dAstous, J. Martinez-Correa, and S. H. Shore (2020). Cross-program differences in returns to education and the gender earnings gap.

Andersson, R., J. M. Quigley, and M. Wilhelmsson (2009). Urbanization, productivity, and innovation: Evidence from investment in higher education. *Journal of Urban Economics 66*(1), 2–15.

Bhuller, M., M. Mogstad, and K. G. Salvanes (2017). Life-cycle earnings, education premiums, and internal rates of return. *Journal of Labor Economics 35*(4), 993–1030.

Bostad (2007). Nou 2007:6 formål for framtida- formål for barnehagen og opplæringen, Norges offentlige utredninger, Norwegian Ministry of Education.

Card, D. (1995). Using geographic variation in college proximity to estimate the return to schooling. *In L.N. Christofides, E.K. Grant, and R. Swidinsky, editors, Aspects of Labor Market Behaviour: Essays in Honour of John Vanderkamp Toronto: University of Toronto Press*.

Card, D. and A. A. Payne (2020). High school choices and the gender gap in stem. *Economic Inquiry*.

Carneiro, P. M., K. Liu, and K. G. Salvanes (2022). The supply of skill and endogenous technical change: evidence from a college expansion reform. *Journal of the European Economic Association*.

Chetty, R., N. Hendren, P. Kline, and E. Saez (2014). Where is the land of opportunity? the geography of intergenerational mobility in the united states. *The Quarterly Journal of Economics 129*(4), 1553–1623.

Clark, T. and Åsa Sohlman (2009). OECD reviews of tertiary education: Norway. Technical report, OECD.

Currie, J. and E. Moretti (2003). Mother's education and the intergenerational transmission of human capital: Evidence from college openings. *The Quarterly journal of economics 118*(4), 1495–1532.

Denning, J. T., B. M. Marx, and L. J. Turner (2019). Propelled: The effects of grants on graduation, earnings, and welfare. *American Economic Journal: Applied Economics 11*(3), 193–224.

Denning, J. T. and P. Turley (2017). Was that smart? *Journal of Human Resources 52*(1).

Evans, B. J. (2017). Smart money: Do financial incentives encourage college students to study science? *Education Finance and Policy 12*(3), 342–368.

Hastings, J. S., C. A. Neilson, and S. D. Zimmerman (2013). Are some degrees worth more than others? evidence from college admission cutoffs in chile. Technical report, National Bureau of Economic Research.

Helland, H. and K. Heggen (2018). Regional differences in higher educational choice? *Scandinavian Journal of Educational Research 62*(6), 884–899.

Johnsen, B. W. (1999). *Fra universitetsvisjon til høyskoleintegrasjon*. Ph. D. thesis, Kristiansand: Høyskoleforlaget.

Kirkeboen, L. J., E. Leuven, and M. Mogstad (2016). Field of study, earnings, and self-selection. *The Quarterly Journal of Economics 131*(3), 1057–1111.

Knutsen, T. K. (2017). Returns to field of study: Evidence from a norwegian reform of college expansion. Master's thesis.

Lindbekk, T. (1998). The education backlash hypothesis: The norwegian experience 1960-92. *Acta Sociologica 41*(2-3), 151–162.

Markussen, S. and K. Roed (2018). The golden middle class neighborhood: Trends in residential segregation and consequences for offspring outcomes.

Norwegian Ministry of Education (1969). St. meld. nr. 63 om virksomheten i statens lånekasse for utdanning i regnskapsåret 1969 [white paper to the parliament about the central government's loans to students].

Norwegian Ministry of Education (1975). St. meld nr. 17 om den videre utbygging og organisering av høgre utdanning [white paper to the parliament about the further expansion of higher education].

Norwegian Nurse Association (2017). History of norwegian nurses.

Ottoesen, K. (1969). St.prp. nr. 136 instilling om videreutdanning for artianere [recommendation on further education for high school graduates].

Røseth, D. (2003). Effektivitet og fordeling: Lånekassen som utdanningspolitisk redskap 1947-2003. Master's thesis.

Stange, K. (2015). Differential pricing in undergraduate education: Effects on degree production by field. *Journal of Policy Analysis and Management 34*(1), 107–135.

Statistics Norway (2000). Classifications of economic regions.

Suhonen, T. and H. Karhunen (2019). The intergenerational effects of parental higher education: Evidence from changes in university accessibility. *Journal of Public Economics 176*, 195–217.

Vangen, T. (2007). Nasjonal utdanningsdatabase dokumentasjonsrapport. *Statistics Norway Reports 2007/54*.

Vassenden, K. (1987). Folke- og bolitellingene 1960, 1970 og 1980 - Dokumentasjon av de sammenlignbare filene. *Reports 87/2, Statistics Norway*.

Appendix

A Year of establishment and municipality

Table A.1. Establishment year and municipality

Nursing		Engineering		Regional college	
Year	Municipality	Year	Municipality	Year	Municipality
1869	Oslo*	1855	Tønsberg	1969	Stavanger
1895	Lillestrøm	1873	Oslo	1969	Kristiansand
1906	Trondheim*	1875	Bergen	1969	Grimstad
1908	Bergen*, Skien/Porsgrunn	1878	Stavanger	1970	Molde
1918	Tønsberg	1884	Skien/Porsgrunn	1970	Volda
1920	Bodø, Grimstad, Stavanger	1897	Follo	1970	Lillehammer
1927	Elverum	1900	Trondheim	1970	Bø i Telemark
1939	Namsos	1955	Hønefoss, Narvik	1975	Sndal
1940	Levanger, Tromsø	1965	Fredrikstad, Ålesund	1977	Alta
1955	Fredrikstad	1966	Gjøvik	1977	Halden
1958	Molde	1967	Grimstad	1979	Rena
1960	Hammerfest	1970	Kongsberg	1980	Steinkjer
1970	Gjøvik	1981	Tromsø	1985	Harstad
1974	Ålesund	1988	Førde, Haugesund		
1976	Kristiansand	1989	Levanger		
1977	Stord				
1979	Førde				
1980	Haugesund				
1982	Harstad				

*In the case of Oslo, Bergen and Trondheim year refers to the year of establishment of the first nursing college as several nursing colleges were established in these cities.

A.1 Areas with college access before 1940

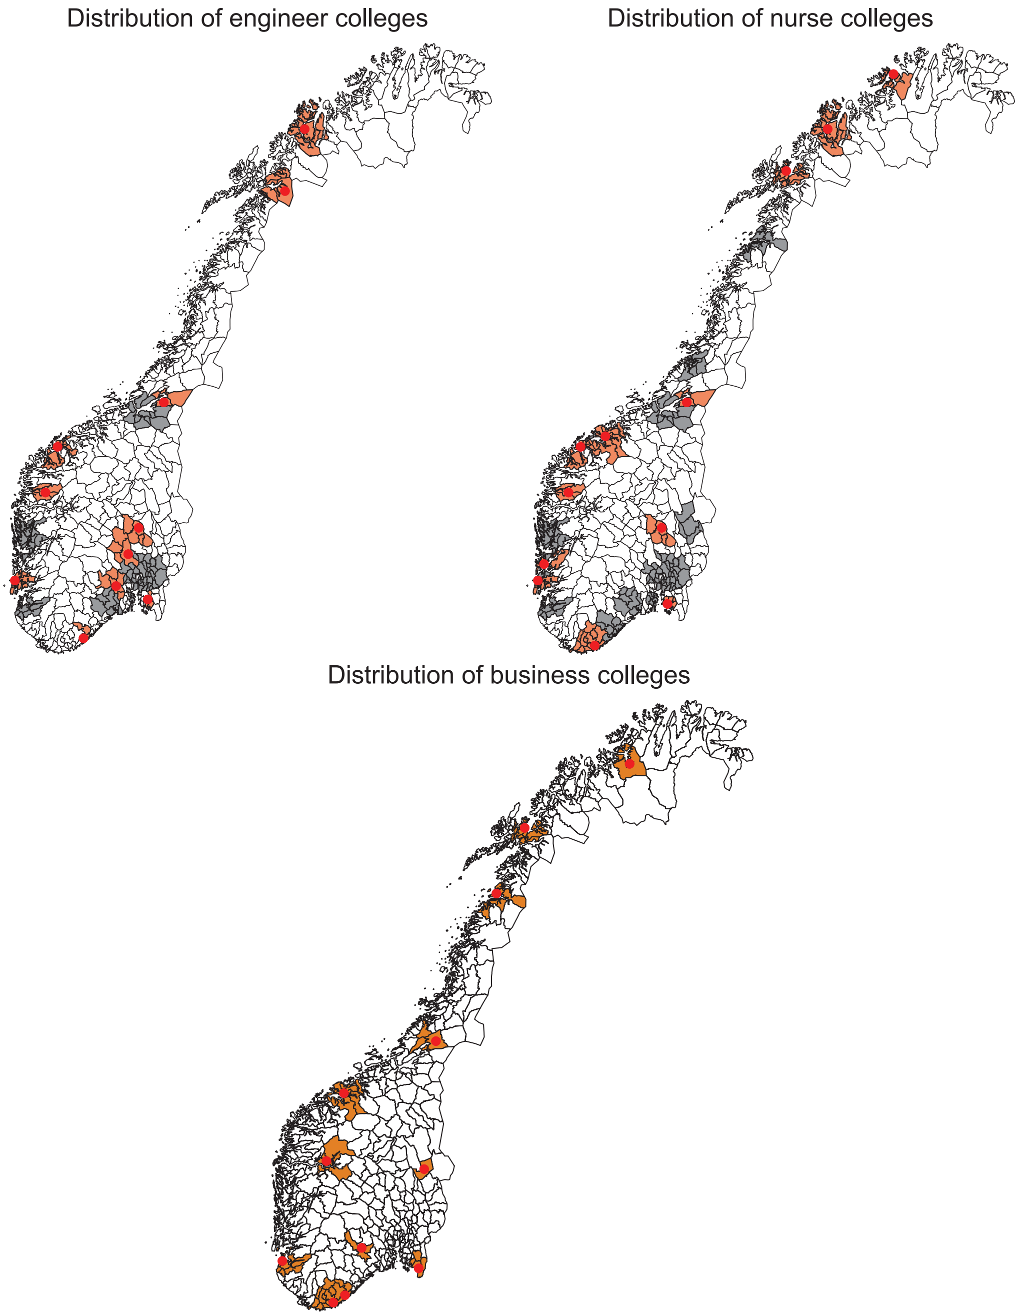

Figure A.1. Maps showing commuting zones for colleges established before and after 1940

Notes: The maps show the commuting zones that got new colleges after 1940 in orange, whereas the red dots indicate the exact college location. The grey areas are the zones with college access before 1940[a]

[a]There were no business colleges before 1940, therefore no grey areas.

B Year of establishment and municipality characteristics

Figure B.1. Timing of reform and municipality baseline characteristics: Nursing colleges

(a) Nursing colleges and 1950 characteristics: Plot of coefficient from eq 4.1

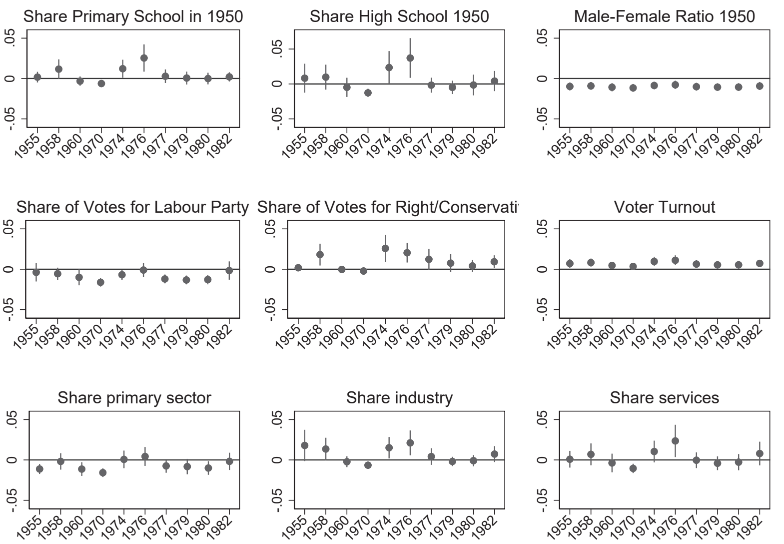

(b) Nursing colleges and 1960 characteristics: Plot of coefficient from eq 4.1

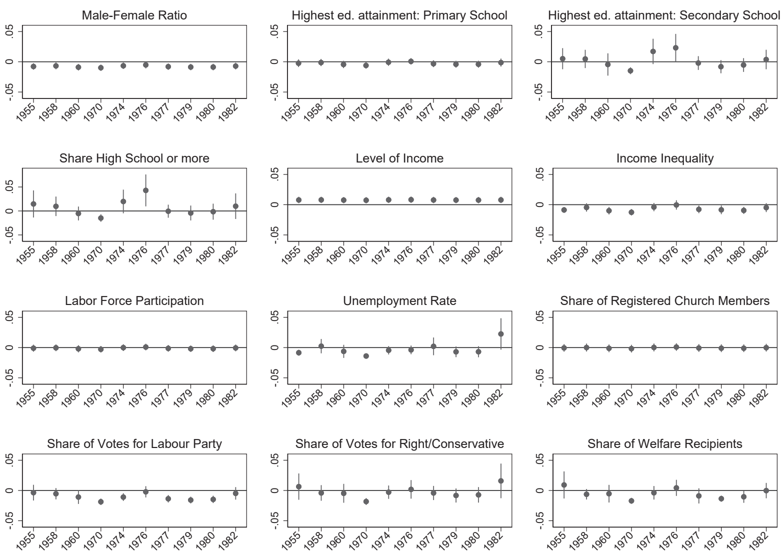

Figure B.2. Timing of reform and municipality baseline characteristics: Engineering colleges

(a) Engineering colleges and 1950 characteristics: Plot of coefficient from eq 4.1

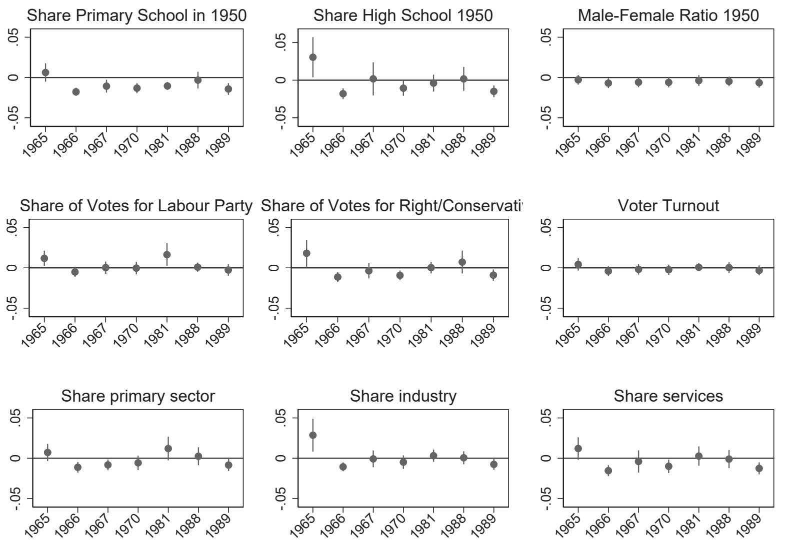

(b) Engineering colleges and 1960 characteristics: Plot of coefficient from eq 4.1

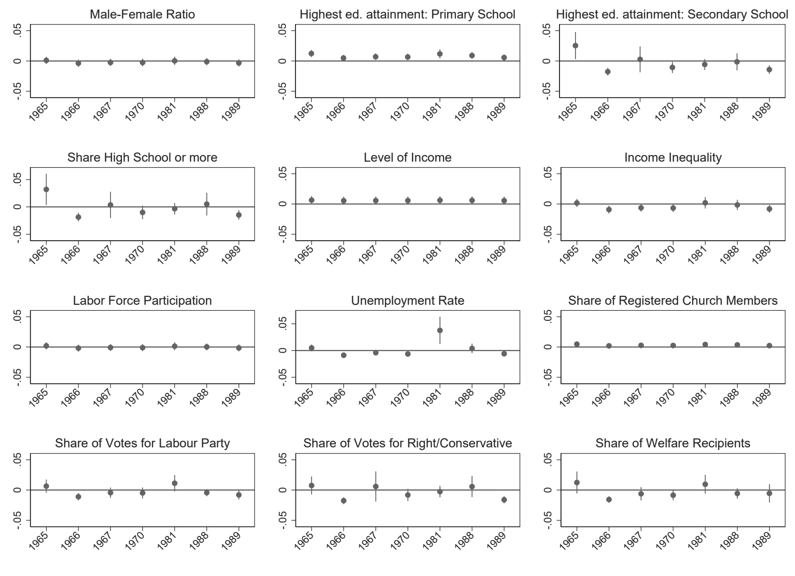

Figure B.3. Timing of reform and municipality baseline characteristics: Business colleges

(a) Business colleges and 1950 characteristics: Plot of coefficient from eq 4.1

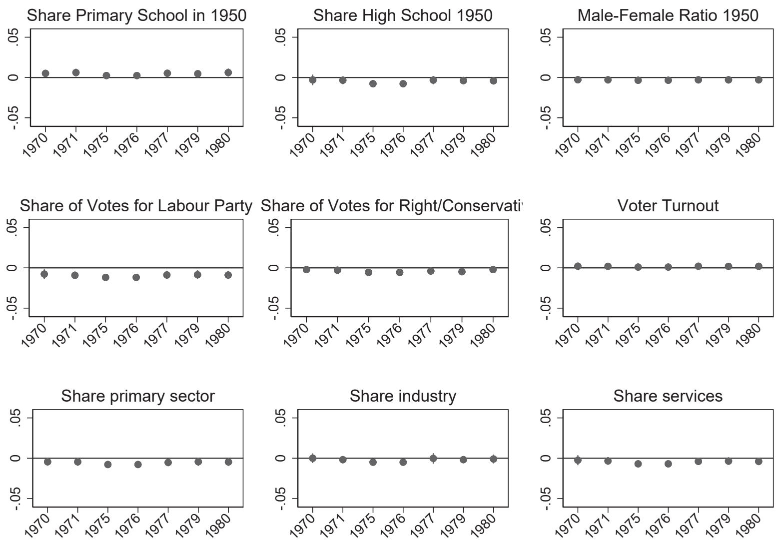

(b) Business colleges and 1960 characteristics: Plot of coefficient from eq 4.1

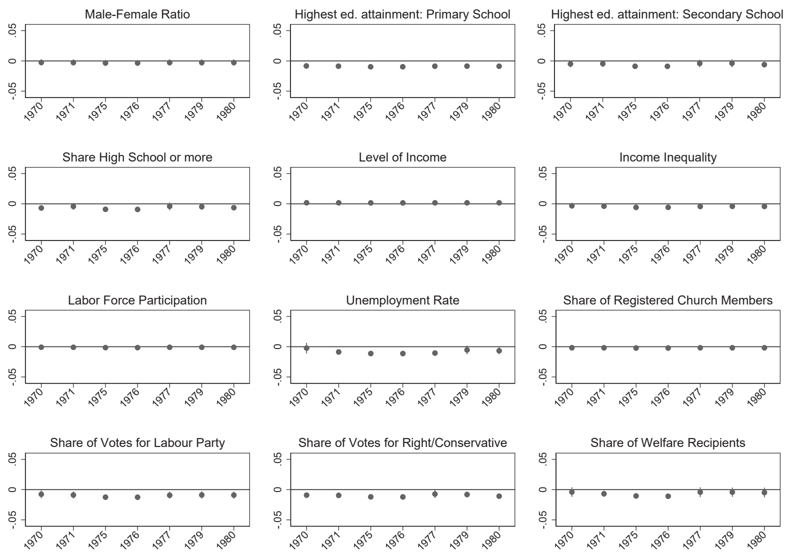

C Results: degree attainment and labour market outcomes

Table C.1. Baseline results for nursing

Age when	All		M low ed		M high ed	
reform	Wom	Men	Wom	Men	Wom	Men
implemented	(1)	(2)	(3)	(4)	(5)	(6)
<=15	-0.006*	-0.000	-0.007	-0.001	-0.004	0.001
	(0.004)	(0.001)	(0.005)	(0.001)	(0.006)	(0.003)
<=16	-0.008*	-0.002	-0.007	-0.002*	-0.006	-0.000
	(0.004)	(0.002)	(0.004)	(0.001)	(0.006)	(0.003)
<=17	-0.002	-0.001	-0.001	-0.002	-0.001	0.001
	(0.004)	(0.001)	(0.005)	(0.001)	(0.006)	(0.002)
<=18	0.006	-0.000	0.004	-0.001	0.009	0.001
	(0.004)	(0.002)	(0.004)	(0.001)	(0.005)	(0.003)
<=19	0.013***	-0.001	0.013***	-0.002	0.011*	0.000
	(0.003)	(0.001)	(0.004)	(0.002)	(0.006)	(0.002)
<=20	0.012***	0.000	0.009**	0.000	0.010	0.000
	(0.004)	(0.001)	(0.004)	(0.002)	(0.007)	(0.002)
<=21	0.009**	0.001	0.007*	0.003**	0.005	-0.001
	(0.004)	(0.001)	(0.004)	(0.001)	(0.007)	(0.002)
<=22	0.008**	-0.000	0.010***	0.001	-0.003	-0.003
	(0.004)	(0.001)	(0.004)	(0.001)	(0.006)	(0.003)
<=23	0.005	0.000	0.004	0.003**	-0.002	-0.004*
	(0.003)	(0.001)	(0.004)	(0.001)	(0.006)	(0.002)
<=24	0.002	0.001	0.001	0.003*	-0.006	-0.003
	(0.003)	(0.001)	(0.004)	(0.001)	(0.006)	(0.002)
<=25	-0.002	0.001	-0.002	0.002	-0.010	-0.001
	(0.004)	(0.002)	(0.004)	(0.001)	(0.007)	(0.002)
<=26	-0.004	0.001	-0.002	0.001	-0.011	-0.001
	(0.004)	(0.001)	(0.003)	(0.001)	(0.007)	(0.002)
N	325275	363664	182422	207922	142853	155742

Notes: Reported are point estimates and corresponding standard errors from Estimating Eq (1) for each age cut-off where we compare individuals below and above the particular age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term. Standard errors are heteroskedasticity robust and clustered at the municipality level. *p<0, 10, **p<0,05, ***p<0,01.

Table C.2. Baseline results for engineering

Age when reform implemented	All		M low ed		M high ed	
	Wom	Men	Wom	Men	Wom	Men
	(1)	(2)	(3)	(4)	(5)	(6)
<=15	0.000	0.010**	-0.001	0.004	0.002	0.016***
	(0.002)	(0.005)	(0.002)	(0.005)	(0.003)	(0.006)
<=16	-0.000	0.010**	-0.002	0.005	0.002	0.013***
	(0.001)	(0.004)	(0.001)	(0.005)	(0.003)	(0.005)
<=17	-0.000	0.011***	-0.002	0.003	0.002	0.019***
	(0.001)	(0.004)	(0.002)	(0.004)	(0.002)	(0.005)
<=18	-0.000	0.011***	-0.002	0.003	0.002	0.019***
	(0.001)	(0.003)	(0.001)	(0.004)	(0.003)	(0.004)
<=19	-0.001	0.010***	-0.002	0.005	0.001	0.016***
	(0.001)	(0.004)	(0.001)	(0.004)	(0.002)	(0.005)
<=20	-0.001	0.007**	-0.001	0.004	-0.001	0.012**
	(0.002)	(0.003)	(0.001)	(0.003)	(0.003)	(0.006)
<=21	-0.000	0.008**	-0.000	0.003	0.000	0.015**
	(0.001)	(0.004)	(0.002)	(0.003)	(0.003)	(0.007)
<=22	-0.001	0.006	0.000	0.004	-0.003	0.006
	(0.001)	(0.004)	(0.002)	(0.003)	(0.003)	(0.008)
<=23	-0.000	0.004	0.001	0.003	-0.002	0.004
	(0.001)	(0.003)	(0.002)	(0.003)	(0.002)	(0.007)
<=24	-0.001	0.002	0.001	0.003	-0.004	0.002
	(0.002)	(0.004)	(0.002)	(0.004)	(0.003)	(0.007)
<=25	-0.002	0.002	-0.001	0.002	-0.003	0.003
	(0.002)	(0.005)	(0.002)	(0.005)	(0.003)	(0.008)
<=26	-0.000	0.000	-0.000	0.002	0.001	-0.001
	(0.002)	(0.005)	(0.002)	(0.004)	(0.003)	(0.010)
N	325275	363664	182422	207922	142853	155742
Mean						

Notes: Reported are point estimates and corresponding standard errors from Estimating Eq (1) for each age cut-off where we compare individuals below and above the particular age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term. Standard errors are heteroskedasticity robust and clustered at the municipality level. *p<0, 10, **p<0,05, ***p<0,01.

Figure C.1. The reduced form coefficients of the college reform on labour market outcomes for individuals with lower educated mother

(a) Nursing and engineering: Formal participation in the labour force (early income above 187 000 nok)

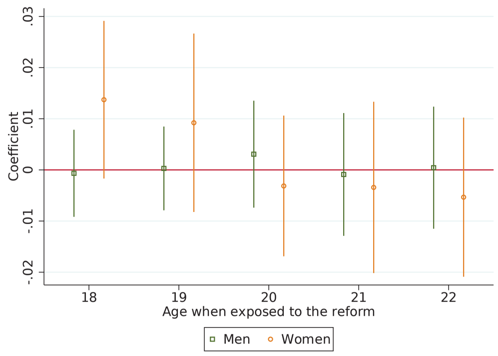

(b) Nursing and engineering: Equivalent full-time salary (yearly income above 467 500 nok)

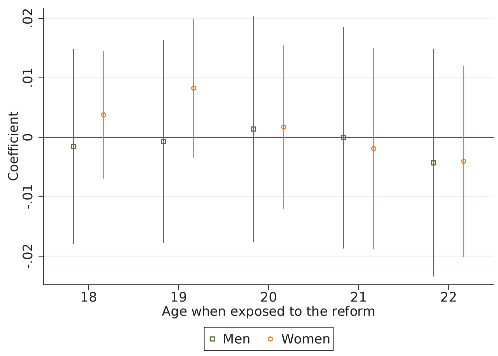

(c) Nursing and engineering: L of income

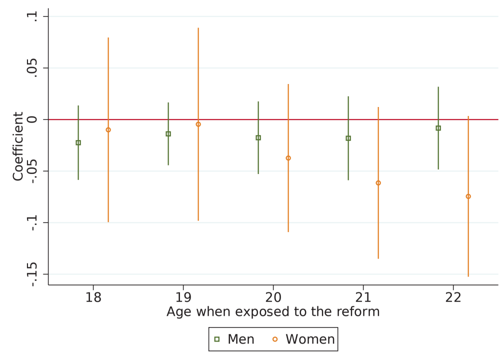

(d) Business: Formal participation in the labour force (early income above 187 000 nok)

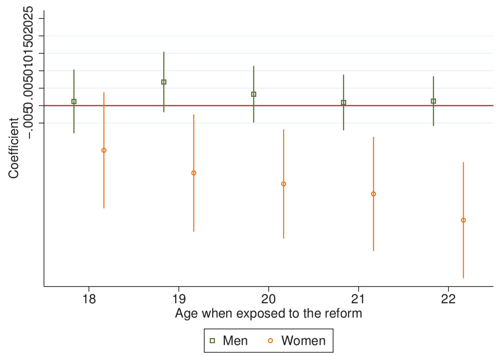

(e) Business: Equivalent full-time salary (yearly income above 467 500 nok)

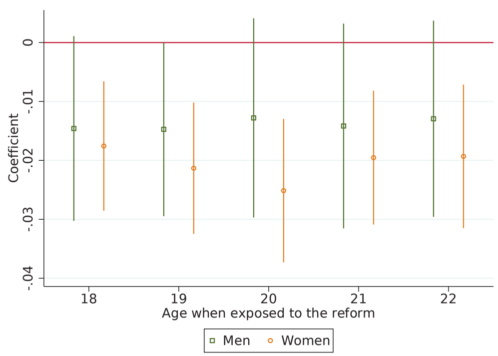

(f) Business: L of income

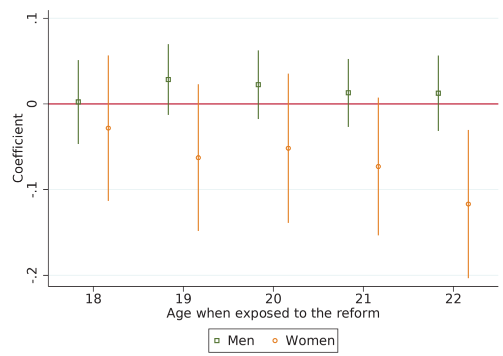

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term. Standard errors are clustered at the municpality level.

Figure C.2. The reduced form coefficients of the college reform on labour market outcomes for individuals with lower educated mother

(a) Nursing and engineering: Formal participation in the labour force (early income above 187 000 nok)

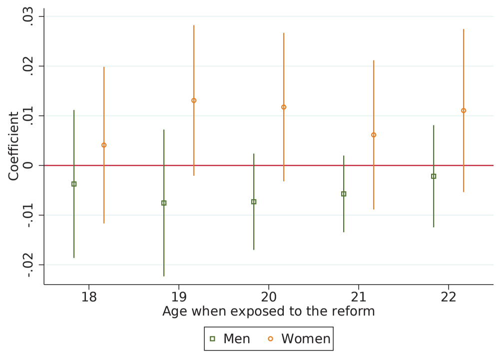

(b) Nursing and engineering: Equivalent full-time salary (yearly income above 467 500 nok)

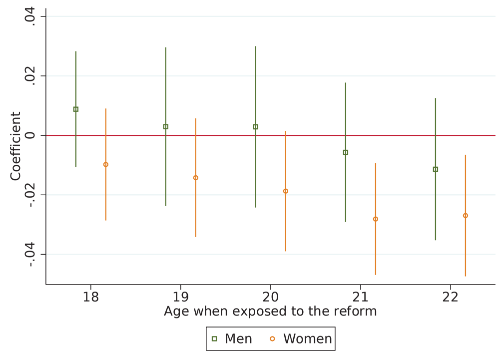

(c) Nursing and engineering: L of income

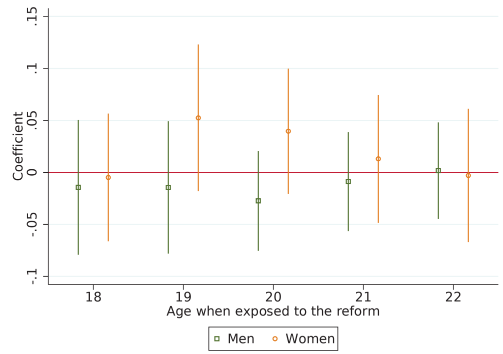

(d) Business: Formal participation in the labour force (early income above 187 000 nok)

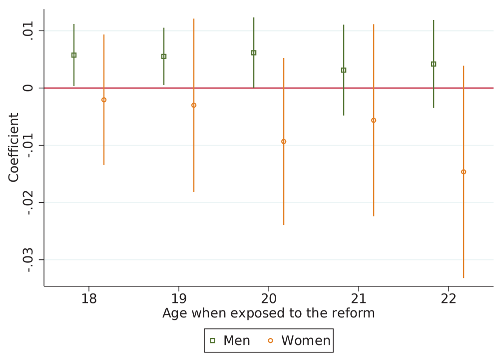

(e) Business: Equivalent full-time salary (yearly income above 467 500 nok)

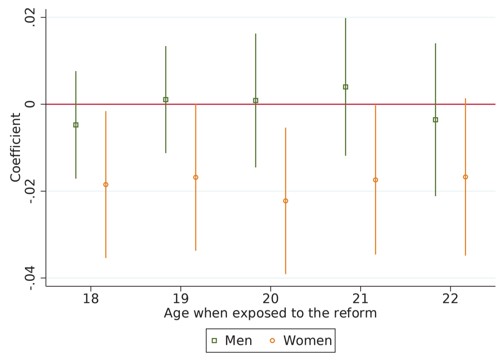

(f) Business: L of income

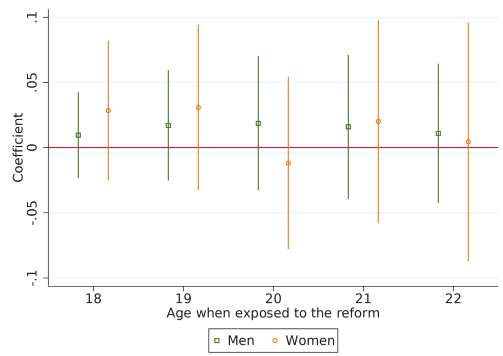

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term. Standard errors are clustered at the municipality level.

Figure C.3. The reduced form effects of the the opening of new nursing and engineering colleges on the probability of staying in the same municipality or region at age 35

(a) Municipality

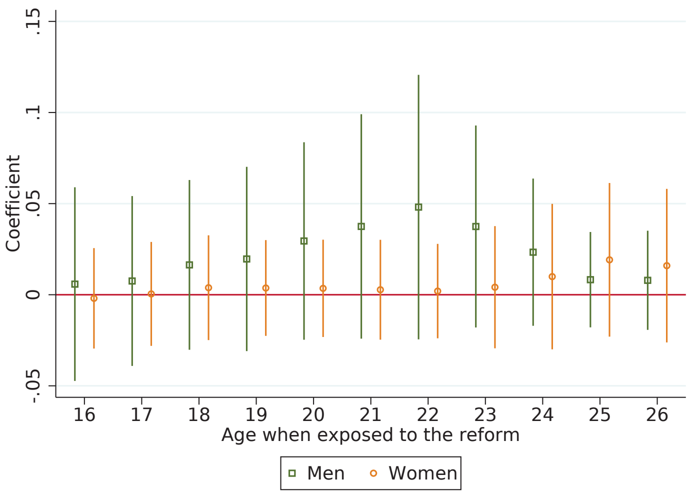

(b) Region

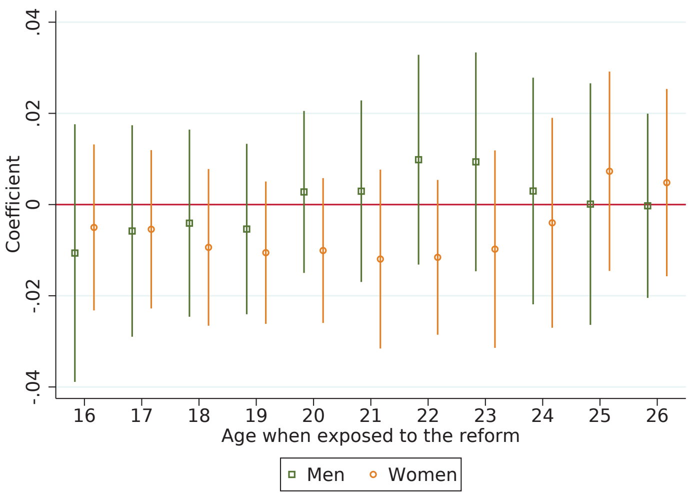

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term.

Figure C.4. The reduced form effects of the the opening of business colleges on the probability of staying in the same municipality or region at age 35

(a) Municipality

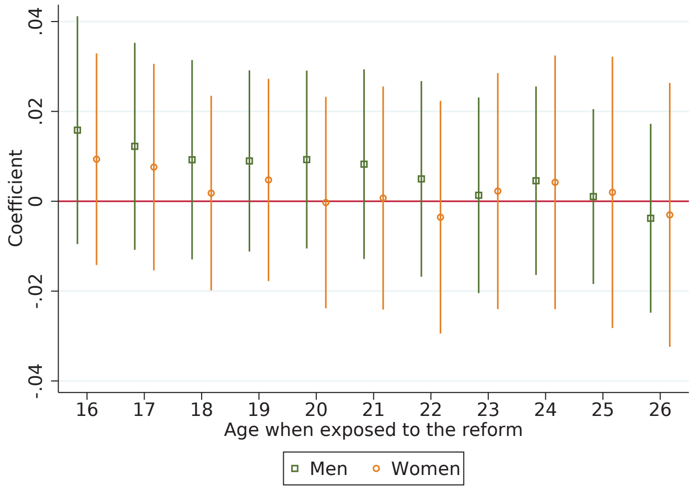

(b) Region

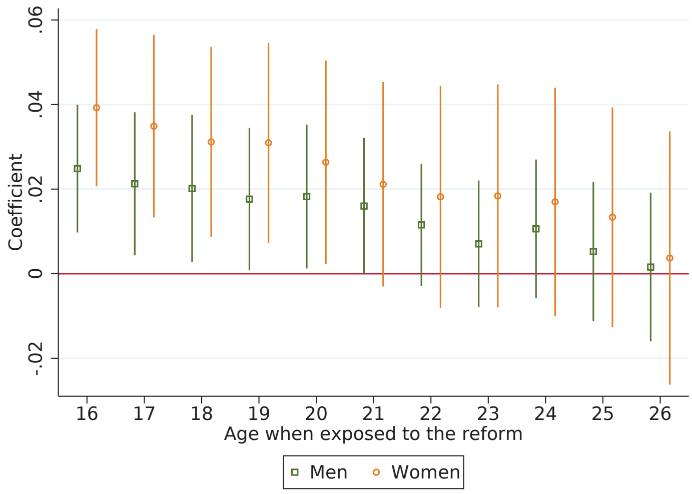

Notes: The x-axis shows the age cut-off, and each line report the point estimate and corresponding 95 percent confidence interval from comparing individuals below and above that age cut off in treated and untreated areas. Included in all specifications are municipality fixed effects, dummy variables for birth year, compulsory schooling reform, and municipality specific time trends (where we interact municipalites with birth year) and a constant term.Standard errors are clustered at the municpality level.

Table C.3. Professions, contingent on having a degree in business administration

	Men	Women
Occupations classified as management positions		
Administrative and mercantile leaders	1245	613
Commodity production and service units	209	97
Occupations classified as academic professions		
ICT advisors	592	649
Realister civil engineers and such	330	116
advisors in finance, administration and sales	96	140
Occupations classified as college careers		
Professions within the culture and sport sector	969	1059
Engineers as such	236	75
employees in finance, administration and sales	85	143
Other occupations		
General office workers.	256	713
Other and unknown	1479	1686
Nr of observations	5497	5291

CHAPTER 3

The effect of introducing industry specific minimum wages

The effect of introducing industry specific minimum wages*

Tora Knutsen

Abstract

What is the effect of extending the coverage of wage floors? In this paper, I study the introduction of legally enforced collective bargaining agreements in three industries in Norway: shipyards, construction, and cleaning. The findings suggest that despite relatively high wage floors, workers in construction and cleaning on average experienced lower wage growth than comparable workers after the change in wage policy. Using rich administrative data on employment, demographics, and benefits, I track individuals' wage and labor market attachment over a period of 15 years. To consider omitted variables at the individual and at the industry level, I use two different data driven methods. Firstly, I run an individual fixed-effects model and add control variables flexibly using a double-post Lasso estimator altering results only marginally. Secondly, I construct synthetic control industries to control for omitted variables at the industry level. Neither adding control variables nor using the synthetic control method change results substantially, but several robustness checks underline the uncertainty in studying a policy implemented in the whole industry at the same time. Exploring the mechanism at the individual level, suggests that the decrease in wages is correlated with experience and union membership. Moreover, I find no evidence that employment or the probability of changing industry are affected by the wage floor policy, but cleaning and construction workers do seem to experience small increases in the probability of receiving disability benefits.

*For this work I acknowledge support from the Norwegian Research Council project "Alternative work arrangements and worker welfare" (grant nr. 314267). In addition, I would like to thank Jo Thori Lind, Andreas Kotsadam, Kalle Moene, Gøril Andreassen and Katarzyna Segiet for helpful comments.

1 Introduction

Minimum wage policies have featured prominently both in policy debates and in economic research. Proponents suggest that the minimum wage is an efficient measure to increase wages for low-wage earners and to decrease poverty. Opponents to wage regulations claim that the minimum wage reduces employment since the price of labor input increases, hence harming those with low incomes.

For workers in several European countries, the relevant wage floor is not a national minimum wage, but the wage set by collective bargaining agreements. In many countries, wage floors stemming from sectoral bargaining are extended by law to non-covered firms. In contrast to national minimum wages, sectoral bargaining is not only designed to establish a wage floor. Wages are set based on age and experience to ensure that workers get regular raises beyond the minimum wage floor. In the Nordic countries, the coverage of such agreements is based on voluntary membership in labor unions and employer confederations. While both voluntary and legally enforced collective bargaining coverage have been studied extensively, less is known about the effects of a transition from a voluntary to a legally enforced wage floor.

In this paper, I study the wages and labor market trajectories in three specific industries that introduced minimum wages in Norway. The results suggest that introducing a wage floor has mixed effects: While workers with experience seem to have less wage growth after its introduction, workers new to the industry have higher wage growth than comparable workers in other industries.[1] However, all results are subject to uncertainty. As the policy was implemented throughout a whole industry at the same time, other factors, unrelated to wage policy, may have affected wages. Even though this is fundamentally difficult to account

[1] As hourly wages are not registered directly in Norway, the main outcome variable is annual earnings. In parts of the analysis I also use hourly earnings with different ways of controlling for substantial measurement error in hours worked.

for, I attempt to control for such factors using post estimation Lasso and the synthetic control method. Yet, conclusions regarding effects of the wage policy remain uncertain. It is likely that the policy did reduce or at least did not raise wages in two of the industries; cleaning and construction. In the third industry, shipyards, wages seem to have increased, but these wage changes seem not to be timed according to the wage floor introduction.

Like the other Scandinavian countries, Norway has never had a national minimum wage. Instead, 70% of the workers in private industries work for employers covered by a collective agreement.[2] These agreements are not legally enforced, implying that paying very low wages is not illegal for employers not covered by a collective bargaining agreement.

Since the enlargement of the European Union in 2004, with new countries included in the European labor market, legally enforced wage floors have been introduced in several industries in Norway to avoid too large wage differentials between native and immigrant workers.[3] These wage floor introductions provide a unique setting for studying the effect of transitioning from voluntary to legally enforced wage floors.

Labor unions in Norway oppose the introduction of a national minimum wage, claiming that it will lead to a downward pressure on wages, by creating a wage "roof" instead of a wage "floor".[4] Despite their opposition to a minimum wage, unions have advocated for collectively

[2]Many non-member firms also follow these agreements so that even workers in non-covered firms may be paid according to the bargained scheme. This number is therefore higher in practice. The coverage of collective bargaining is not high compared to other European countries. 13 countries have higher collective bargaining coverage, including all Nordic countries, France, Italy, and Spain (https://ilostat.ilo.org/topics/collective-bargaining). A requirement for claiming a collective agreement is that 50% of workers are union members, but 10% is sufficient when the employer is member of an employer's organization. Although 73% of workers in the private sector work for an organized employer, only 38% are themselves a member of a union. In the public sector, all workers are covered by collective agreements and 80% are labor union members. Information on membership and coverage is available at https://www.arbeidslivet.no/Lonn/Fagorganisering/tariffavtaledekning/

[3]Among the new members states eight were countries with wages significantly lower than in Western Europe: Estonia, Latvia, Lithuania, Poland, Czech Republic , Slovakia, Slovenia, Malta, Cyprus, and Hungary. As EU member states, they became part of the inner market, entailing free movement of labor, goods and services. Many in Norway and Western Europe feared that increased labor migration would undermine wages and labor rights. Appendix A.1 provides more details on the institutional setting and background for minimum wage introductions.

[4]An example of their resistance to minimum wages in general is Norwegian union's opposition to the EU

bargained wage floors to be extended to non-covered firms in sectors experiencing increased labor immigration. This was considered as a way of getting rid of “very low” wages without weakening the importance of wage bargaining and leveling the playing field for native and non-native workers.

As a result of union advocacy, the wage determined through bargaining between the employer confederation and the labor union has been extended to cover all workers and employers in certain sectors subject to large inflows of foreign workers. I will focus on the first three industries introducing wage floors: construction, shipyards, and cleaning.[5] In these, minimum wages were introduced in 2006/2007, 2009 and 2012, respectively. The study of the different industries offers insight into the introduction of a mandatory wage floor in very different parts of the labor market. While cleaning has a large proportion of low-wage work, this share is negligible in the shipyard industry and low in construction. The minimum wage to median wage ratio is between 0.5 and 0.65, highest in cleaning, and lowest in the shipyard sector. In an international context, these ratios are fairly high.[6]

The study of wage floors in Norway is one among few studies of the introduction of such extensions. The same reform has been studied in Norwegian publications with different methodological approaches as well as their research questions. Bratsberg and Holden (2015) find a positive effect on wages in the construction industry, using geographical variation in the timing of the minimum wage in this industry, while Skjerpen et al. (2015), in a policy

Directive on statutory minimum wages where they together with Swedish and Danish unions have aligned with the confederation of European Business and Nordic Governments in opposing this directive. In response, the EU has promised not to force countries with a high degree of collective bargaining to implement a national minimum wage (https://fagbladet.no/nyheter/full-splittelse-om-minstelonn-frykter-alvorlig-skade-pa-den-nordiske-modellen-6.91.730179.f242d77c64).

[5]Legally enforced wage floors were also introduced for a number of petroleum plants in 2004, but were taken away in 2008 as it was considered that they were no longer needed. As this covered only a few employers, it would be difficult to identify the workers, and therefore petroleum plants were not included in the study. In 2010, mandatory wage floors were introduced in the agricultural sector. Studying agriculture proved difficult since the share of seasonal workers is large and employers were before 2015 not required to register short-term jobs in the employer-employee registry (Statistics Norway, 2016).

[6]The minimum wage to the median wage ratio is often called the Kaitz index. France has among the highest ratios in the OECD (0.61) and the U.S has the lowest ratio (0.32).

evaluation report, find positive effects on the shipyard industry and the construction industry, but no effect on wages in the cleaning industry.

For a large number of workers in many European countries, the coverage of collective bargaining remains high. As the degree of union membership has declined, countries such as Germany have introduced national minimum wages to compensate for the decline in sectoral bargaining coverage. In several European countries, including France, Germany, Italy and Spain, collective agreements are extended by the country's authorities beyond their subscribing employer associations, unions, and their affiliated workers. Thus, a collective agreement is extended to cover all workers in the industry for which the agreement was signed. In these countries, there may be a large discrepancy between union density and union coverage. In France, for example, 98% are covered by a collective agreement, but only 10% of workers are union members.

While collective bargaining and minimum wages are well-studied, less is known about the wage impacts of extending such collective agreements.[7] An exception is the Portuguese setting where the scattered timing of agreement extensions have been used as variation in agreement coverage. Martins (2021) and Card and Cardoso (2022) both find that wages increase following an extension. Card and Cardoso (2022) find a 50% pass-through rate of floor increases, but no significant effect on employment, while Martins (2021) find negative effects on employment.

The traditional view that the minimum wage automatically lowers employment has been challenged by a range of studies that find little or no effect on the number of jobs.[8] Yet,

[7]The effect of collective bargaining on wages and other outcomes in Norway is well-studied. For example, Holden (1998) find that wage drift does not offset the effect of central negotiations and Calmfors and Nymoen (1990) find that wages are responsive to unemployment in the Nordic countries. More recent papers have shown that influential unions mitigate the inequality inducing effect of performance pay (Barth et al., 2012). (Barth et al., 2020) use changes in tax subsidies for union members to study the causal link between union density, wages and productivity. They find that increasing union density at the firm level raise wages and productivity.

[8]This has been called "New minimum wage research" starting with Card and Krueger (1994). Some examples of recent studies include Cengiz et al. (2019) and Dube et al. (2010) in the U.S, Ahlfeldt et al.

there are studies detecting effects for groups of workers including those with below-median experience (Jardim et al., 2018), low-skilled (Clemens and Wither, 2019) and teenagers (Neumark and Wascher, 1992). The magnitude of the effect is less prominent when comparing workers within low wage industries (Neumark and Shirley, 2021). As the effect, if any, has often been found for particular sub-groups such as teenagers, Manning (2021) states that the employment effect is "elusive."

Using Norwegian administrative register data including all employees in Norwegian companies, I trace workers' wages and labor market trajectories over a period of 15 years. This enables me to control for all time-invariant unobserved individual characteristics. To study the effect of introducing a minimum wage, I first use a difference-in-difference strategy with individual fixed effects to compare wage growth among workers in the targeted sectors to that of other workers. Surprisingly, workers in the construction and cleaning sectors experience lower annual wage growth after the minimum wage introduction compared to other workers. While individual fixed effects control for the correlation between the outcome and any background variable that is constant over time, it does not capture whether segments of workers experience specific trends over time, not linked to changes in wage policy.

To capture such underlying trends both at the individual and at the industry level, I combine two data-driven methods used in previous minimum wage research; a post Lasso-estimator (Allegretto et al., 2017) and the synthetic control method.[9] The post Lasso-estimator applies Lasso to select the most important predictors for the outcome and the treatment.[10] Belloni et al. (2014) suggest running a simple OLS regression of the outcome on the treatment and the double-selected set of controls (hence the term "post-LASSO").

(2018) and Dustmann et al. (2021) in Germany and Engbom and Moser (2021) in Brazil.

[9]See for example Allegretto et al. (2013, 2017), Dube and Zipperer (2015), Jardim et al. (2017), Neumark and Wascher (2017), Powell (2021), and Reich et al. (2017).

[10]The Lasso regression is machine learning method that obtains a subset of predictors minimizing the prediction error by imposing a penalization on included predictors(see e.g. Angrist and Frandsen, 2022; Athey and Imbens, 2019; Mullainathan and Spiess, 2017).

Norwegian administrative data includes a very rich set of background variables on educational field and length and immigrant background, well-suited to control for such trends. However, adding all (potentially) relevant controls leads to overfitting. Therefore, I use post-estimation Lasso to select a set of covariate-specific trends. Adding these does not change the results qualitatively. Moreover, the 'wage-floor industries' may form part of a subsection of the economy that has had less wage growth. To find similar industries, I construct synthetic control industries consisting of a weighted combination of non-affected industries with similar wages prior to the minimum wage introduction. Applying the synthetic control method, the effect on wage growth in the shipyard industry turns positive, but several robustness checks question the validity of this result. For the cleaning and the construction industries, the changes in wage growth remain negative.

In the last part of the paper, I explore which groups experienced less wage growth after the minimum wage introduction by comparing workers in the targeted industries to workers in synthetic control industries. The results at the individual level are subject to larger uncertainty due to the selection process of the sample. Yet, the results do indicate that the workers experiencing less wage growth in cleaning and construction are full-time workers. In cleaning, the negative effect is driven by low-wage workers while this is not the case in construction. Decline in wage growth is also associated with union membership and experience. Thus, the zero or negative change in wage growth for construction and cleaning workers could come from a lower relative growth (or decline) in the return to experience. Moreover, I find no evidence that employment is affected by the wage floor policy. At the individual level, the probabilities of exiting the labor market or changing industry seem not to change.

The rest of this paper is structured as follows. Before turning to the data analysis, I briefly discuss the data sources and variable definitions in Section 2. Moreover, I present some descriptive statistics on the industries subject to the wage policy. In Section 3, I discuss the empirical implementation. Section 4, the result section, is divided into three parts: In Section 4.1 I present the baseline results with and without control variables from the Lasso estimation and Section 4.2 the results from the synthetic control method. In Section 4.3 I compare workers in the synthetic control industry to workers in the wage floor industries to explore mechanisms for the findings.

2 Data and descriptive statistics

To study the introduction of a wage floor, ideally one would use hourly data since the wage floor is linked to hourly rates. However, as described below, hourly data in Norway is not very reliable because hourly wages are not recorded directly. Instead, individuals are recorded with regard to their employer, their annual salary, and a rich set of demographics. In this section, I describe how I link individuals to an employer, how I calculate their annual wages, and how I calculate hourly wages in the analysis. Moreover, I show that the three industries included in this study have starkly different wage levels and demographics.

2.1 Data

2.1.1 Data Sources

To study the introduction of minimum wages, I use data linking every individual to an employer and a annual salary. More specifically, the main data source is administrative data records including all employer-employee relationships in Norway from 2003 to 2018.[11] In 2015 and onward there has been major changes in the way the data is collected at the administrative level, but that does not have large consequences for this study.[12] Appendix Figure A3 shows a discontinuity in the standard deviation of hourly and annual wages in the three industries in 2015. Since all minimum wage introductions happen before 2015

[11]Data is available from 1999 onward, but due to major data improvements, in particular for hourly wages, I use data starting in 2003.

[12]In particular, the number of people appearing in the data records changes and the measurement of hours worked seems to change, resulting in changes in hourly wages. Figure A4 in the Appendix shows the number of observations by year and the increase between 2014 and 2015. However, I also show that the CVLASSO as described below and in detail in Appendix A.2.1 corrects in part for this increase.

and all estimates include year-fixed effects, this is not causing substantial problems in the analyses. However, some graphs on hourly wages therefore stop in 2014, as changes in the median of hourly wages after that year might be caused by a change in the administrative data collection and therefore does not represent actual changes in wages.

All employers are coded with an industry code (NACE) which is crucial for identifying the affected industries as well as following them over time. This classification was changed between 2008 and 2009. In Appendix A.2.2 I list the conversion between the new and the old NACE classification.[13]

Using unique identifiers for each individual, I link this data to several other data sources. Most importantly, to the Central Population Records which contains entries for all residents in Norway. The registries include background variables such as age, country of origin, immigration status, gender and education. Information on educational enrollment and degrees is collected from the Norwegian Education Data Base (NUDB). This data base contains individual-level data on all education undertaken since 1970. Educational choices and attainment are reported by the education institutions directly to Statistics Norway thereby minimizing any measurement error due to misreporting. However, immigrant's educational background is not fully reported on. While information on education is missing for 4% of all workers, it is missing for 23% of those born abroad and 40% of those from countries that entered the EU countries in 2004. In the regressions, I impute average education levels for these individuals and add a dummy for missing education.

The coding of education has a high level of detail, allowing the identification of educational field. In addition, the data base includes data on high school track. In high school, students choose between an academic or a vocational track. In the vocational track, students obtain a certificate/diploma (”Fagbrev/svennebrev”) after having completed two years in

[13] I made the conversion for the purpose of this study, so there might be errors making it less suitable for other studies

school and a two year apprenticeship. This is a variable of interest because bargained wages in industries that employ workers with vocational training depend on the completion of a "fagbrev". In addition individuals are linked to the Social Security database for unemployment and disability status. Variable definition and data sources are listed in Appendix Table A1.

2.1.2 Sample Selection

The baseline sample includes all workers in Norway between 25 and 62 years old.[14] One individual may have several observations per year, either because they have more than one job, they change job over the year, or, most often, because their contract or working hours have changed. This means that there are many duplicates in terms of employer-employee combinations.

To reduce the data to one observation per individual per year, I aggregate the salary and number of days worked for each employer-employee combination. For those individuals who still have more than one observation per year, I choose the job with the highest salary and, next, with the highest number of contracted hours. Lastly, if there are still duplicates, I choose the job that was valid on October 1st of that year. The large majority of workers match with a single employer in a given year.

2.1.3 Hourly wages

In Norway, wages are recorded annually at the employee-employer level. This means that hourly wages are not registered directly and must be calculated using annual wages and the total number of hours worked. Hours worked is however measured with errors, resulting in less reliable data on hourly wages. As the minimum wage increase happens at the margin of hourly wages, I still use this variable in parts of the analysis with some modifications: I identify the observations for which the hourly wage is measured with a large error and

[14] Appendix Figure A5 shows that the overall trend in the industries subject to a minimum wage is a decreasing proportion of workers under 25 years and an increasing proportion of workers 62 and older. The trends are similar to that of the whole workforce.

include dummies for wages that are potentially poorly measured.

To estimate the degree of measurement error, I first predict log hourly wages using cross-validation Lasso. This means that the data is repeatedly partitioned into training and validation data. The model is fit to the training data and the validation data is used to calculate the prediction error (Ahrens et al., 2020; Athey and Imbens, 2019; Chetverikov et al., 2021). Next, I use the residuals to identify observations that have a residual more than 3 standard deviations from the mean residual. With this measure, around 5% of the observations of hourly wages are measured with substantial error. As an alternative, I also use a Mincer-like OLS-regression with only education and experience as predictors. With this method, around 3% of observations have residuals more than 3 standard deviations from the mean. Both methods suggest that most wages measured with substantial error are biased downwards.

Although the degree of measurement error seems not be linked to the wage floor implementation, Figure A2 shows that the proportion with error increases over time in the cleaning and construction industries. Thus, the degree of substantial measurement error seems not to be orthogonal to year nor to industry. The share with error is high in cleaning, average in construction and low in the shipyard industry. Moreover, the degree of substantial error is larger for immigrants than for natives: For immigrants the proportion of observations with error is 20% with Lasso and 3% with OLS, for natives it is around 2% with both methods. It is important to note that in the OLS estimation, country background is not included, while it is part of the admissible set of predictors in Lasso. Ultimately, these methods result in a large decrease of the variation in the outcome variable as shown in Appendix Figure A3. In the case that measurement error is less than these methods suggest, decreasing variation in an outcome variable, might lead to false positives

In this paper, I focus on annual and not hourly earnings since hourly earnings seem not to be a reliable measure. In the few instances where I do use hourly earnings, I control for measurement error in a flexible way; I make a dummy for measurement error with OLS, one for measurement error with Lasso, and an interaction term. I also check that all results are robust to dropping observations potentially measured with error. Moreover, when calculating median hourly wages, I exclude wages with substantial measurement error (with Lasso). Appendix Section A.2.1 includes more details on measurement error in hourly wages and how this is handled.

2.2 Descriptives

Minimum wages were introduced in three very different parts of the labor market. As shown in Figure 1b, the median hourly wage in the shipyard industry is 20% above the national one, while it is 20% below in the cleaning industry. Despite the low industry specific median wage as shown in Figure 1b, Figure 1a shows that cleaning has the *highest* minimum wage to median wage ratio, while the shipyard industry has the *lowest*.

In all industries, the 2004 enlargement of the European Union led to a stark increase in the share of workers from the new member countries. As shown in Figure 2a, this increase was most prominent in the cleaning industry, going from less than one percent in 2003 to almost one third 15 years later. The share of workers born outside Norway was already at a much higher level in cleaning than in the other two industries (see Figure 2b). In the time period included in this study, employment increased in construction and cleaning, but decreased in the shipyard industry (see Figure 2c), while the average number of hours worked per year stayed fairly stable (see Figure 2d).

Low wage workers, defined here as those earning less than 60% of the national median income, are particularly prevalent in the cleaning industry with almost half of workers in this category.[15] Although low pay is linked to not working full time, around 30% of full-time

[15] Low income status is often measured relative to the national median wage. A commonly used threshold is having an income less than 60% of the average median income. This is the standard measure in the European Union, while the OECD uses 50% instead. In Norway, a household is usually defined as low income if the low-income status persists for at least three years. This is not taken into account here.

workers are also low-waged. The proportion of low-wage workers in the construction industry is fairly similar to that of the whole population and slightly higher than in the shipyard industry. For the entire work force, about 10% of all workers and 5% of those working full-time are low-waged according to the same definition. Appendix Figure A6 shows that share of low wage workers among immigrant employees has increased in all industries over time.

Figure 1c shows that for a full-time worker, minimum wages in the construction and cleaning industry would give a annual wage of slightly more than 60% of the median wage, whereas it would not in the shipyard industry. However, this is before any additional pay for overtime or inconvenient working hours is included. Thus, the minimum wages could potentially move workers out of a low-wage status. For the entire work-force, the proportion of low-wage workers has remained stable over time, with a small increase for immigrant workers (see Appendix Figure A6).

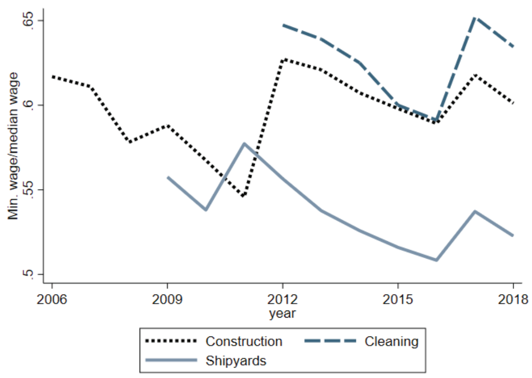

(a) Minimum wage relative to national median hourly wage

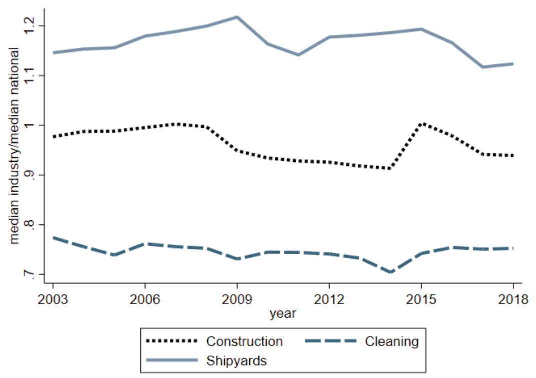

(b) Median wage in industry relative to national median hourly wage

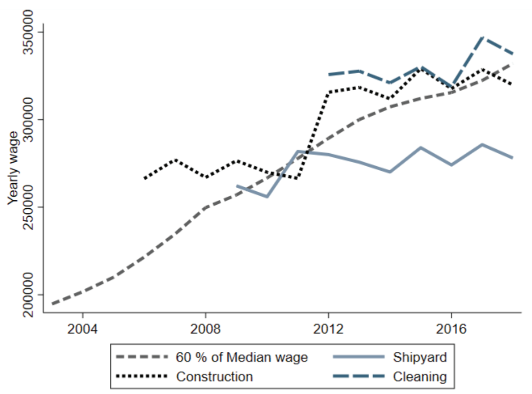

(c) Minimum wages compared to 3/5 of the median wage

Figure 1: Wages: The development of median wages and minimum wages

Notes: The upper left and upper right panel show how the national median wage compares to the industry specific wage floor (left) and the industry specific median wage (right). The bottom panel shows how the yearly wage (working 1950 hours) corresponding to the industry specific minimum wages compares to 60% of the national median wage, which is a measure for low income status. Lines related to the minimum wage start in the respective year of the wage floor introduction: 2006/2007 (Construction), 2009 (Shipyards) and in 2012 (Cleaning) All wages are CPI-adjusted with base year 2015.

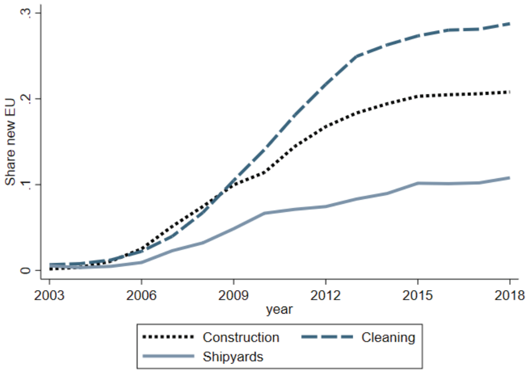

(a) Share immigrant: Only new EU

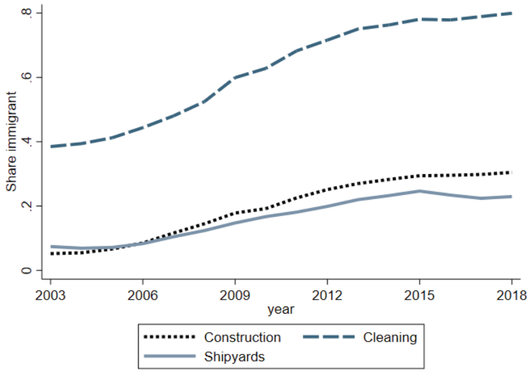

(b) Share immigrant: All

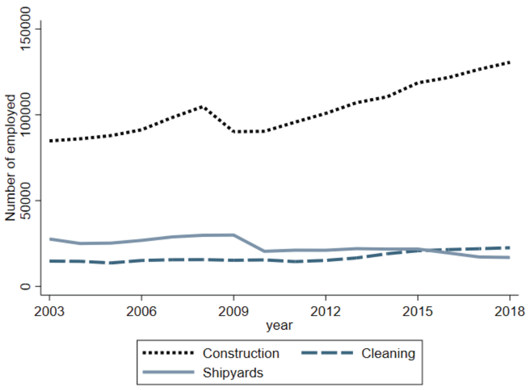

(c) Number of employed

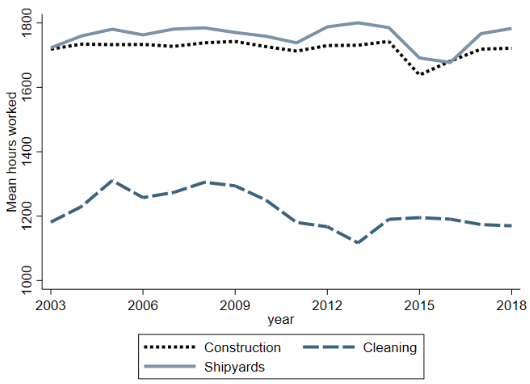

(d) Mean hours worked

Figure 2: Employment, share immigrants and hours worked in the minimum wage industries

3 Empirical Strategy

To study the effect of a wage floor, I start out by checking whether the wage floor introduction coincided with a decrease in the share of workers earning less than the minimum wage. Figure 3 shows that the proportion of workers earning below the minimum wage seems to *increase* slightly in the construction industry, *decrease* in the cleaning industry, and not change in the shipyard industry.

In the minimum wage literature, comparing workers just below the minimum wage to those earning just above has become increasingly common (Cengiz et al., 2019; Dustmann et al., 2021). A visual inspection of within-industry distribution (Figure 4) and the share earning below the minimum wage (Figure 3) suggests the following: Wages in construction and cleaning seem to become *more* dispersed after the bargained wage is extended to cover the whole industry.[16] While these exercises are not very promising in terms of finding an effect of the wage floor, an underlying problem interpreting these figures is the unreliable measure of hourly wages. The remaining part of the analysis will therefore instead focus on annual wages and the comparison of wage floor industries to industries where wage floors were not introduced .

3.1 Estimation

The goal of the analysis is to investigate the development of wages for workers working in the wage-floor industries before the policy introduction. Restricting the analysis to these individuals avoids any concerns regarding whether the composition of workers change as a result of the minimum wage.

[16]Cleaning and construction have smaller left tales after the minimum wage introduction, with skewness from - 0.72 to -0.68 in construction, and from -0.3 to -0.18 in cleaning. A less negative skewness means that less of the density is concentrated right of the mean. In other words the density increases below the mean. Another take on this, which confirms this observation, is shown in Appendix Figure A7 where the distribution after the wage floor introduction is compared to how the distribution would have looked if all wages had grown at the same rate. In the case of cleaning and construction, the distributions are more dispersed after the minimum wage introduction than the counterfactual distribution with fixed growth.

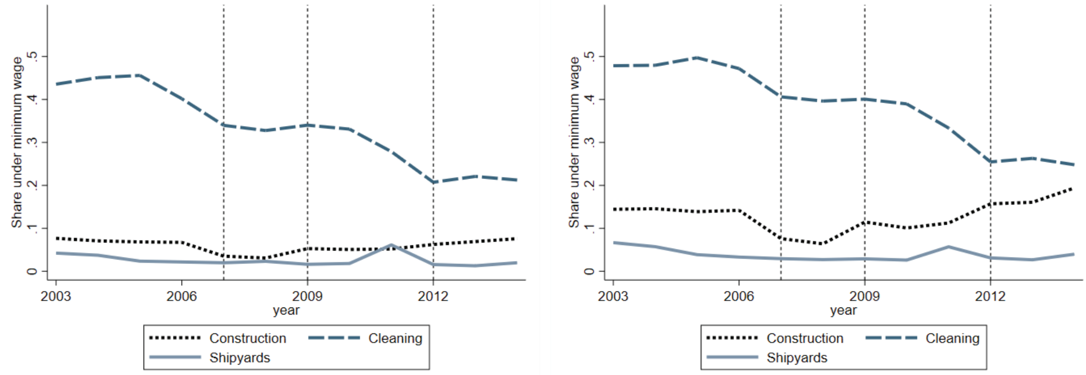

(a) All (b) Immigrant workers

Figure 3: Graph that shows the share of workers with hourly wages measured as less than the minimum hourly wage by industry and separate for immigrants.

Notes: Graph that shows the share of workers with hourly wages measured as less than the minimum hourly wage by industry, for all workers and separately for immigrants. For the years before the minimum wage was introduced, the first price adjusted minimum wage is used. The share of workers earning less than the minimum wage is calculated excluding those whose hourly wages are measured with error (as described in Section 2.1.3).

First, I run a baseline individual fixed-effects model to test whether those working in cleaning, construction or shipyards the year before the wage floor introduction did experience a different annual wage growth in the years after the introduction than before. To do this, I estimate the following time-varying difference-in-difference specification:

$$Outcome_{it} = \alpha_i + \theta_t + \beta Treated_{it} + \epsilon_{i,s} \tag{1}$$

Where $Treated_{it}$ is a dummy variable equal to 1 for individuals who worked in the targeted industries the year before the minimum wage introduction. $Treated_{it}$ is time-varying implying that β measures the average difference in wage for individual i working in a minimum wage industry before and after the minimum wage introduction, compared to workers who are employed in industries not affected by the policy. Including individual fixed-effects, α_i, accounts for time-invariant individual characteristics that are correlated with wage growth. On average individuals are observed approximately 8 years in the data. All regressions also include year-fixed effects and standard errors clustered at the individual and at the industry

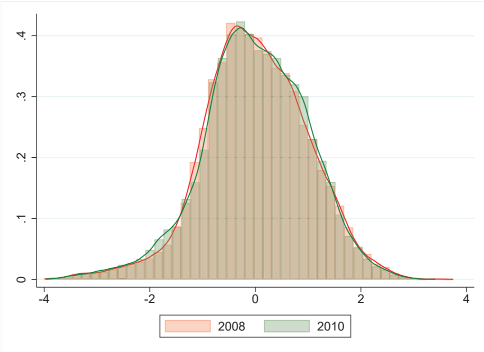

(a) Shipyards

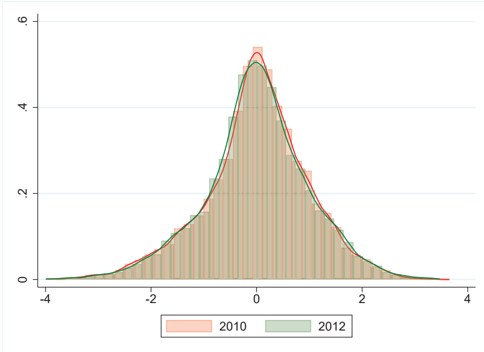

(b) Cleaning

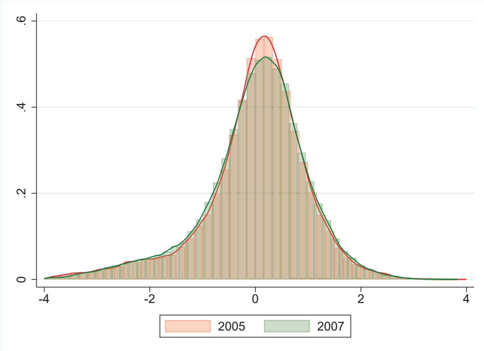

(c) Construction

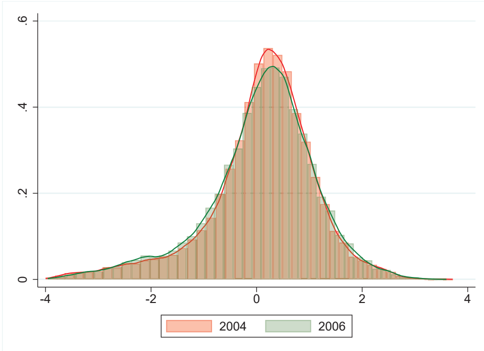

(d) Construction in Oslo and Hordaland

Figure 4: Histogram showing density and estimated kernel density of normalized wages before and after minimum wage introduction in the whole country and in counties Oslo/Hordaland as the minimum wage was introduced there before.

Note: All graphs show hourly wages (excluded those measured with error as described in Section 2.1.3) normalized by industry-year, thus the x-axis shows standard deviations from the industry mean in the respective year.

level. This allows for the error term to be correlated among individuals and within industries.

Several recent papers (Callaway and Sant'Anna, 2020; Goodman-Bacon, 2021) have pointed out the issues emerging in interpreting causal parameters when units receive treatment at different times. These papers show that the estimator in staggered difference-in-difference regression designs cannot necessarily be interpreted as the average treatment effect on the treated, but a weighted average of all treatment effects. The weights on the different units depend on the size of the control group and the variation in the treatment status resulting in units receiving treatment in the middle of the panel having higher weights than those treated earlier or later. For this reason, and because there is no reason to believe that there will be a homogeneous effect, all regressions are done separately for cleaning, construction and shipyard workers.[17]

Similarly, to understand the timing of the change in wages I estimate an individual fixed-effects model with year-specific dummies for the treated populations instead of a time-variant $Treated_{it}$. Thus, in Equation 2, β_t measures the difference between individuals working in a minimum wage industry compared to the rest of working individuals in year t.

$$Outcome_{it} = \alpha_i + \theta_t + \beta_t Treated_i + \epsilon_{i,s} \tag{2}$$

3.2 Omitted variables at the individual and industry level

A worry in interpreting the β's as the effect of the mandatory wage floor introduction, is that workers in these industries may be systematically different in terms of for example education level, study field, country background or age compared to other workers. Although including individual-fixed-effects controls for the correlation between the outcome and background variables that are constant over time, it does not pick up whether different groups experience specific trends over time causing omitted variable bias.

The data set is rich in background variables on characteristics such as for example experience, education, educational field and country of birth, but we do not know which one of these are omitted variables. With many potential control variables, the "right" set of controls is not known. This presents a choice between using too few controls, or the wrong ones, and omitted variable bias will be present, or using too many, and the model will suffer from overfitting. Moreover, by allowing one-self to choose control variables there is a risk of cherry-picking them to get a certain result. Using a Lasso-procedure to select control variables has the potential to solve both issues of overfitting and cherry-picking (Angrist and

[17] Another approach that would give similar results is running a regression with a industry specific treatment dummy. While this design would be slightly less clean, it would also not make it possible to later use the post-double-selection Lasso procedure.

Frandsen, 2022).

Lasso, which abbreviates the "least absolute shrinkage and selection operator," is a form of penalized regression that improves out-of-sample prediction by discarding some regressors and shrinking the coefficients on those retained. The post-double-selection Lasso, introduced by Belloni et al. (2014), implies a two-step procedure where Lasso picks the covariates that explain the outcome variable, and in a second step the covariates that explain the treatment variable. Finally, one can run an OLS regression of the outcome on the treatment including the union of the controls selected in the two steps. This means that I estimate the regression:

$$Outcome_{ti} = \alpha_i + \beta Treated_{it} + \theta_t + t \times (\delta \mathbf{X_i}) + t^2 \times (\delta \mathbf{X_i}) + \epsilon_{i,s} \tag{3}$$

Where X is a vector of potential covariates and t is a linear and t^2 a quadratic time term. The Lasso process generates a list of $t \times X$ and $t^2 \times X$ categories.[18]

A second worry is potential deviations in wage growth between workers in the targeted industries and the rest of the economy, caused by idiosyncratic shocks or shocks affecting the rest of the economy, but less so the targeted industries. As the industry specific minimum wage was introduced to the whole industry (except in construction) at the same time, it is difficult to capture the development in the same industry had the minimum wage not been introduced. While, the control variables selected by Lasso can control for trends specific to observational covariates, it does not control for industry specific trends. Running a regression with industry-specific trends however, may also partly control for the effect.[19] Industry specific effects will be absorbed by the individual fixed-effects for those not changing industry. However, with this reform the rest of the economy was not affected by the minimum wage so there are many potential control industries.

To control for shocks specific to the industries studied, the most likely control industries

[18]Details on covariates included are found in Appendix D. This regression set-up resembles that in Clemens and Wither (2019), who instead use an author-chosen set of covariates.

[19]I also check whether results are robust to adding industry specific time trends in section B1.

are the ones which experienced similar wage trends in the years leading up to the minimum wage introduction.[20] To identify comparable industries I therefore use the synthetic control method (Abadie et al., 2010, 2015; Abadie and Gardeazabal, 2003).

Synthetic controls have become widely applied in empirical research in economics, including in minimum wage research (Allegretto et al., 2017; Jardim et al., 2017; Neumark and Wascher, 2017; Powell, 2021; Reich et al., 2017), and has been described as "arguably the most important innovation in the policy evaluation literature in the last 15 years"(Athey and Imbens, 2017). The synthetic control method allows me to calculate weights for a (sparse) small number of industries that have similar development in median wages in the years prior to the wage floor introduction.

To explore potential mechanisms for diverging wage developments in the wage floor industry and the synthetic control industry, I investigate trajectories of individuals working in the control unit and the treated unit the year before the wage floor introductions. The synthetic control individuals are weighted with the synthetic control weight, w_s, multiplied the inverse of $industry_i$'s share of workers in the synthetic control unit, N_{tot}/N_s, i.e. total number of workers in the synthetic control divided by number of workers in $industry_i$. This is to make sure that average wages are similar at the individual level to the weighted synthetic control average at the industry level.

Comparing wages between individuals in the affected industries to the weighted individuals in the synthetic control industry, I try to identify which type of workers are driving the effect. Doing this, I encounter one major methodological issue. When drawing a non-random sample of people working in a few industries, it is no longer obvious at which level to cluster as the number of industries is too small for conventional two-way clustering. With too few

[20] Appendix Table B2 shows the results when using industries where minimum wages will be introduced in the future as controls. Although the results are not very different from using a synthetic control method, it is not clear whether future minimum wage industries are necessarily the most suitable control industries. Although they all share the fact that immigrant workers earn less than domestic workers, they do not necessarily respond similarly to economic shocks or share the same underlying trend in wages.

groups Wald tests generally over-rejects. Moreover, with the number of observations varying across clusters the number of effective clusters may be reduced (Carter et al., 2017; Imbens and Kolesar, 2016).

I approach this issue in two ways: For the main analysis I cluster at the level of $industry_i \times year_t$, where $industry_i$ refers to the industry in which individual i works in the year before the policy intervention. The rationale behind is that the treatment and the control group are chosen based on the industry they worked in the year before the policy intervention and clustering at this level will take into account clustering in the assignment of the policy intervention as suggested by Abadie et al. (2017). With the synthetic control industry based on 3-4 industries, the number of clusters is then approximately $4 \times 12 = 48$ clusters which is close to the minimum suggested number of clusters (Angrist and Pischke, 2008).

A second approach is to use the wild bootstrap (Cameron and Miller, 2015). However, in the case of only one or few treated clusters, inference with the wild bootstrap is known to fail (Roodman et al., 2019). MacKinnon and Webb (2018) propose the sub-cluster bootstrap, meaning bootstrapping at a finer level than the level of clustering that would be ideal for the research design. They use the example of bootstrapping at the state-year level, when standard errors are clustered at the state level. In a similar approach, I run a regression clustering at the industry level and bootstrap at the industry× year level. MacKinnon and Webb (2018) suggest to compare p-values for the unrestricted and restricted bootstrap and recommend not trusting the results when these two are very different. In the case of comparing the workers in the synthetic control industry to the minimum wage industry, p-values are different for the restricted and unrestricted bootstrap, so I do not use this for the main analysis in the paper. However, although bootstrapped confidence intervals are asymmetric in the expected direction, they all include zero. Therefore, any results comparing synthetic control to treated workers at the individual level are subject to uncertainty and should only be interpreted as suggestive mechanisms. The results from bootstrap are discussed in more depth in Appendix C.2.

4 Results

4.1 Baseline Results

Table 1 shows the results regressing Equation (1) on all working individuals in Norway aged 25-62. The estimates measure the change in the outcome variable before and after the wage floor introduction controlling for time-invariant variation at the individual level with individual fixed effects, as well as time-variant factors common to all individuals with year fixed effects.

Workers in the shipyard and the construction industries both experienced a decrease in hours worked. Together with increasing growth in hourly earnings, for shipyard workers this resulted in no net-effect on annual earnings. In contrast, construction workers, who had a smaller increase in hourly earnings, the effect on annual earnings is negative. Cleaning workers experienced a declining wage growth both in terms of hourly and annual wages. However, the coefficient on annual wage is only significant at the 10% level. A coefficient of -0.069 implies that the annual wage is on average 7% lower after the minimum wage was introduced than before. The probability of being a union member increases for those in the cleaning industry but decreases by 2% for construction and shipyard workers. Appendix Table B1 shows that adding firm or industry fixed effects does not lead to substantially different results for annual earnings, indicating that workers changing firm or industry experienced similar developments compared to workers who did not.[21]

Figure 5 shows the year-specific β coefficients from Equation (2). The coefficients measure the change in wage growth relative to the year before the minimum wage was introduced. In

[21] Adding industry specific trends do change some results: The coefficient on yearly wage is negative and significant in the shipyard and the construction industry, while that for cleaning is no longer significant. Although, this result questions whether the negative wage growth in cleaning reflects a negative trend over time or a change occurring around the year of the minimum wage introduction, adding linear trends can also potentially control away the effect.

Table 1: Fixed-effects estimation (whole population).

	Log yearly wage	Log hourly wage	Union member	Full-time	Hours worked
Cleaning	-0.067*	-0.049*	0.050***	0.0090	24.0**
	(0.036)	(0.026)	(0.0076)	(0.0070)	(9.97)
Construction	-0.068***	0.015**	-0.031***	-0.033***	-68.0***
	(0.011)	(0.0075)	(0.0064)	(0.0061)	(11.0)
Shipyard	-0.016	0.046**	-0.020***	-0.019***	-70.5***
	(0.038)	(0.023)	(0.0073)	(0.0063)	(20.5)
No. of observations	29819791	29579531	29819791	29819791	29605248
No. of individuals	3014159	2988139	3014159	3014159	2989626

Notes: All coefficients come from a separate regression of a time-varying, industry-specific treatment variable on the outcome variable with individual fixed effects and year dummies. The treatment variable is equal to 1 for those who work in the targeted industry one year before the the introduction of the minimum wage. Robust SEs clustered at the individual and industry level in parentheses. * p<0.1, ** p<0.05, *** p<0.01.

cleaning and construction, coefficients are smaller in the years after than in the years before the policy intervention, although this is more prominent for annual than for hourly wages. This indicates that there was a change in wage growth after the wage floor introduction and that the results presented in Table 1 are not solely capturing a declining wage growth over time. For the shipyard industry coefficients become larger in size over time, showing that workers had more wage growth after the policy introduction.

Although individual fixed effects control for the correlation between the outcome and any background variable that is constant over time, it does not capture whether different groups experienced specific trends. The post-double-selection Lasso uses the Lasso estimator to select control variables.[22] Table 2 shows the estimated coefficients from Regression 3, that is the baseline regression with control variables chosen using Lasso. Running a memory demanding analysis like Lasso on very comprehensive data set, forces a reduction in the number of observations due to memory constraints. The Lasso is therefore estimated on a random sample of the population where those treated are over-sampled.[23]

[22]The background variables included in the Lasso regressions are listed in Appendix D.

[23]More specifically, the sample consists of 1/30 of the total number of individuals in the full data set plus half of individuals working in the targeted industries one year before the minimum wage was introduced.

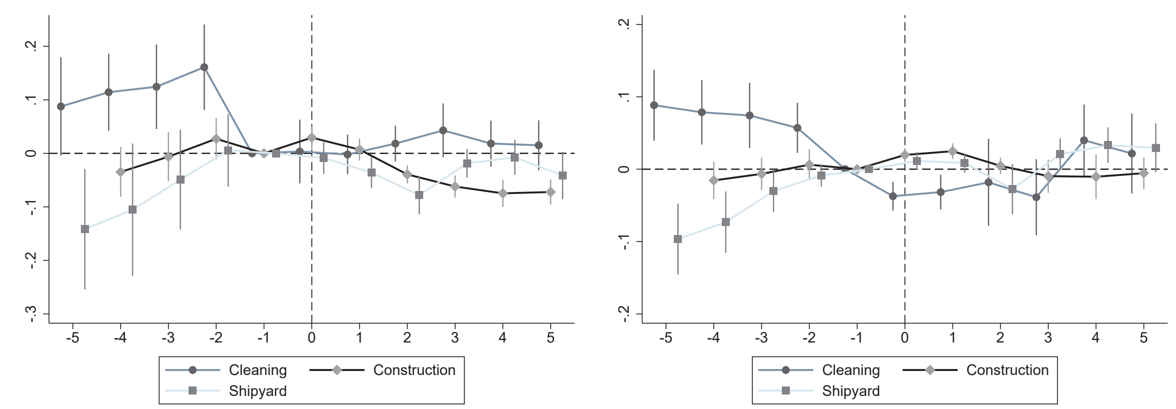

(a) Log yearly wages (b) Log hourly wage

Figure 5: Plot of year-specific treatment dummies (β_t-coefficients) for those working in a minimum wage industry one year before the policy was implemented

Note: All graphs shows treated*year dummies from regression 2 with base year the year before implementation. The x-axis is expressed with reference to the year that the minimum wage is introduced.

In Appendix D, I check how the estimated coefficient on log yearly wage changes with the choice of tuning parameter *lambda* (λ). A higher λ implies a higher penalization for including explanatory variables and therefore results in fewer controls and vice versa a lower λ results in more selected controls. The default choice of λ for the post-double-selection procedure is the theory-driven ('rigorous') λ. The theory-driven λ has been shown to provide the basis of causal inference in the presence of many control variables (Belloni et al., 2012, 2014; Chernozhukov et al., 2015). Figure D1 shows that the coefficient on annual wages in the construction industry drops in size, but remains significant even with a large increase in control variables.[24] In the cleaning industry, the coefficient on annual earnings is only significant at the 10% level in Table 2, but the effect increases in size and becomes more significant the more controls are included. Adding more controls does not alter the coefficient

In finding Lasso controls only the random sample and those working in the targeted industry are included. The reason for using a small sample is due to capacity constraint running a Lasso with a large number of potential covariates. In the next section, including only those that work in the synthetic control industries, all workers in these industries are included.

[24]In practice, I change the λ value by varying the parameter γ. A larger γ results in a lower λ which again increases the number of control variables.

Table 2: Fixed-effects estimation with controls (random sample of the population)

	Yearly wage		Hourly wage	
	Baseline sample	With controls	Baseline sample	With controls
Cleaning	-0.068*	-0.060*	-0.066*	-0.067**
	(0.037)	(0.033)	(0.035)	(0.033)
Adjusted Rsq	0.5649	0.5737	0.4412	0.4434
No. of individuals	118879	116456	118879	116456
Construction	-0.064***	-0.052***	0.014	0.0066
	(0.010)	(0.0090)	(0.013)	(0.012)
Adjusted Rsq	0.5706	0.5772	0.4634	0.4656
No. of individuals	125682	123261	125682	123261
Shipyard	-0.023	-0.020	0.047**	0.047**
	(0.038)	(0.031)	(0.023)	(0.021)
Adjusted Rsq	0.5818	0.5910	0.4796	0.4814
No. of individuals	118476	116058	118476	116058

Notes: All coefficients come from a separate regression for each industry of a time-varying, industry-specific treatment variable on the outcome variable with individual fixed effects and year dummies for a random sample of the population where workers in the affected industries are over-sampled. The treatment variable is equal to 1 for those who work in the targeted industry one year before the the introduction of the minimum wage. Controls are chosen using post-double selection Lasso. More details on this method and the variables used is found in Appendix D. Robust SEs clustered at the individual and industry level in parentheses. * $p<0.1$, ** $p<0.05$, *** $p<0.01$.

in the shipyard industry.

4.2 Synthetic Control Industries

4.2.1 Results

The double-selection Lasso does not necessarily control for trends specific to certain segments in the labor market. For example, did the service sector experience less wage growth than other industries? To find comparable industries, I use the synthetic control method to find weights for industries experiencing similar development in median wages prior to the minimum wage introduction.

The method searches for weights so that the resulting synthetic control best resembles the treated unit in terms of the outcome variable in the year prior to the intervention. Weights are found using a minimum distance approach, meaning the weights $\mathbf{W}^*$ that minimizes

$||\mathbf{X_1} - \mathbf{X_0W}||$, combined with the restriction that the resulting weights are non-negative and sum to one (Abadie et al., 2010). In other words, the synthetic control weights are designed so that the average outcome in the pre-period is similar for affected industries to the weighted average for control units.

First, I calculate predictors as well as the outcome variable at the year×industry level.[25] Including the other covariates rarely matter when the lagged outcome variable is included (Athey and Imbens, 2006). Next, I implement the synthetic control method to find a synthetic control unit for each of the affected industries; construction, cleaning, and shipyards. The synthetic control unit for construction is a combination of the production of metal goods, car repair, and labor services, while the shipyard industry is compared to a weighted combination of oil extraction and the production of beverages, metal and metal goods. For cleaning, the synthetic control consists of agriculture, restaurants and real estate services, with restaurants having the highest weight. Appendix Table C1 shows the list of industries in the synthetic control units as well as their weights.

Figure 6 shows that the annual wages in the construction and the cleaning industries have increased less than in the synthetic control units. In the shipyard sector, it seems as if wages have been growing slightly more than in the synthetic control unit. Thus, comparing the targeted industries to industries with similar trends in wages prior to the minimum wage introduction, does not change the main impression from the fixed-effect estimation: Wages in the cleaning and the construction industries seem to have grown less, at least not more, compared to other industries. Appendix Figure C2 shows that median hourly wages and log employment developed in similar ways in the targeted industries compared to a synthetic control unit, when using the synthetic control method with these variables as outcomes.[26]

[25] Predictors include: Percent females, percent immigrants, log of employees, and share of full-time workers. Moreover, industries with less than 3000 employees per year (on average) are dropped from the pool of industries.

[26] Appendix Table C3 shows which industries constitute the synthetic control unit for median hourly wage and log employment.

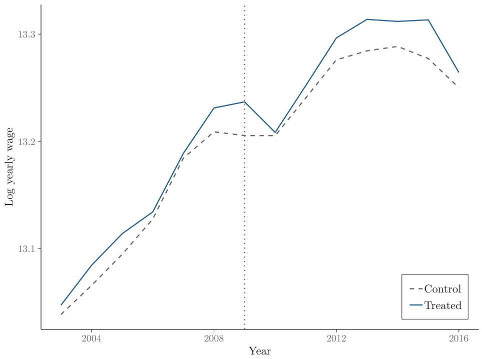

(a) Shipyards

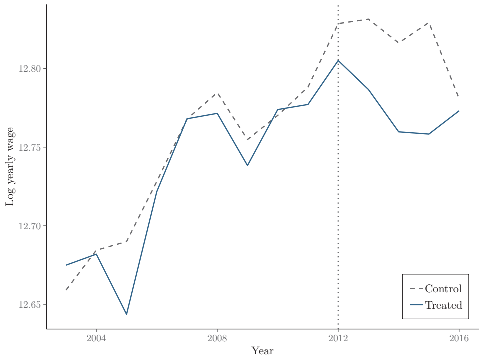

(b) Cleaning

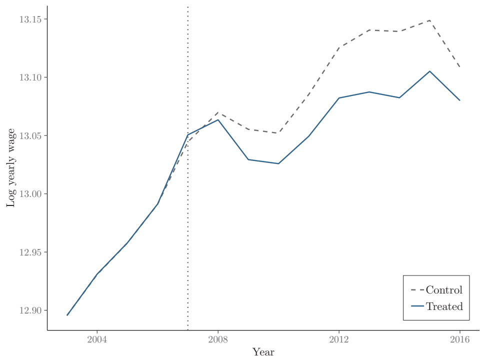

(c) Construction

Figure 6: Plot of median yearly wage in the industry compared to its synthetic control industry

Note: All graphs show the development in the logarithm of the median yearly wages in the industries subject to a minimum wage compared to a synthetic control industry. The synthetic control industry consists of a weighted average of industries chosen using the synthetic control method to approximate the industries affected in the pre-period. The median yearly wage is calculated using full-time workers employed 360 days.

4.2.2 Robustness checks

The synthetic control method has been criticized for being a "black box" (Manning, 2021). First, the choice of pre-periods may arbitrarily affect weights and the chosen control industries. Appendix C.1 reports the results using different pre-periods. While, the choice of pre-period seems to matter for the weights and which industries that are in the synthetic control, the evolution in log wages in the synthetic control is qualitatively the same regardless of the pre-period for construction and cleaning, while for shipyards there are choices of pre-period that would give a different result. Similarly, as suggested by Abadie et al. (2015), by dropping one-by-one industry included in the synthetic control I show that the evolution of wages is not driven by a particular industry (see Appendix Figure C1).

Abadie et al. (2010) suggest to evaluate the significance of the estimates by asking whether the results could be driven entirely by chance. Therefore, I apply the synthetic control method to industries where the minimum wage was not introduced. Figure 7 shows the difference between all industries and their synthetic control industry. The shipyard industry provides a good example on how to interpret these graphs: As the difference between the shipyard industry and its synthetic control seems to be centered in the middle of all the lines, this may indicate that there is nothing particular about this industry. Construction, on the other hand, follows its synthetic control industry before diverging after 2007. Cleaning is less clear; wages seem to decrease after having increased before the year of the minimum wage introduction. However, we cannot expect all placebo states to evolve exactly as their synthetic control industry, both because they might not be very similar to the synthetic control in the pre-period, but also because other shocks happen to the labor market. Nevertheless, the graphical comparison of the actual difference for the treated and the placebo state may still illuminate how the path of the treated state differs from those of the placebo states.

4.3 Exploring mechanisms at the individual level

Within the context of legally enforced wage floors, several mechanisms for how wages for affected workers may decline have been suggested. In economics, a common worry with minimum wages is that the most vulnerable groups might fall out of the labor market (Clemens and Wither, 2019) and that total employment might fall (Neumark and Shirley, 2021). Moreover, an increase in hourly wages might be followed by a decrease in hours worked (Jardim et al., 2017). Finally, labor unions in Norway have voiced concerns than minimum wages can decrease their bargaining power.

In this part of the analysis, I construct a control group consisting of individuals working in the synthetic control industry from Section 4.2 and attempt to identify what types of worker

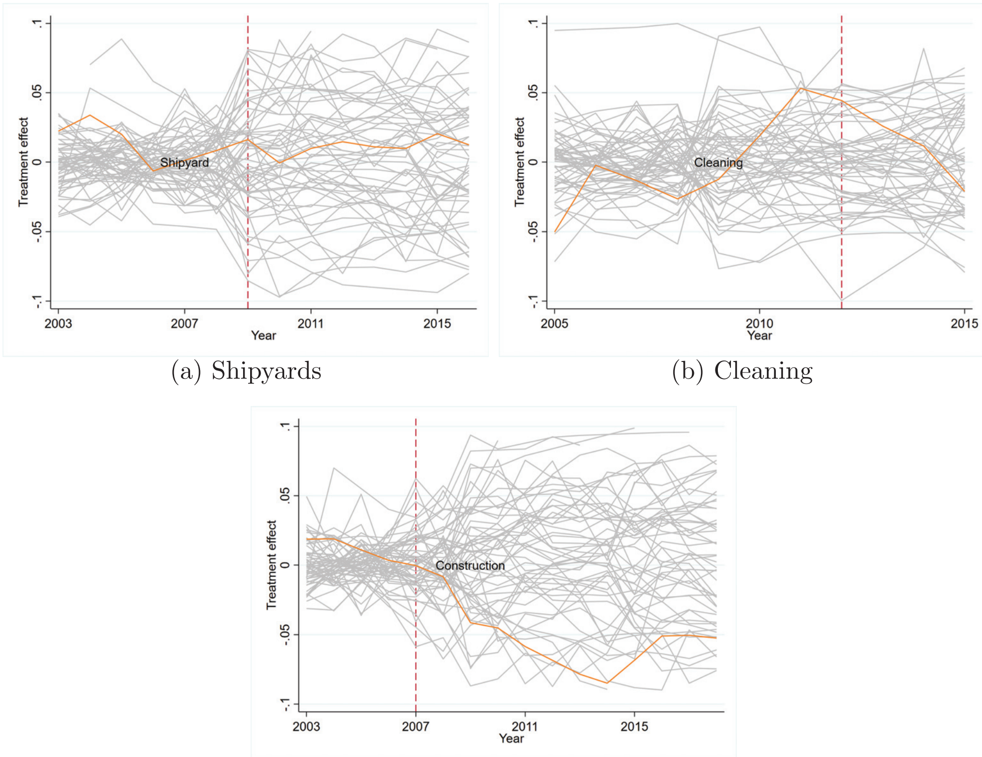

(a) Shipyards (b) Cleaning

(c) Construction

Figure 7: Difference in synthetic and actual log wage for minimum wage industry and 58 Placebo industries

Note: All graphs show the difference in yearly log wage between a industry and its synthetic control industry. The orange line shows the difference between the minimum wage industry and its synthetic control industry

are driving the differences in wages between them. Notice that this is an exploratory analysis trying to explain the difference between the wage floor industries and their synthetic control industry, not necessarily claiming that all differences are due to the wage floor introduction. For the individual level analysis, the treated group consists of individuals who worked in the affected industries the year before the minimum wage introduction, while control group individuals are those working in the control industries in the same year. Figure 8 shows that the evolution of wages for these individuals also follows the same trend in the years before the minimum wage was introduced.

In terms of demographics, Table C4 shows that the affected industries are similar to the synthetic control industries across a range of demographics, but they do not balance perfectly.

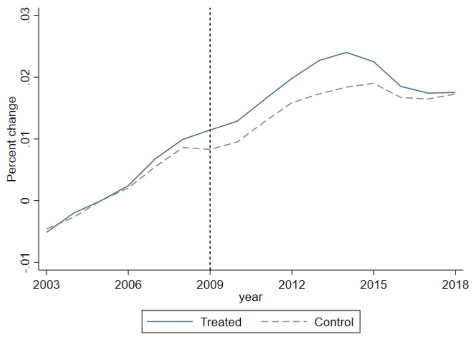

(a) Shipyards

(b) Cleaning

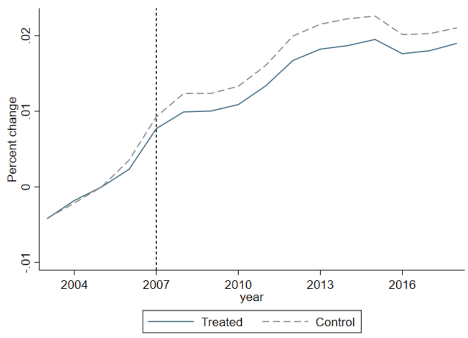

(c) Construction

Figure 8: Percent change in median yearly wage workers in minimum wage industry compared to workers in synthetic control unit

Note: All graphs shows the increase in median yearly wage with 2005 as the base year. The treated group consist of individuals working in the targeted industry the year before the minimum wage introduction, while the control group consists of individuals working in a control industry the same year. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s .

For example, the cleaning and the shipyard industries have higher shares of immigrants than the control industries. To deal with potential differences in wage trajectories experienced by different demographic groups, I again use double-selection Lasso to choose a set of covariates.

Appendix Table C5 shows that a fixed effect model comparing individuals in the treated industries to those in the synthetic control industries, suggest that the effect is no longer significant in the cleaning industry while it is positive in the shipyard industry. Adding control variables, however, changes the results somewhat: Now the effect is again significant for the cleaning industry while the coefficient size drops in absolute value both for workers in the construction and the shipyard industry. When adding industry-specific trends together with controls, the positive coefficient for shipyard workers is no longer significant implying that the change in wages is not timed with the minimum wage introduction or simply that a trend controls away the the effect.

MacKinnon and Webb (2018) recommend to compare p-values for the restricted (imposing null hypothesis) and the unrestricted bootstrap. When the two are very different, this suggest that inference is less reliable. Checking this in practice provides no definite answer, but as confidence intervals include zero, we cannot reject that there was no change in wages at the individual level. Confidence intervals for cleaning and construction are still concentrated to the right (negative effect) while the confidence interval for shipyards is almost symmetric around zero. Thus, the overall pictures looks the same: construction and cleaning workers most likely had negative or no change in wage growth while shipyard workers most likely had positive or no growth. However, since the analysis consist of few clusters and only one treated cluster the bootstrap procedure suggests that findings in this section are subject to uncertainty. This uncertainty comes from the fact that we cannot rule out that the results are caused by other shocks than the wage floor introduction.

4.3.1 Heterogeneity in wage developments

The minimum wage was not introduced to raise wages overall in a industry, but to avoid that immigrants fall behind native workers in terms of wage growth. Figure 9 shows the coefficients from estimating Equation (1), the baseline individual-effects model, separately for those born in Norway ("native"), those born abroad ("immigrants") and for those born in countries that entered the European Union in 2004 ("New EU") including only individuals in the minimum wage- and the synthetic control industry. The results suggest heterogeneity in the average effect. In cleaning, the average effect is negative for natives and zero for workers born abroad. In the construction industry it is negative for both natives and immigrants, while it is positive for all groups in the shipyard industry. Interestingly workers from new

EU member states in cleaning and construction, earning below the industry median, have positive wage developments compared to workers from the same countries in the synthetic control industries.

In the cleaning industry, the negative effect is driven by low-waged workers of all backgrounds. However, for immigrants, whose earnings are in the range between being above 2/3 of the national median and below the median in the cleaning industry, seem to have experienced higher wage growth. This does suggest some effect of the minimum wage since the minimum wage in the cleaning industry was set close to the industry's median wage, as shown in Figure 1c. In the construction industry, there is a similar development suggesting that immigrants earning below the industry median have higher wage growth, while those who are low-waged do not.

Furthermore, I check how the change in wage growth depends on experience, low income status and union membership. Table 3 in the Appendix shows the results from the baseline regression with interaction effects:

$$\begin{aligned} LogYearlyWage_{it} = \alpha_i + \beta Treated_{it} + \theta Treated_{it} \times BackgroundVariable_i \\ + \gamma Year_t + \delta t \times \mathbf{X_i} + \delta t^2 \times \mathbf{X_i} + \epsilon_i \end{aligned} \tag{4}$$

where θ measures the difference in wage growth between the part of the treatment group with the specified background compared to the rest of the treatment group. All background variables are coded so that they take the value as it is in the year before the minimum wage was introduced. This means that variables such as union membership and experience, which may vary over time, only has one value per individual. $\mathbf{X_i}$ refers to a vector of control variables chosen using Lasso. t is a linear time term and t^2 a squared time term. As before all regressions include individual and year fixed effects with standard errors clustered at the individual and year$\times$industry level (the year before minimum wage introduction). Results are shown in Appendix Table 3.

Figure 9: Coefficients from a baseline regression and regression with interaction effects by background (individuals in treated and synthetic control industry)

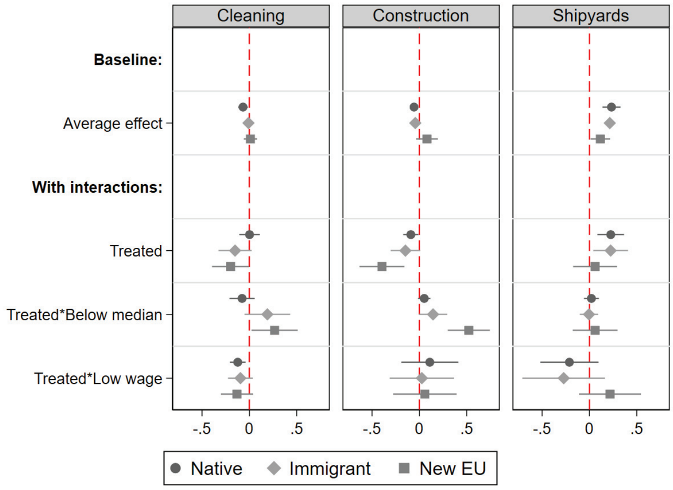

Notes: All plots show the coefficients on treated from (1) a baseline regression done separately by native, immigrant and New EU and (2) from the same regressions with added interaction terms where *low wage* is a dummy equal to 1 if an individual earned less than 60% of the national median wage and below median is a dummy equal to 1 if an individual earned less than the industry *median wage* the year before the minimum wage introduction. All regressions include year and individual fixed effects. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Standard errors are clustered at the individual $industry_i \times year_t$, where $industry_i$ refers to the industry individual i works in the year before the policy intervention.

A minimum wage is a wage floor, thus the effect of minimum wages is expected to fade higher up in the wage distribution (Cengiz et al., 2019). However, the results indicate In the cleaning industry workers who were low-waged, earning less than two-thirds of the median wage in Norway, experience a greater decline in wages than those earning more. In the construction industry, the effect seems to be reverse, but the coefficients are not significant. Moreover, experience with the same employer, working full-time and being a union member are correlated with a greater decline in wage growth in both industries. Also in the shipyard industry, union membership and experience are correlated with lower wage growth. In the shipyard industry, wage growth seems to be concentrated for those earning above the median wage, which suggests that wage growth is caused by other factors than the wage floor, as the shipyard industry had the lowest minimum to median wage ratio.

Table 3: Interaction effects by experience, wage- and union status (individuals in treated and synthetic control industry)

	Construction	Shipyards	Cleaning
(1) Low wage:			
Treated	-0.035	0.17***	-0.011
	(0.023)	(0.050)	(0.024)
Treated*low wage	0.13	-0.17	-0.099**
	(0.17)	(0.17)	(0.047)
(2) Below median:			
Treated	-0.065	0.16**	-0.0033
	(0.039)	(0.068)	(0.063)
Treated*Below median	0.063	0.015	-0.038
	(0.046)	(0.039)	(0.084)
(3) Union member:			
Treated	-0.0088	0.20***	0.018
	(0.019)	(0.042)	(0.024)
Treated*Union member	-0.053***	-0.060**	-0.14***
	(0.0059)	(0.023)	(0.018)
(4) Full-time:			
Treated	0.13***	0.23***	0.087
	(0.038)	(0.065)	(0.053)
Treated*Full-time	-0.16***	-0.060	-0.19**
	(0.038)	(0.077)	(0.075)
(5) No. years with employer:			
Treated	0.072**	0.26***	0.11***
	(0.029)	(0.045)	(0.038)
Treated*Year w/employer	-0.030***	-0.063***	-0.060***
	(0.0035)	(0.020)	(0.015)
Treated*Year w/employer2	0.0011***	0.0023**	0.0022***
	(0.00012)	(0.0011)	(0.00072)
No. of observations	2804920	2278890	642184
No. of individuals	211693	177194	66748

Notes: All columns represent separate regressions on working in a minimum wage industry the year before the minimum wage introduction with an interaction term. Regressions are done separately by the following interaction term: *low wage*, earning less than 60% of the median wage in Norway (1); *below median*, earning less than the median wage in the industry (2); being a union member the year before (3); working full-time the year before the minimum wage introduction (4); and years with current employer measured the year before the minimum wage introduction (5). All regressions include individual fixed effects, year dummies and controls chosen by Lasso. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Standard errors are clustered at the individual $industry_i \times year_t$, where $industry_i$ refers to the industry individual i works in the year before the policy intervention. * p <0.10, ** p<0.05, *** p<0.01.

4.3.2 Movement in the labor market

In the literature increases in the minimum wage are often linked to a decline in employment (see e.g. Neumark et al., 2004; Neumark and Wascher, 1992), although this phenomena has been question in a range of studies.[27] Figure C2 in the Appendix shows that using the synthetic control method, log employment in the cleaning and the construction industry did not decline compared to a synthetic control group. Employment in the shipyard industry has gone down in both absolute numbers as shown in Figure 2c, and compared to the synthetic control industry. However, declining employment in this industry might be linked to the financial crisis affecting the shipyard sector particularly hard.[28][29]

Table 4 shows the coefficients on being "treated" separately for three labor market outcomes. Except for significant coefficients on changing firm; a small increase for workers in the cleaning industry and a decrease for workers in the construction and shipyard industry, these results do not indicate any changes in the probability of exiting the labor market or changing industry.[30]

Although these result suggest no significant changes in labor market attachment for individuals working in the targeted industries before the minimum wage was introduced, looking at the probability of either unemployment or disability as outcome variables suggests a slightly different picture. The results, shown in Appendix Table C7 shows that the probability of receiving disability benefits increased by 1% for affected workers in the cleaning industry compared to workers in the synthetic control group, although the two in Appendix Figure C4 show a similar trend in the years before the minimum wage was introduced. Workers in the construction industry also have a small increase in the probability of

[27]See e.g. Card and Krueger (1994), Cengiz et al. (2019), Dube et al. (2010), and Dustmann et al. (2021)

[28]The value of orders in the Norwegian shipyard industry was for example 90% lower in 2010 than in 2005: https://e24.no/boers-og-finans/i/mRzreO/ordre-ras-paa-90-prosent-for-verftene

[29]For whoever reads this first, I will offer a bottle of port wine.

[30]Figure C3 in the Appendix, showing the treated×year (β_t-coefficients) confirms that the probability of changing firm changes after the wage floor introduction, while the probability of exiting the labor market or changing industry does not.

Table 4: Labor market outcomes (individuals in treated and synthetic control industry)

	Change firm		Change sector		Exit labor market	
	Baseline	w/controls	Baseline	w/controls	Baseline	w/controls
Cleaning	0.03	0.04**	0.03	0.04	0.01	0.01
	(0.02)	(0.02)	(0.03)	(0.03)	(0.02)	(0.02)
Construction	-0.04***	-0.04***	0.02	0.02	-0.010	-0.004
	(0.009)	(0.009)	(0.01)	(0.01)	(0.01)	(0.006)
Shipyards	-0.03***	-0.03***	0.01	0.01	-0.01	-0.005
	(0.01)	(0.01)	(0.03)	(0.03)	(0.01)	(0.006)

Notes: All coefficients come from a separate regression of a time-varying, industry-specific treatment variable on the outcome variable with individual fixed effects and year dummies. The treatment variable is equal to 1 for those who work in the targeted industry one year before the the introduction of the minimum wage. Control variables are chosen using Lasso. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Standard errors are clustered at the individual $industry_i \times year_t$, where $industry_i$ refers to the industry individual i works in the year before the policy intervention. * p <0.10, ** p<0.05, *** p<0.01.

unemployment (.7%) and disability (.3%). Thus the results do not suggest any changes in total employment due to the wage floor introduction, but we cannot rule out that the wage floor affected disability rates.

Next, I turn to checking whether those changing industry, those that are new to the labor market, or change employer, experience a different wage development than those who do not. This is done to understand if there are any links between the wage development in the minimum wage industries and movement in the labor market. Table 5 shows the interaction effect between treatment and movement in the labor market.

Specifically, Table 5 shows the results from the following regression:

$$\begin{aligned} LogWage_{it} = \alpha_i + \beta Treated_{it} + \delta Treated_{it} * Movement_i \\ + \theta_t + \theta t * \mathbf{X_i} + \theta t^2 * \mathbf{X_i} + \epsilon_i \end{aligned} \tag{5}$$

Where $Movement_i$ is a dummy variable equal to one if $individual_i$ changes industry (1), or changes firm (4) 1-2 years after the minimum wage introduction or is new to the labor

Table 5: Labor market movement (individuals in treated and synthetic control industry)

	Construction	Shipyards	Cleaning
(1) Changes sector 1 or 2 y. after min. wage:			
Treated	-0.005	0.07***	-0.002
	(0.02)	(0.02)	(0.04)
Treated*Changes sector	-0.06*	-0.03	-0.07
	(0.03)	(0.02)	(0.1)
(2) Enters labor market 1-2 y before min. wage			
Treated	-0.04**	0.05***	-0.07***
	(0.02)	(0.02)	(0.02)
Treated*New market	0.2***	0.08*	0.2***
	(0.05)	(0.04)	(0.05)
(3) New in sector 1-2 y before min. wage			
Treated	-0.03	0.03	-0.08***
	(0.02)	(0.02)	(0.03)
Treated*New sector	0.1	0.2***	0.2***
	(0.1)	(0.04)	(0.06)
(4) Changes employer 1-2 y after min. wage			
Treated	-0.02	0.06***	-0.05
	(0.02)	(0.02)	(0.03)
Treated*Change employer	-0.02	-0.03	0.05
	(0.02)	(0.03)	(0.06)
No. of observations	2710849	2497829	745995
No. of individuals	211693	192224	66551

Notes: All columns represent separate regressions on working in a minimum wage industry the year before the minimum wage introduction and an interaction term. The interaction term consist dummy variable equal to one for the following outcomes where "before" and "after" are years relative to the minimum wage: Changing industry 1 or 2 years after (1), enters labor market 1 or 2 years before (2), New in current industry 1 or 2 years before (3), changing employer one or two years after (4). All regressions include individual fixed effects, year dummies, and Lasso controls. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Standard errors are clustered at the individual $industry_i \times year_t$, where $industry_i$ refers to the industry individual i works in the year before the policy intervention. * p <0.10, ** p<0.05, *** p<0.01.

market (2) or new to industry (3) 1-2 years before the reform. δ thus measures the difference in log wage between treated individuals experiencing $movement_i$ to treated individuals who do not. It is important to note that all regressions in table 5 are done separate by industry and outcome.

Workers entering the cleaning or construction industry right before the minimum wage introduction seem to experience more wage growth than those who have more experience. This is in line with the results seen in the previous section where more experience was linked to a larger wage growth decline. In the construction industry, those staying in the industry experience less wage decline than those who leave. This suggest that in the construction industry, the decline in wage growth is partly driven by workers leaving the construction industry. In the shipyard industry, staying in the industry is linked to more wage growth compared to leaving the industry. Appendix Table C6 shows the mean of the dependent variable by minimum wage industry and their synthetic control group.

5 Conclusion

In this paper, I study the effect on wages from the legal enforcement of a collectively bargained wage floor. In Norway, wage levels are in general not regulated by law, but a majority of workers are covered by collectively bargained agreements. As immigration increased from new EU member states after 2004, labor unions feared that workers from countries with lower wage levels would outperform domestic workers by accepting lower wages. In response, unions advocated for the legal enforcement of collectively bargained wage floors in certain industries subject to large increases in the share of non-domestic workers.

The results do not suggest that the wage floors in general have raised wages. Instead, workers in two of the targeted industries, cleaning and construction, seem to have had decreased wage growth in comparison with a synthetic control industry. Shipyard workers, on the other hand, may have increased wages, but the results indicate that this is not timed with the wage floor introduction.

However, explorative analyses comparing workers in the synthetic control industry to those in the targeted sectors suggest that workers new to the labor market or the industry did experience wage gains and that immigrant workers may have experienced higher wage growth compared to other immigrant workers. Thus, the wage floor policy might partly have had some effect by avoiding that immigrant workers earn substantially less than domestic workers in these industries.

As the policy was introduced at the same time in the whole industry, except for construction where it was introduced one year earlier in parts of the country, there is a significant amount of uncertainty in interpreting the direct effects of the wage policy. While I attempt to control for the effects of both differential wage trends by covariates and by industry, I cannot rule out the effect of other industry-specific shocks. The effect of extending wage floors to non-covered firms is an interesting venue for future research since the wage floor policy has been extended to new industries and the data on hourly earnings has improved since 2015. In particular, it will be interesting to understand whether wage floors have any effect on the distribution of hourly wages.

References

Abadie, Alberto, Susan Athey, Guido W Imbens, and Jeffrey Wooldridge (2017). *When should you adjust standard errors for clustering?* Tech. rep. National Bureau of Economic Research.

Abadie, Alberto, Alexis Diamond, and Jens Hainmueller (2010). "Synthetic control methods for comparative case studies: Estimating the effect of California's tobacco control program." In: *Journal of the American statistical Association* 105.490, pp. 493–505.

— (2015). "Comparative politics and the synthetic control method." In: *American Journal of Political Science* 59.2, pp. 495–510.

Abadie, Alberto and Javier Gardeazabal (2003). "The economic costs of conflict: A case study of the Basque Country." In: *American economic review* 93.1, pp. 113–132.

Ahlfeldt, Gabriel M, Duncan Roth, and Tobias Seidel (2018). "The regional effects of Germany's national minimum wage." In: *Economics Letters* 172, pp. 127–130.

Ahrens, Achim, Christian B Hansen, and Mark E Schaffer (2020). "lassopack: Model selection and prediction with regularized regression in Stata." In: *The Stata Journal* 20.1, pp. 176–235.

Allegretto, Sylvia, Arindrajit Dube, Michael Reich, and Ben Zipperer (2013). "Credible research designs for minimum wage studies." In: *IZA Discussion Paper No. 7638.*

— (2017). "Credible research designs for minimum wage studies: A response to Neumark, Salas, and Wascher." In: *ILR Review* 70.3, pp. 559–592.

Alsos, Kristin and Line Eldring (2008). "Labour mobility and wage dumping: The case of Norway." In: *European Journal of Industrial Relations* 14.4, pp. 441–459.

Angrist, Joshua and Brigham Frandsen (2022). "Machine Labor." In: *Journal of Labor Economics* 40.S1, pp. 97–140.

Angrist, Joshua D and Jörn-Steffen Pischke (2008). *Mostly harmless econometrics*. Princeton university press.

Athey, Susan and Guido W Imbens (2006). "Identification and inference in nonlinear difference-in-differences models." In: *Econometrica* 74.2, pp. 431–497.

— (2017). "The state of applied econometrics: Causality and policy evaluation." In: *Journal of Economic Perspectives* 31.2, pp. 3–32.

— (2019). "Machine learning methods that economists should know about." In: *Annual Review of Economics* 11, pp. 685–725.

Barth, Erling, Bernt Bratsberg, Torbjørn Hægeland, and Oddbjørn Raaum (2012). "Performance pay, union bargaining and within-firm wage inequality." In: *Oxford Bulletin of Economics and Statistics* 74.3, pp. 327–362.

Barth, Erling, Alex Bryson, and Harald Dale-Olsen (2020). "Union density effects on productivity and wages." In: *The Economic Journal* 130.631, pp. 1898–1936.

Belloni, Alexandre, Daniel Chen, Victor Chernozhukov, and Christian Hansen (2012). "Sparse models and methods for optimal instruments with an application to eminent domain." In: *Econometrica* 80.6, pp. 2369–2429.

Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen (2014). "Inference on treatment effects after selection among high-dimensional controls." In: *The Review of Economic Studies* 81.2, pp. 608–650.

Bertrand, Marianne, Esther Duflo, and Sendhil Mullainathan (2004). "How much should we trust differences-in-differences estimates?" In: *The Quarterly journal of economics* 119.1, pp. 249–275.

Bratsberg, Bern and Mari Holden (2015). "Effekter av allmenngjøring i byggebransjen." In: *Samfunnsøkonomen* 2.

Callaway, Brantly and Pedro HC Sant'Anna (2020). "Difference-in-differences with multiple time periods." In: *Journal of Econometrics*.

Calmfors, Lars and Ragnar Nymoen (1990). “Real wage adjustment and employment policies in the Nordic countries.” In: *Economic Policy* 5.11, pp. 397–448.

Cameron, A Colin and Douglas L Miller (2015). “A practitioner’s guide to cluster-robust inference.” In: *Journal of human resources* 50.2, pp. 317–372.

Card, David and Ana Rute Cardoso (2022). “Wage Flexibility Under Sectoral Bargaining.” In: *Journal of the European Economic Association*.

Card, David and Alan B Krueger (1994). “Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania.” In: *American Economic Review* 84.4, pp. 772–93.

Carter, Andrew V, Kevin T Schnepel, and Douglas G Steigerwald (2017). “Asymptotic behavior of at-test robust to cluster heterogeneity.” In: *Review of Economics and Statistics* 99.4, pp. 698–709.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer (2019). “The effect of minimum wages on low-wage jobs.” In: *The Quarterly Journal of Economics* 134.3, pp. 1405–1454.

Chernozhukov, Victor, Christian Hansen, and Martin Spindler (2015). “Valid post-selection and post-regularization inference: An elementary, general approach.” In: *Annu. Rev. Econ.* 7.1, pp. 649–688.

Chetverikov, Denis, Zhipeng Liao, and Victor Chernozhukov (2021). “On cross-validated lasso in high dimensions.” In: *The Annals of Statistics* 49.3, pp. 1300–1317.

Clemens, Jeffrey and Michael Wither (2019). “The minimum wage and the Great Recession: Evidence of effects on the employment and income trajectories of low-skilled workers.” In: *Journal of Public Economics* 170, pp. 53–67.

Dube, Arindrajit, T William Lester, and Michael Reich (2010). “Minimum wage effects across state borders: Estimates using contiguous counties.” In: *The review of economics and statistics* 92.4, pp. 945–964.

Dube, Arindrajit and Ben Zipperer (2015). "Pooling multiple case studies using synthetic controls: An application to minimum wage policies." In: *Available at SSRN 2589786*.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge (2021). "Reallocation effects of the minimum wage." In: *The Quarterly Journal of Economics* 137.1, pp. 267–328.

Engbom, Niklas and Christian Moser (2021). *Earnings inequality and the minimum wage: Evidence from Brazil*. Tech. rep. National Bureau of Economic Research.

Goodman-Bacon, Andrew (2021). "Difference-in-differences with variation in treatment timing." In: *Journal of Econometrics*.

Holden, Steinar (1998). "Wage drift and the relevance of centralised wage setting." In: *Scandinavian Journal of Economics* 100.4, pp. 711–731.

Imbens, Guido W and Michal Kolesar (2016). "Robust standard errors in small samples: Some practical advice." In: *Review of Economics and Statistics* 98.4, pp. 701–712.

Jardim, Ekaterina, Mark C Long, Robert Plotnick, Emma Van Inwegen, Jacob Vigdor, and Hilary Wething (2017). *Minimum wage increases, wages, and low-wage employment: Evidence from Seattle*. Tech. rep. National Bureau of Economic Research.

— (2018). *Minimum wage increases and individual employment trajectories*. Tech. rep. National Bureau of Economic Research.

MacKinnon, James G and Matthew D Webb (2017). "Wild bootstrap inference for wildly different cluster sizes." In: *Journal of Applied Econometrics* 32.2, pp. 233–254.

— (2018). "The wild bootstrap for few (treated) clusters." In: *The Econometrics Journal* 21.2, pp. 114–135.

Manning, Alan (2021). "The elusive employment effect of the minimum wage." In: *Journal of Economic Perspectives* 35.1, pp. 3–26.

Martins, Pedro S (2021). "30,000 minimum wages: The economic effects of collective bargaining extensions." In: *British Journal of Industrial Relations* 59.2, pp. 335–369.

Mullainathan, Sendhil and Jann Spiess (2017). "Machine learning: an applied econometric approach." In: *Journal of Economic Perspectives* 31.2, pp. 87–106.

Neumark, David, Mark Schweitzer, and William Wascher (2004). "Minimum wage effects throughout the wage distribution." In: *Journal of Human Resources* 39.2, pp. 425–450.

Neumark, David and Peter Shirley (2021). "Myth or measurement: What does the new minimum wage research say about minimum wages and job loss in the United States?" In: *Industrial Relations: A Journal of Economy and Society.*

Neumark, David and William Wascher (1992). "Employment effects of minimum and subminimum wages: panel data on state minimum wage laws." In: *ILR Review* 46.1, pp. 55–81.

— (2017). "Reply to "credible research designs for minimum wage studies"." In: *ILR Review* 70.3, pp. 593–609.

Powell, David (2021). "Synthetic control estimation beyond case studies: Does the minimum wage reduce employment?" In: *Journal of Business & Economic Statistics* 40.3, pp. 1302–1314.

Reich, Michael, Sylvia Allegretto, and Anna Godoey (2017). "Seattle's minimum wage experience 2015-16." In: *Available at SSRN 3043388*.

Roodman, David, Morten Ørregaard Nielsen, James G MacKinnon, and Matthew D Webb (2019). "Fast and wild: Bootstrap inference in Stata using boottest." In: *The Stata Journal* 19.1, pp. 4–60.

Skjerpen, Terje, Tom Kornstad, and Marina Rybalka (2015). *Virkninger allmenngjøring av tariffavtaler*. Tech. rep. Center for Wage Formation.

Statistics Norway (2016). *Nærmere om forholdet mellom gammel og ny statistikk*. URL: `https://www.ssb.no/arbeid-og-lonn/naermere-om-forholdet-mellom-gammel-og-ny-statistikk`.

Appendices

A Data

A.1 Institutional Background

In Norway, a law enabling minimum wages in some sectors had already been adopted in 1993 as Norway entered the European Economic Area. The lack of minimum wage laws in Norway, combined with large wage differences between Norway and some EU member countries, led to concerns that European workers would offer their labor at much lower prices than domestic workers. According to the law, bargained tariff wages could be introduced as minimum wages to protect foreign workers from unreasonable wages and working conditions. This reasoning for a minimum wage law does not exist in other European countries (Alsos and Eldring, 2008).

Until 2006, the law had not been tried out in any sector. The dormant law on collective bargaining extension gained appeal among Norwegian labor unions in 2004, when ten countries, mostly in Eastern Europe, joined the European Union. As EU member states, they were to be part of the internal market entailing free movement of labor, goods and services. Since wage levels in the new member states were lower than in Norway and Western Europe, many feared that increased labor migration would undermine wages and labor rights.

According to the law, a collectively bargained agreement can be extended to cover non-organized firms in particular sectors. For such an extension to take place, the labor union has to document that migrant workers in an industry receive lower wages and experience worse working conditions than domestic workers.[31] These claims are assessed by a tribunal consisting of one representative for unions, one for employers and three other members.[32]

[31] Another party in the tripartite collaboration such as an employer organization can also do this, however in practice it is the labor union that has upheld such claims.

[32] In addition, for each case one member of the party that upholds the claims (for example the worker union in the industry) and one member from their counterpart (such as the employer confederation for that

In general, the labor unions, and the biggest one Landsorganisasjonen (LO), in particular, have supported sector specific minimum wages despite concerns that free-riding might decrease union membership. However, in most cases only minimum wages in the agreements have been legally binding, so that other aspects of the collective agreements are still not mandatory for non-organized employers. Thus, there are still differences between workers covered by collective agreements and those who are not.

On the employer side, the reactions differed by sector. The employers in the construction industry did not support the continuation of minimum wages in 2008, claiming that unjust differences in wages between foreign and domestic workers were not very prevalent.[33] In 2005, however, the employer organization was divided. The part of the organization representing construction companies strongly opposed because they wanted competition from suppliers, while the part of the organization representing painters and wallpaper firms were in support as they feared competition from firms with lower wage levels.[34] In 2018, the employer side again supported the continuation of the minimum wage regime.[35]

The minimum wage in the shipyard industry remains the most controversial. In 2009, eight shipyards, backed by their employer organization, sued the labor union and the state for the minimum wage regime claiming that it was not legal and in conflict with the Norway's EEA obligations. The employers finally won the case in the EFTA Surveillance Authority (ESA).[36] In particular, shipyards disagreed that foreign workers who do not live in Norway should get the same wage as workers who live in Norway.[37]

industry). Other aspects of collective agreements, such as coverage of board and lodging while working are in many cases also a part of the law.

[33]https://frifagbevegelse.no/article-6.158.3744.032ab7ecd0

[34]https://kulturplot.no/sjefene-mener-150000-er-akseptabel

[35]https://www.bnl.no/artikler/2018/tariffnemdas-vedtak-2018/

[36]https://www.nrk.no/norge/staten-og-lo-vant-verftssaken-1.10936323

[37]Shipyards use foreign workers for temporary needs and therefore employ eastern European workers for shorter periods and recruit them directly from their country of residence. Thus, the minimum wage and other working conditions were to cover all who worked in Norway, irrespective of where they lived. In addition the required salary included diet, travel and lodging. The employer side took this to EFTA surveillance Authority (ESA) that monitors compliance with the Agreement on the European Economic Area. ESA

The shipyard industry is among the most organized and less than 10% of workers are working in firms not covered collective bargaining.[38] This pattern is seen in many countries; manufacturing firms are more often covered by collectively bargained agreements than firms in the service sector. Over time, the level of organization on the employer side has has increased from 50% of private industry firms in the 1980's to 70% today.

In the cleaning industry, which has a much lower level of organization on the employer side than the two others, the main employer organization, was very positive to set a binding minimum wage, claiming in their annual report that it will benefit their members.[39] Around 55 to 60 percent of employees in in the cleaning industry worked in a firm that was assumed to be covered by a collective agreement.[40]

The main employer organization NHO has several times suggested to introduce a national minimum wage.[41] However, the current system has political support both from the parties to the left and the biggest right-wing party Høyre and only the smaller right-wing parties, Venstre and FRP, have suggested a national minimum wage.

A.2 Variables

A.2.1 Hourly wages

In order to identify observations whose wages measured with error, I estimate predicted wages given covariates and use the residual to find potential outliers. As the degree of measurement error is not known, I use two very different approaches to predict wages:

First, I predict wages using a classic Mincer equation including only years of schooling and experience as predictors. Normalizing residuals, I identify those whose residuals are more

disagreed with the Supreme Court and concluded that the costs associated with travels workers who come directly from another country to work in Norway, could not be a part of the "allmenngjøring". Thus, the employer did not have to cover their travel, lodging and diet costs. However, such costs are still to be covered for work travels within Norway.

[38] https://www.fafo.no/images/pub/2018/10285.pdf

[39] https://issuu.com/nhoservice/docs/aarsrraport2010-lett

[40] https://evalueringsportalen.no/evaluering/til-renholdets-pris/20209.pdf/@@inline

[41] https://www.aftenposten.no/norge/i/1k9JJ/nho-foreslaar-lovfestet-minsteloenn-i-norge

Table A1: Variable definition

Variable	*Description*
Labor market outcome	Source: Employer-employee relationships
Main employment	Job with the highest salary and next with the highest number of contracted hours.
Yearly Earnings	Total earnings per year with main employer
Hours worked	Total hours worked with main employer
Hourly earnings	Yearly earnings/hours worked. Corrected for possible measurement error (Section A.2.1)
Change industry	Main employment in different industry than year before
Change firm	Main employment at different firm from year before
Exit labor market	Not in the data in a given year
	Source: FD-Trygd (Social Security database)
Unemployed	Registered as unemployed or partly unemployed
Disability	Registered with reduced working capacity
Worker characteristics	Source: National education database, Central Population registry
Age	Age as measured as age January 1st in a given year
Immigrant	Indicator of having a country of birth other than Norway
Education level	First digit NUS-code
Educational field	Second digit NUS-code
Academic high school track	Completion of the academic high school track
Vocational certificate	Certificate of having completed the test for vocational subjects ("fagbrev")
Union member	Having union member fee deduction in individual tax returns
Region	Region/county ("fylke") where worker is employed

than 3 standard deviations from the mean as containing measurement error. If the residuals had been normally distributed, this would have removed .03%, instead 3% of observations are potentially measured with error. With the amount of variables in the dataset, it might seem strange including only years of schooling and experience as predictors. However, which variables to include is not known and including them all lead to over-fitting. For example, including dummies for industry, birth year, educational field and country background would lead to 40% of observations being classified as containing measurement error.

In a second approach, I include a rich set of covariates using cross validated Lasso to predict wages and identify those whose residual are more than 3 standard deviations from the mean as outliers.[42] This procedure is done using STATAs CVLASSO, which is a part of "lassopack" developed by Ahrens et al. (2020). This means that the data is repeatedly

[42]These include industry, birth year and educational level and educational field dummies, industry, immigrant background, country group background, years in Norway, vocational certificate, experience, experience squared, age, county, immigrant background interacted with education level as well as age and experience

partitioned into training and validation data. The model is fit to the training data and the validation data is used to calculate the prediction error. According to this measurement, around 5% of the observations of hourly wages are potentially measured with error. The correlation between the OLS method and Cross-Validation Lasso is around 0.6. The histogram in Figure A1a shows that the two methods yield similar distributions of hourly earnings. However, while OLS identifies measurement error only in the tails of the earnings distribution, Lasso also suggest that there is measurement error closer to the mean as shown in Figure A1b. Figure A1c shows how wage distributions are similar by year, but that the change in data collection occurring in 2015 results in more measurement error as suggested by the predicted values. Although the degree of measurement error seems not be linked to the wage floor implementation, Figure A2 shows that the share with error increases over time in the cleaning and construction industry. Thus, the degree of substantial measurement error seems not to be orthogonal to year nor to industry. The share with error is high in cleaning, average in construction and low in shipyards.

A.2.2 Conversion industry classification SN2002-SN2007:

The standard for industry classification is based on the standard of the EU called NACE. This standard was changed in 2007, and the new standard was taken into use from 2009. The classification has up to 5 digits per industry. The industries where minimum wages were introduced are defined at the most detailed level that was possible, while for other industries only the 2-digit classification is used. The table below shows a conversion between the industry classification used until 2008 and the one used from 2009. For most industries the conversion is straight forward, while for others it might be less good. Industry 33 in the newest classification, "Repair and installation of machinery and equipment", is not accounted for in the old standard.

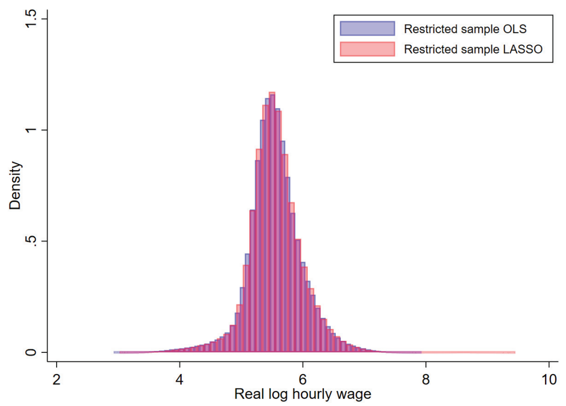

(a) Histogram of log hourly wages: Raw data and the data excluding observations that are very different from their predicted value

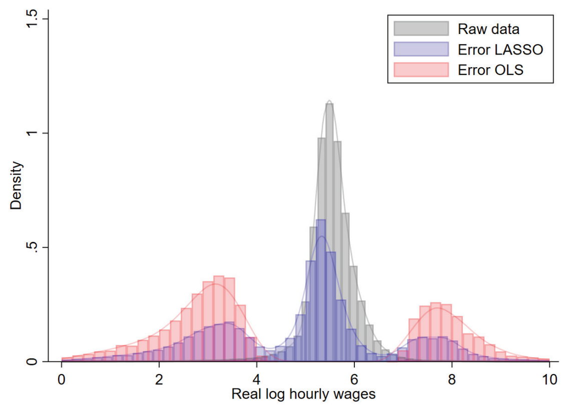

(b) Histogram of log hourly wages: Raw data and observations identified as measurement error using Lasso and ols

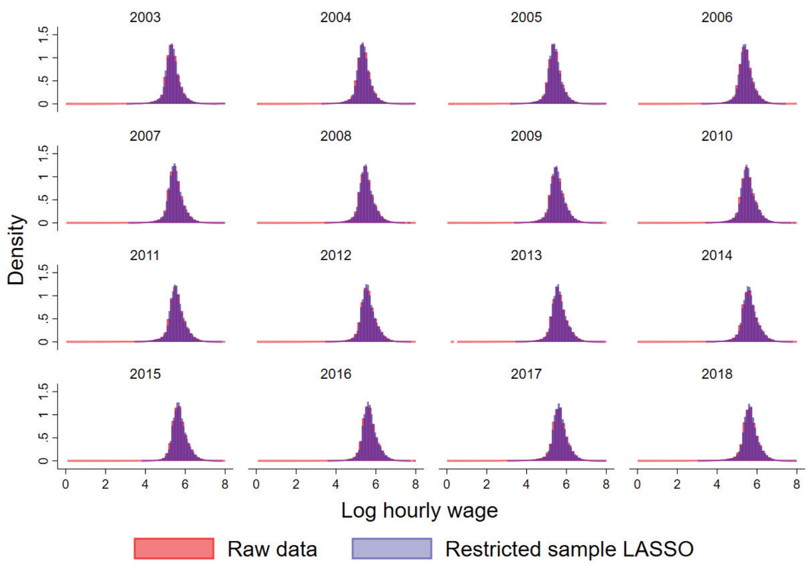

(c) Histogram of log hourly wages showing the raw data and the data excluding observations that are very different from their predicted value by year

Figure A1: Figures showing log hourly wages with and without outliers (identified usingLasso or OLS)

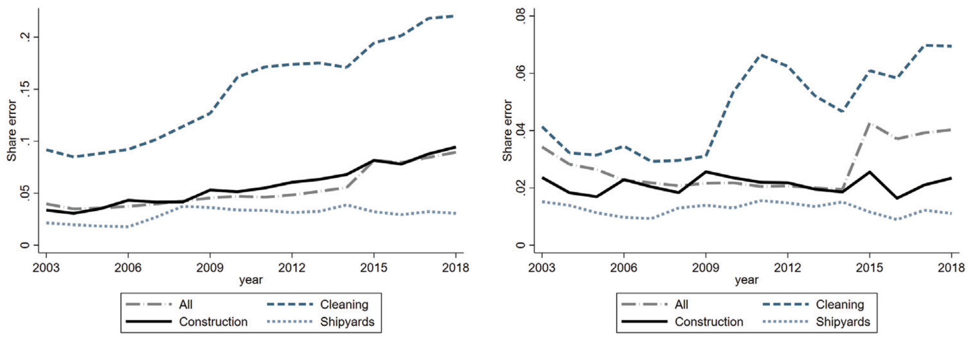

(a) Lasso (b) OLS

Figure A2: Share of observations with measurement error according to a) Lasso b) OLS

Note: These graphs shows the share of individuals whose hourly wages are measured with error according to two different ways of identifying measurement error as described Section A.2.1 a) Lasso and b) OLS. All refers to all working individuals.

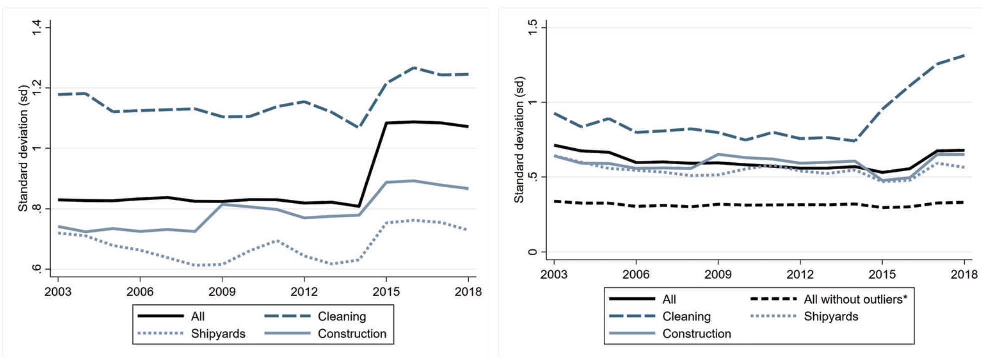

(a) Yearly wages (b) Hourly wages

Figure A3: Development of standard deviation in log total earnings by year.

Note: This graph shows the development in the standard deviation in log total wages by year for all workers and for workers in minimum wage industries. There is a large increase in 2015 due to a change in data collection at the administrative level. (*) Without outliers refers to the exclusion of those whose hourly wages are measured without error as described Section A.2.1 (with CVLASSO).

	Before 2009 (SN 2002)	After 2009 (SN 2007)	Name sector(s)	Before 2009 (SN 2002)	After 2009 (SN 2007)
Agriculture	01	01	Air transport	62	51
Forestry	02	02	Warehouse and support activites for transportation	63 (-633)	52
Fish and fish farming	05	03	Postal and courries services	64	53
Mining of coal, extraction of peat	10	05	Accommodation	551, 552	55
Extraction of oil and gas	11	06	Food and beverage service activities	553, 554, 555	56
Mining of metal ores	13	07	Publishing activities	221	58
Other mining and quarrying	14	08	Motion picture, videao and television production	921	59
Service activities oil and gas	112	09	Programming and broadcast activities	922	60
Manufacture of food products	15 (- 159)	10	Telecommunications	642	61
Manufacture of beverages	159	11	Information service activites	72	63
Manufacture of tobacco products	16	12	Financial service activites, except insurance	65	64
Manufacture of textiles	17	13	Insurance	66	65
Manufacture of wearing apparel	18	14	Activites auxiliary to financial and insurance services	67	66
Tanning and dressing of leather	19	15	Real estate activities	70	68
Manufacture of wood products except furniture	20	16	Legal and accounting activites	7411, 7412	69
Manufacture of pulp and paper products	21	17	Management and consultancy activities	7413, 7414	70
Publishing and reproduction of recorded media	22	18	Architectural and engineering activities	722, 723	71
Manufacture of coke and refined petrolium products	23	19	Scientific research and development	73	72
Manufacture of chemicals and manmade fibres	24 (- 244)	20	Advertising and market research	744	73
Manufacture of pharmaceuticals	244	21	Other professional, scientific and technical activites	748	74
Manufacture of rubber and plastic products	25	22	Veterinary activities	852	75
Manufacture of non-metallic mineral products	26	23	Rental and leasing activities	711	77
Manufacture of basic metals	27	24	Employment activities	74	78
Manufacture of metal products except machinery	28, 296	25	Travel agency and tour operator	633	79
Manufacture of computer and electronic products	297, 32, 30, 33	26	Security and investigation acitivites	746	80
Manufacture of electrical equipment	31, 297	27	Service to buildings and landscape activities	747, 813, 70322	81
Manufacture of machinery and equipment	29 (- 297 and 296)	28	Office administrative and office support	748 (- 74810)	82
Manufacture of motor vehicles	34	29	Public administration and defence	75	84
Manufacture of other transport equipment*	35	30	Education	80	85
Manufacture of furniture	361	31	Human health activities	851	86
Other manufacturing	36 (- 361)	32	Residental care activites	8531	87
Electircity, gas and steam supply	40	35	Social work activities without accomodation	8532	88
Water collection and supply	41	36	Creative, arts and entertainment activites	92	90
Sewerage	9001	37	Libraries, archives and museums	925	91
Waste collection and disposal activities	37, 9002	38	Gambling and betting activities	9271	92
Remediation acitivities	9003	39	Sport, amusement and recreation activities	9272	93
Construction**	45, 42, 43	41	Activities of membership organisations	91	94
Wholesale, trade and repair of motor vehicles	50	45	Repair of computers and household goods	527, 725	95
Wholesale trade, except motor vehicles	51	46	Other personal service activities	93	96
Retail trade, except motor vehicles	52	47	Households as employers of domestic personel	95	97
Land transport and transport via pipelines	60	49	Minimum wage sectors		
Water transport	61	50	Construction	45 (- 4512,-4531)	41, 43 (- 4313, -4321)
*except shipyards			Cleaning	747	8121
**except sectors covered by min. wage			Shipyards	351	301, 3315

Table A2: Industry code conversion

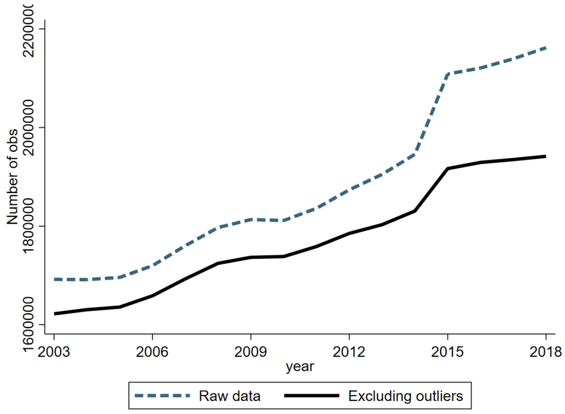

Figure A4: Number of observations per year

Note: The graphs shows the number of observations per year in 1) the raw data (including only those aged 25-61) and 2) Excluding outliers as measured using CVLASSO as described in Appendix Section A.2.1. This means that many individuals appearing in the data in 2015 due to changes in data collection are actually captured by controlling for measurement error in hourly earnings.

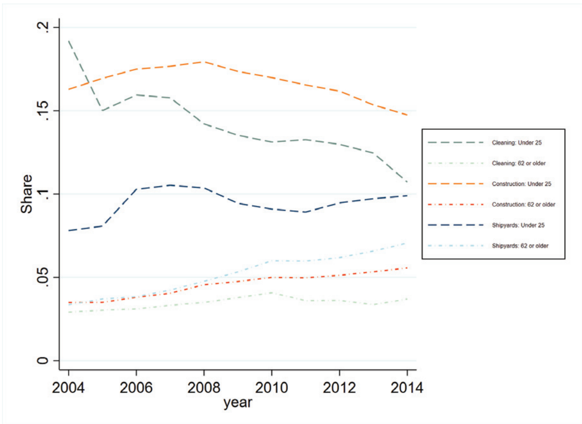

Figure A5: Share of workers younger than 25 and older than 61

Note: This graph stops in 2014 because the change in data collection because the the share of workers above 61 is not comparable before and after the change in data collection that happened in 2015. Before 2015, workers who transitioned to retirement were registered too late so the share of older workers was therefore relatively large (Statistics Norway, 2016).

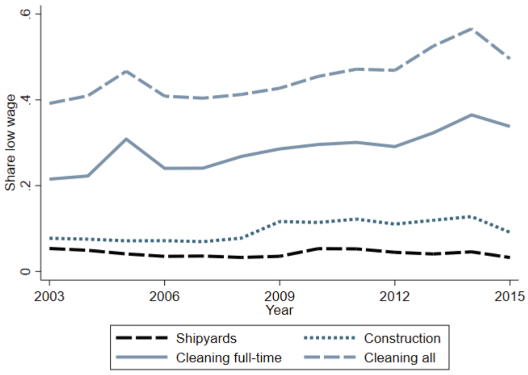

(a) All

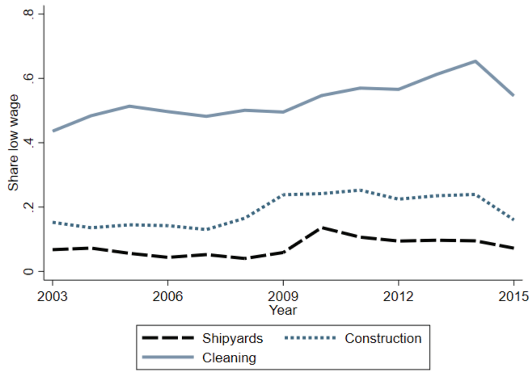

(b) Immigrant workers

Figure A6: Share of low wage workers in the minimum wage industries

Note: Graphs show the share of low wage workers by year separately for 1) All full-time workers (and part-time workers in cleaning) and 2) Workers born outside of Norway. Low wage is defined as having yearly earnings less than 60% of the median wage of a full-time worker. The graphs include only workers that were employed the whole year (more than 360 days)

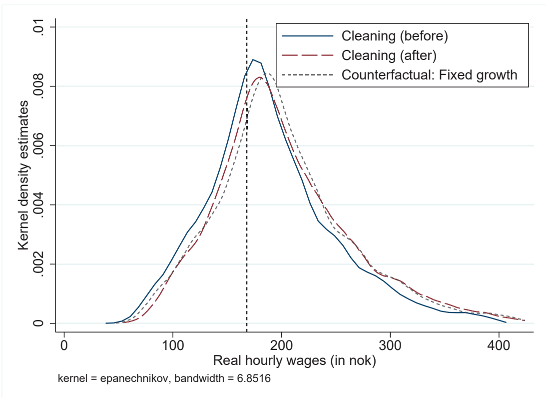

(a) Cleaning

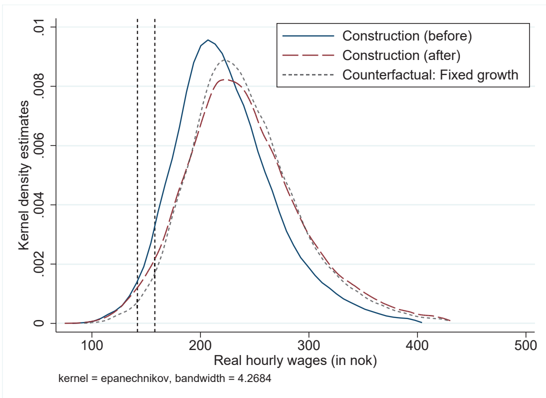

(b) Construction

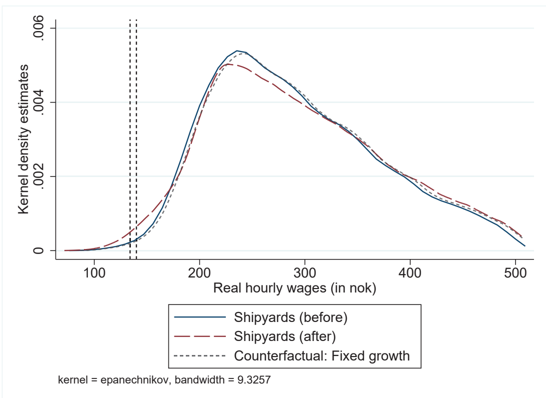

(c) Shipyard

Figure A7: Kdensity plots hourly wages before and after policy introduction

B Results

Using those working in future minimum wage industries as control: The minimum wage was introduced in only specific industries characterized by for example a high share of workers from new EU countries, low education prerequisites and most importantly where the labor union has evidence of substantial wage differences between immigrant workers and native workers. Thus, these industries are different than other industries and may have experienced different wage developments for example due to the influx of labor from other countries or more generally lower wages for workers without higher education degrees. A potential control group, working in industries that share these features, might be those working in industries where the minimum wage will be introduced in the "future". These industries are the fish industry (2015), transport (2015) and the hotel/restaurant industry (2018). Table B2 shows the results from regression 2 including only those working in industries where the minimum wage will be introduced in the future.

Table B1: Alternative specifications baseline regression (the whole working population)

	Log yearly wage			Log hourly wage		
	Firm FE	Sector FE	Sector trend	Firm FE	Sector FE	Sector trend
	(1)	(2)	(3)	(4)	(5)	(6)
Cleaning	-0.12***	-0.091*	0.017	-0.0058	-0.016***	0.024***
	(0.042)	(0.046)	(0.014)	(0.0043)	(0.0034)	(0.0054)
Construction	-0.052***	-0.059***	-0.045***	0.012**	0.014**	0.014***
	(0.014)	(0.018)	(0.016)	(0.0055)	(0.0062)	(0.0051)
Shipyard	-0.035*	-0.029	-0.063**	0.020***	0.029***	0.0015
	(0.021)	(0.031)	(0.029)	(0.0069)	(0.0042)	(0.012)
No. of observations	29777636	29819791	29819791	23658808	23614883	23658808
No. of individuals	3011506	3014159	3014159	2728937	2724192	2728937

Notes: All coefficients come from a separate regression of a time-varying, industry-specific treatment variable on the outcome variable with individual fixed effects and year dummies. The treatment variable is equal to 1 for those who work in the targeted industry one year before the the introduction of the minimum wage. Robust SEs clustered at the individual and industry specific level level in parentheses and * p<0.10, ** p<0.05, *** p<0.01.

Table B2: Average effects (individuals in treated industry compared to individuals in industries where wage floors introduced after 2012)

	Log yearly wage	Log hourly wage	Full-time	Hours worked
Cleaning	-0.0086	-0.051***	0.025***	38.6***
	(0.0070)	(0.0052)	(0.0035)	(4.25)
Construction	-0.039***	0.010***	-0.041***	-57.7***
	(0.0030)	(0.0023)	(0.0016)	(1.91)
Shipyard	0.033***	0.062***	-0.049***	-58.7***
	(0.0035)	(0.0026)	(0.0021)	(2.25)
No. of observations	1422634	1414663	1422634	1416373
No. of individuals	121573.00	121570.00	121573.00	121573.00

Notes: All coefficients come from a separate regression of a time-varying, industry-specific treatment variable on the outcome variable with individual fixed effects and year dummies. The treatment variable is equal to 1 for those who work in the targeted industry one year before the the introduction of the minimum wage. The control group consist of other industries where wage floors have been introduced after 2012. Robust SEs clustered at the individual and industry level in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

C Synthetic control

C.1 Robustness checks and other outcomes

Robusness checks: The synthetic control method may be sensitive to the arbitrary choice of pre-period or the inclusion or exclusion of particular units in the pool of potential control units. Table C2 shows how the weights change for various ranges of pre-periods. Regardless of the pre-period, the restaurant industry is highly weighted as a control industry for the cleaning industry. Other industries seem to fluctuate more. For the construction industry, real-estate services, production of metal goods and oil services form part of the synthetic control industry for all choices of pre-period. Other industries are included for some periods and not for others. Although there seems to be a high level of consistency in the industries chosen, the small fluctuations may affect the result. Figure C1 shows how wages in the synthetic control industry for the different choices of pre-period. The main interpretation of the result does not change for the cleaning and construction industry, while there seem to be choices of pre-period that would change the result for the shipyard industry.

Finally, Figure C1 shows how the synthetic control industry would evolve if industries are dropped one-by-one from the pool of potential control industry. This means that for example for cleaning, first restaurants are dropped and in a second attempt real-estate services are dropped. There is no indication that the exclusion/inclusion of a particular industry would alter the results.

Other outcomes: Table C3 shows the synth weights for alternative outcome variables, log hourly wages and log employment. For hourly wages, naturally, the synthetic control industry is similar to the synthetic control industry for yearly wages. Employment, however, seems to be different in particular for construction. Figure C2 shows that the median hourly wage (corrected for measurement error as described in Appendix A.2.1) does not seem to change compared to the synthetic control industry after the wage floor introduction, while

Table C1: Industries in the synthetic control unit

Industry (1)	Synth Weight (w) (2)	N (3)	Share (α) (4)	Ind. weight: $\frac{w}{\alpha}$ (5)
Cleaning		14,429		
Agriculture	.002	6,247	.12	.017
Restaurants	.828	26,248	.49	1.69
Real estate services	.171	20,153	.39	.44
Total synthetic	1	52,648	1	
Construction		78,953		
Production metal goods	.583	18,072	.14	4.29
Car repair	.183	35,814	.27	.68
Labor services	.235	79,161	.59	.395
Total synthetic	1	133,047	1	
Shipyards		29,682		
Oil extraction	.191	18,750	.11	1.74
Beverage production	.01	4,125	.025	.40
Metal production	.185	10,197	.06	3.08
Production metal goods	.364	19,894	.12	3.03
R&D	.035	12,416	.075	.49
Construction	.214	97,996	.61	.35
Total synthetic	1	163,378	1	

Notes: The two first columns show the industries in the synthetic control unit and their weights as calculate by the synthetic control method. In the explorative analysis people working in the synthetic control unit prior to the minimum wage are used as a control group to the people working in the minimum wage industry. In the third and fourth column, N refers to the number of people working in the industry the year before the minimum wage introduction and α refers to the share of total workers in the synthetic control unit working in the specified industry. To ensure that the mean wage is the same at the individual level as the industry level, all individuals are weighted by the synthetic weight w of the industry they work in multiplied by the inverse of the share α of that industry.

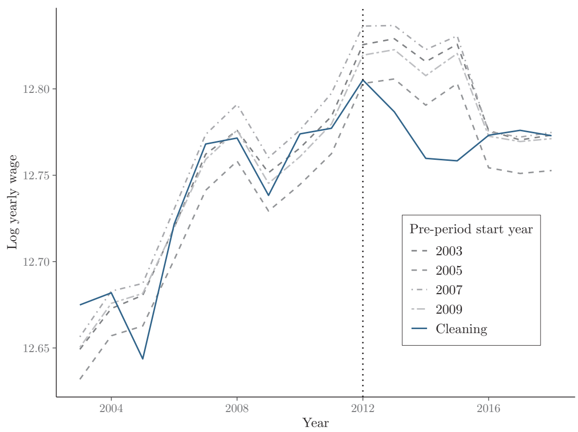

(a) Alternative pre-period: Cleaning

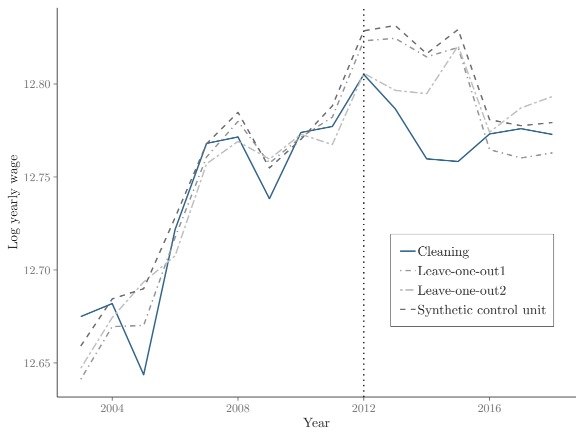

(b) Leave-out one industry: Cleaning

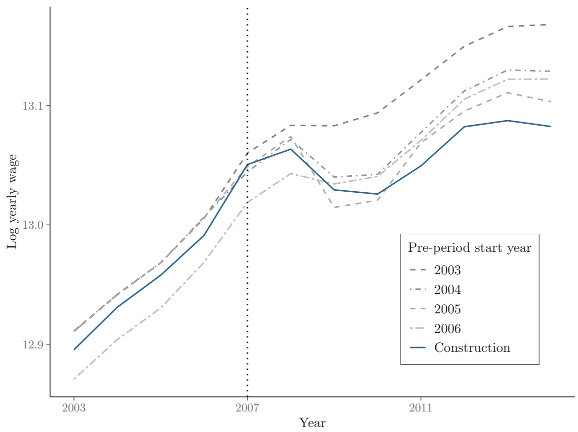

(c) Alternative pre-period: Construction

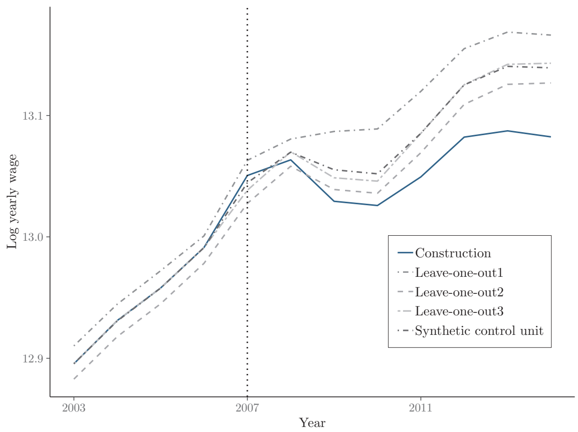

(d) Leave-out one industry: Construction

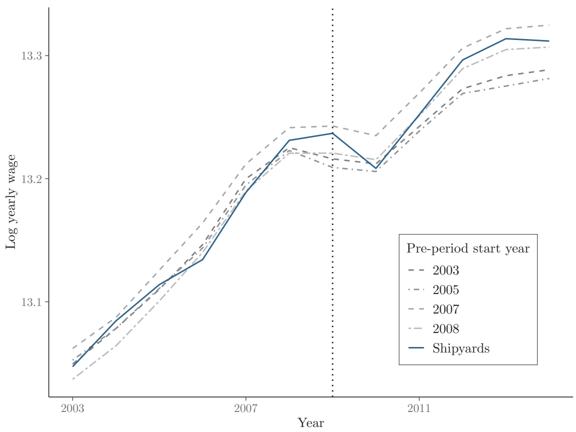

(e) Alternative pre-period: Shipyards

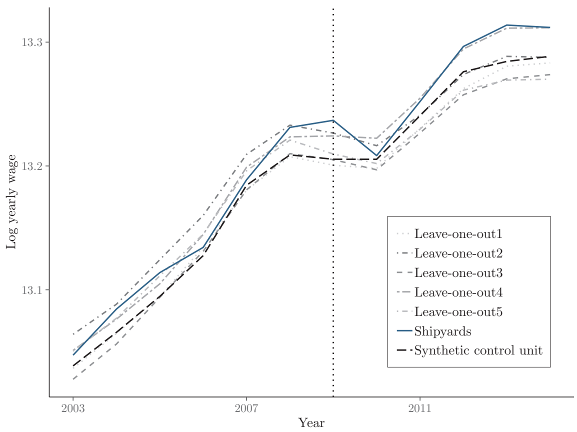

(f) Leave-out one industry: Shipyards

Figure C1: Robustness checks synth

Note: All figures show how the synthetic control industry varies with a) choice of pre-period and b) when dropping industries from the pool of control units one-by one. For a) A set of other possible choices, as indicated on the figures, are used to construct a synthetic control unit. For b) The synthetic control weights are re-calculated when dropping industries one-by-one. The rationale behind this exercise is to see how if results are sensitive one particular industry.

(a) Cleaning

Industry	2003 -2010	2004 -2010	2005 -2010	2006 -2010	2007 -2010	2008 -2010	2009 -2010	2010 -2010
Restaurants	0.81	0.83	0.84	0.82	0.89	0.81	0.81	0.86
Mail distribution	0.12	0.00	0.00	0.05	0.00	0.00	0.00	0.00
Real estate services*	0.07	0.17	0.15	0.13	0.01	0.19	0.18	0.13
Oil extraction	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.01
Transport	0.00	0.00	0.01	0.00	0.00	0.00	0.00	0.00
Aviation	0.00	0.00	0.00	0.00	0.03	0.00	0.00	0.00
Hotels	0.00	0.00	0.00	0.00	0.05	0.00	0.00	0.00
Labor services	0.00	0.00	0.00	0.00	0.03	0.00	0.00	0.00
Care in institution	0.00	0.00	0.00	0.00	0.00	0.00	0.01	0.00

(b) Construction

Industry	2003 -2006	2004 -2006	2005 -2006	2006 -2006
Waste management	0.22	0.00	0.00	0.00
Sewage/water distr.	0.21	0.00	0.00	0.28
Production of metal goods	0.17	0.51	0.47	0.35
Transport	0.13	0.00	0.00	0.11
Real estate services*	0.12	0.37	0.31	0.18
Fishing	0.08	0.00	0.00	0.00
Services oil	0.08	0.12	0.07	0.07
Rental and leasing.	0.00	0.00	0.00	0.00
Agriculture	0.00	0.00	0.04	0.00
Oil extraction	0.00	0.00	0.02	0.00
Labor services	0.00	0.00	0.09	0.00

(c) Shipyards

Industry	2003 -2009	2005 -2000	2006 -2009	2007 -2009
R&D	0.01	0.047	0.053	0.004
Oil extraction	0.02	0.039	0.085	0.098
Services oil	0.15	0.14	0.14	0.191
Production of metal	0.136	0.31	0.24	0.054
Production of metal goods	0.00	0.11	0.321	0.564
Technical consultancy	0.09	0.07	0.03	0.00
Electricity	0.157	0.002	0.00	0.00
Construction	0.435	0.38	0.17	0.00
Beverage production	0.00	0.00	0.052	0.09

Table C2: Table showing how synthetic control weights change for various choices of pre-period

only in the shipyard industry does employment seem to decrease compare to the synthetic control industry. However, this particular result is difficult to interpret as the shipyard industry was hit by the financial crisis immediately after the wage floor introduction and was possibly more hardly hit than any other industry in Norway.

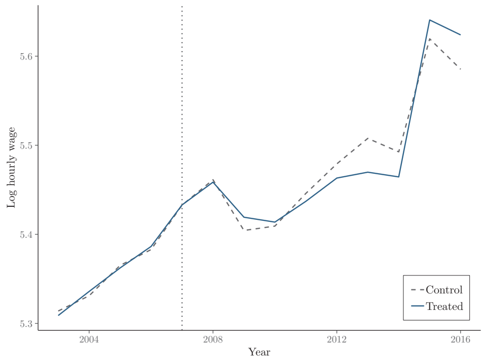

(a) Log hourly wages: Construction

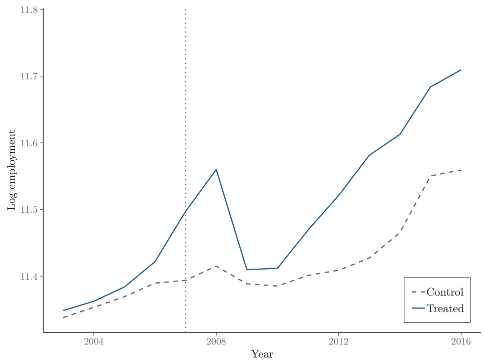

(b) Log employment: Construction

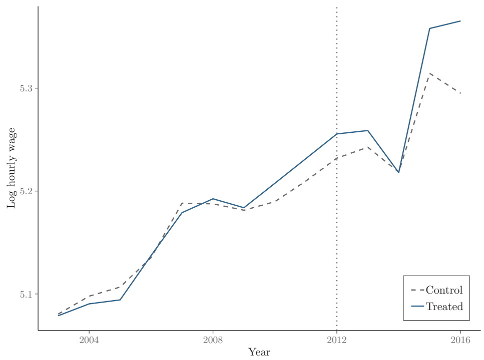

(c) Log hourly wages: Cleaning

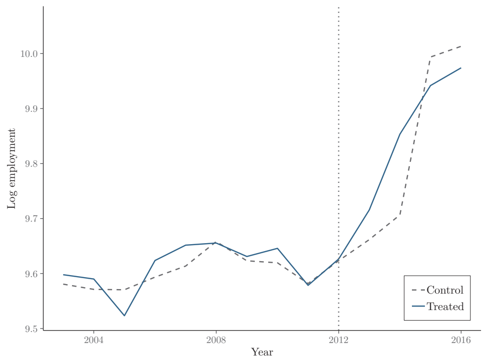

(d) Log employment:: Cleaning

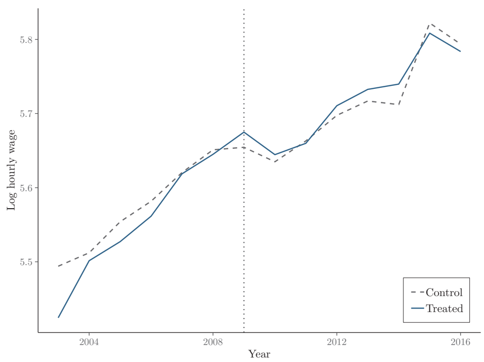

(e) Log hourly wages: Shipyards

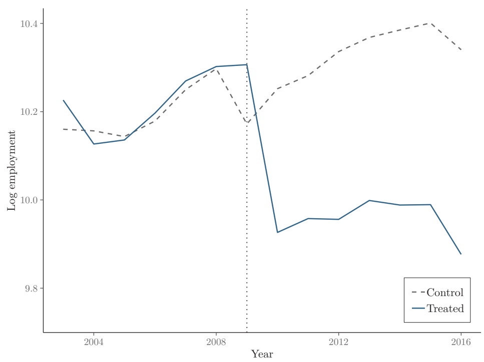

(f) Log employment: Shipyards

Figure C2: Plot of log employment and hourly wages in industry compared to its synthetic control industry

Note: All graphs shows the development in the logarithm of the outcome in the industries subject to a minimum wage compared to a synthetic control industry. The synthetic control industry consists of a weighted average of industries chosen by stata's synth program to approximate the pre-trend in the outcome for the industries affected. This means that the synthetic controls industry in this case contains the industries and weights as shown in Table C3. The log median hourly wage is calculated excluding individuals whose hourly wage is measured with error as described in Section 2.1.3.

Table C3: Table showing synth weights for hourly wages and log employment as the outcome variable

Log hour		Log employment	
industry	Weight	industry	Weight
Construction			
Production Metal Goods	0.792	Fishing and aquaculture	0.025
Labor services	0.18	Oil Extraction	0.011
Shipyard	0.028	Retail	0.383
		Restaurants	0.001
		Technical testing and analysis	0.214
		Education	0.365
Cleaning			
Transport	0.06	Agriculture	0.22
Restaurants	0.94	Restaurants	0.19
		Hotels	0.59
Shipyards			
Production Metal Goods	0.497	Production Metal Goods	0.426
Services Oil	0.319	Services Oil	0.278
R&D	0.083	Public Administration	0.123
Construction	0.08	Labor services	0.088
Manufacture of Basic Metals	0.021	Agency Trading	0.056
		Oil Extraction	0.016
		Transport	0.012

Table C4: Balancing tests treated individuals and individuals in synthetic control group

	Treated		Control		Difference	
Variable	*Mean*	*sd*	*Mean*	*sd*	b	*P-value*
Construction						
Immigrant	0.065	0.247	0.068	0.252	0.00	(0.68)
EU 2004	0.011	0.103	0.006	0.077	0.00	(0.00)
Years of schooling	11.770	1.452	12.702	2.156	-1.05	(0.00)
Full-time	0.950	0.217	0.912	0.284	0.05	(0.00)
Male	0.919	0.273	0.709	0.454	0.24	(0.00)
Observations	78953		133047		212000	
Cleaning						
Immigrant	0.688	0.463	0.435	0.496	0.30	(0.00)
EU 2004	0.184	0.387	0.052	0.221	0.10	(0.00)
Years of schooling	11.208	2.220	11.593	2.054	-0.42	(0.00)
Full-time	0.544	0.498	0.671	0.470	-0.15	(0.00)
Male	0.418	0.493	0.498	0.500	-0.12	(0.00)
Observations	14229		52648		66877	
Shipyards						
Immigrant	0.205	0.404	0.138	0.345	0.07	(0.00)
EU 2004	0.072	0.259	0.062	0.242	0.02	(0.00)
Years of schooling	12.632	2.151	12.012	1.891	0.35	(0.00)
Full-time	0.981	0.138	0.959	0.198	0.02	(0.00)
Male	0.816	0.388	0.883	0.321	-0.05	(0.00)
Observations	29682		163378		193060	

Notes: Table showing mean and standard deviation of demographic variables by treated and control group weighted by their synthetic control weights. The second last column shows the mean difference and the last column the two-sided p-value from a t-test of the difference between the two means.

Table C5: Reduced form coefficients w/controls (individuals in treated and synthetic control industry)

	Baseline	With controls	With trends	Trends+controls
Cleaning	-0.014	-0.031**	-0.012	-0.032**
	(0.020)	(0.014)	(0.020)	(0.014)
Adjusted Within Rsq	0.0000	0.0217	0.0116	0.0332
Adjusted Rsq	0.4411	0.4573	0.4475	0.4573
No. of individuals	66808	66748	66808	66748
Construction	-0.047**	-0.035**	-0.039**	-0.042***
	(0.023)	(0.018)	(0.017)	(0.016)
Adjusted Within Rsq	0.0004	0.0180	0.0181	0.0319
Adjusted Rsq	0.5005	0.5096	0.5094	0.5110
No. of individuals	211699	211693	211699	211693
Shipyard	0.063***	0.049**	0.051***	0.017
	(0.021)	(0.019)	(0.018)	(0.015)
Adjusted Within Rsq	0.0005	0.0271	0.0363	0.0545
Adjusted Rsq	0.5753	0.5872	0.5905	0.5501
No. of individuals	192271	192224	192271	192224

Notes: All columns represent separate regressions on working in a minimum wage industry the year before the minimum wage introduction. All regressions include individual fixed effects, year dummies, synth weights. Controls chosen by Lasso are in column 2. Column 3 includes industry-specific trends and the last column (4) includes both controls and trends. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Standard errors are clustered at the individual $industry_i \times year_t$, where $industry_i$ refers to the industry individual i works in the year before the policy intervention. * p <0.10, ** p<0.05, *** p<0.01.

Table C6: Summary statistics labor market outcomes (individuals in treated and synthetic control industry)

	Treated	Control
	Mean(sd)	*Mean(sd)*
Cleaning		
Changes employer 1-2 years after	0.35	0.34
	(0.48)	(0.47)
Exits labor market 1-2 y after	0.37	0.28
	(0.48)	(0.45)
In same sector 1 year after	0.55	0.64
	(0.50)	(0.48)
In same sector 2 years after	0.50	0.57
	(0.50)	(0.50)
Entered labor market 1-2 y before	0.37	0.30
	(0.48)	(0.46)
Observations	14,429	52,648
Construction		
Changes employer 1-2 years after	0.22	0.22
	(0.41)	(0.41)
Exits labor market 1-2 y after	0.17	0.17
	(0.38)	(0.37)
In same sector 1 year after	0.79	0.72
	(0.41)	(0.45)
In same sector 2 years after	0.58	0.51
	(0.49)	(0.50)
Entered labor market 1-2 y before	0.08	0.07
	(0.28)	(0.26)
Observations	87,317	123,555
Shipyards		
Changes employer 1-2 years after	0.21	0.18
	(0.41)	(0.38)
Exits labor market 1-2 y after	0.15	0.17
	(0.35)	(0.38)
In same sector 1 year after	0.54	0.73
	(0.50)	(0.45)
In same sector 2 years after	0.49	0.67
	(0.50)	(0.47)
Entered labor market 1-2 y before	0.16	0.15
	(0.37)	(0.36)
Observations	29,802	170,336

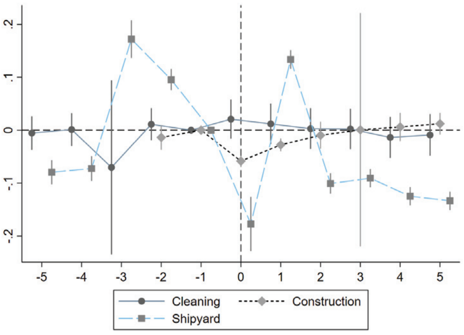

(a) Probability of changing industry

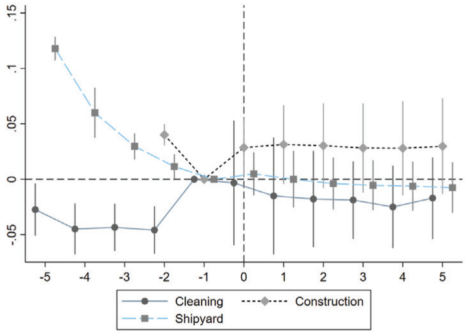

(b) Probability of exiting labor market

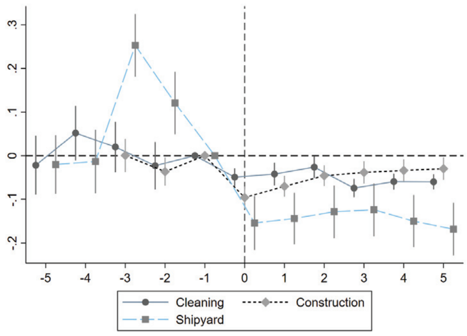

(c) Probability of changing firm

Figure C3: Plot of year-treated dummies from a regression with individual and year fixed effects (individuals in treated and synthetic control industry)

Notes: All graphs show year*treated coefficients. All coefficients come from a regression of the outcome on treated*year dummies and year dummies with individual fixed effects.Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Standard errors are clustered at the individual $industry_i \times year_t$, where $industry_i$ refers to the industry individual i works in the year before the policy intervention.

Table C7: Average effects on unemployment and disability benefits (individuals in treated and synthetic control industry)

	Unemployed		Disability	
	Baseline	W/controls	Baseline	W/controls
Cleaning	-0.002	0.003	0.01***	0.01**
	(0.003)	(0.003)	(0.004)	(0.004)
Construction	0.008**	0.007**	0.002	0.003**
	(0.003)	(0.003)	(0.001)	(0.001)
Shipyards	-0.01	-0.01	-0.008***	-0.007**
	(0.009)	(0.008)	(0.003)	(0.003)

Notes: All coefficients come from a separate regression of a time-varying, industry-specific treatment variable on the outcome variable with individual fixed effects, Lasso-chosen controls and year dummies. Unemployment is a dummy equal to 1 f an individual has been registered as unemployed in in year t. Disability is a dummy equal to 1 if an individual has been registered as fully or partly disabled in year t. The treatment variable is equal to 1 for those who work in the targeted industry one year before the the introduction of the minimum wage. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Standard errors are clustered at the individual $industry_i \times year_t$, where $industry_i$ refers to the industry individual i works in the year before the policy intervention. * p<0.10, ** p<0.05, *** p<0.01.

Figure C4: Share of individuals who undergo unemployment/disability by year in treated and synthetic control group

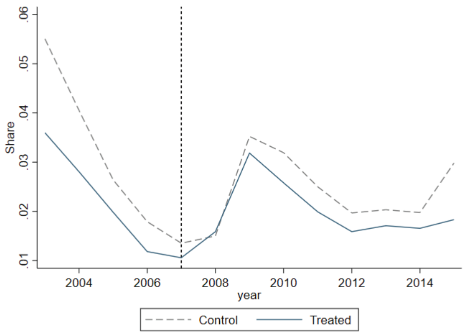

(a) Unemployment: Construction

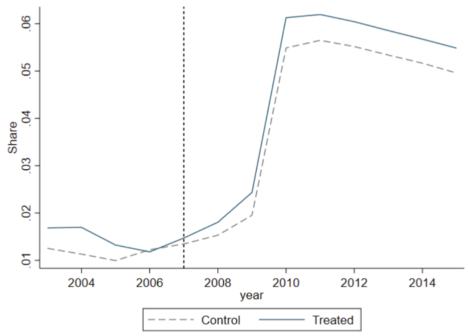

(b) Disability: Construction

(c) Unemployment: Cleaning

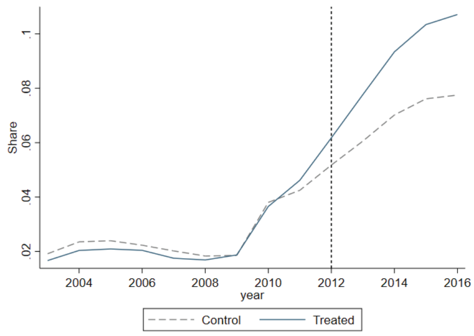

(d) Disability: Cleaning

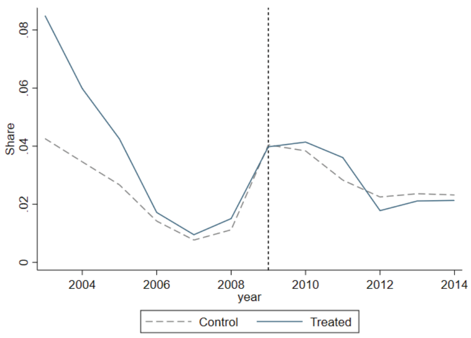

(e) Unemployment: Shipyards

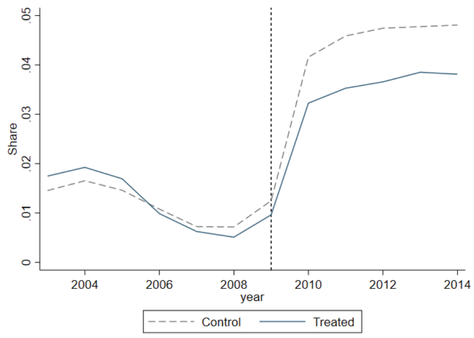

(f) Disability: Shipyards

Note: All graphs shows share of workers receiving unemployment and disability benefits. The treated group consist of individuals working in the targeted industry the year before the minimum wage introduction, while the control group consists of individuals working in a control industry the same year. This means that all workers are in the labor market in this given year, therefore the share of disabled and unemployed is lowest the year before the wage floor introduction. Control individuals are weighted with the weight for the industry they work in the year before policy implementation, w_s as calculated by synth and the inverse of $industry_s$'s share of workers in the synthetic control unit: w_s/α_s. Unemployment is a dummy equal to 1 f an individual has been registered as unemployed in in year t. Disability is a dummy equal to 1 if an individual has been registered as fully or partly disabled in year t.

C.2 Bootstrapped confidence intervals

Developments in econometrics suggest that inference based on cluster-robust standard errors in linear regressions using either Student's t distribution or the wild cluster bootstrap, is known to fail when the number of treated clusters is very small. Cluster robust standard errors tend to over-reject, often very severely, and wild bootstrap tests based on restricted residuals tend to under-reject just as severely MacKinnon and Webb (2017).

According to MacKinnon and Webb (2018) the standard wild bootstrap can yield valid interference provided cluster sizes are the same and intra-cluster correlation is the same for all clusters, assumptions that are likely not to be met in a difference-in-difference set-up. Failure of these assumptions can cause serious errors of inference when the number of treated groups is one as is the case in the regressions in this paper.

Furthermore, MacKinnon and Webb (2018) suggest that the sign of distortions caused by its failure are known and that we can be confident these assumptions are not seriously violated whenever bootstrap p-values for the restricted and unrestricted bootstrap are similar, with the former larger than the latter. For the restricted bootstrap, the null hypothesis is imposed when bootstrapping, while for the unrestricted bootstrap it is not. In practice, it can be tested if restricted and unrestricted differ substantially. However, when the two p-values are similar it does not guarantee that they are entirely reliable. MacKinnon and Webb (2018) findings suggest that the sub-cluster bootstrap may in some cases with only one treated group perform better than the ordinary bootstrap. This means bootstrapping standard errors at a finer level than the cluster. They recommend that in cases with few treated clusters, standard wild bootstrap fails with restricted bootstrap giving large p-value and unrestricted giving small p-value, the best is to combine cluster robust standard errors with bootstrap.

They also note that when each cluster contains M obs that can be evenly divided into S equal-sized subclusters. In this case unless intra-subcluster correlations are large relative to remaining intra-cluster correlations, potential gain from using actual subclusters instead of ordinary bootstrap is likely to be modest.Table C8 shows that as predicted, with industry-year clustering the p-value for the restricted and unrestricted bootstrap are similar and most likely intra-subcluster correlations are not large relative to remaining intra-cluster correlations. However as pointed out by Bertrand et al. (2004) clustering at year$\times$treated can be problematic as that assumes away serial correlation within industry over time. The current regressions do include a rich set covariate trends as well as individual fixed effects which may have the potential to capture such serial correlation (Cameron and Miller, 2015). Below I explore alternative ways of clustering using bootstrap and sub-cluster bootstrap.

Next, I cluster at the industry level. This means clustering at the industry that $individual_i$ works in the year before wage policy implementation. It is important to notice that the synthetic control industry for shipyards consists of six industries while cleaning and construction are compared to three. According to MacKinnon and Webb (2018) the bootstrap performs more consistently as the number of groups exceeds 4. Bootstrapping at the industry level reveals inconsistency: Now the unrestricted and restricted bootstrap have very different p-values. As predicted by MacKinnon and Webb (2018), by sub-clustering at the year level the two probabilities converge somewhat, but are still far apart for cleaning and construction. Therefore, the interference may be less reliable.

All confidence intervals include zero, but the probabilities are in the case of construction and cleaning concentrated to the left (negative) and to the right (positive) for shipyards. However, when using the sub-cluster bootstrap for shipyards, the confidence interval is almost symmetrical around 0. This is illustrated in Figure C5 graphing the confidence function of the estimate.

Table C8: Alternative bootstrap ordinary and subcluster

	Construction		Cleaning		Shipyards	
	Synthetic control:		Synthetic control:		Synthetic control:	
	3 industries		3 industries		6 industries	
Cluster level and method	Estimate: -.035		Estimate -.032		Estimate: .063	
	P-val	Conf. int.	P-val	Conf. int.	P-val	Conf int.
Cluster robust industry*year	.029	[-.11, -.006]	.023	[-.059, -.004]	.01	[.019,11]
Restricted: Bootstrap: indutry*year	.058	[-.11, .002]	.029	[-.060, -.003]	.02	[.013, .11]
Unrestricted: Bootstrap: industry*year	.056	[-.14, .001]	.025	[-.059, -.004]	.02	[.012, .12]
Cluster robust industry	.26	[-.20, .085]	.33	[-.12,.056]	.02	[.024, .064]
Restricted: Bootstrap by industry	.50	[-1.90, 1.60]	.72	[-.26, .12]	.30	[-.34, .40]
Unrestricted: Bootstrap industry	.37	[-.90, .67]	.35	[-.21, .11]	.00	[.016, .095]
Restricted: Sub-cluster bootstrap industry, year	.43	[-.71, .50]	.46	[-.25, .12]	.39	[-.54, .49]
Unrestricted: Sub-cluster bootstrap industry, year	.36	[-.59, .35]	.38	[-.20, .089]	.29	[-.069, .18]

Notes: The table shows the confidence intervals and p-values from different clustering and bootstrap procedures. Restricted refers to the bootstrap imposing the null-hypothesis, while unrestricted refers to the bootstrap not imposing the null-hypothesis. Moreover, industry refers to the industry an individual works in the year before policy implementation and the estimate is from a regression with individual and year fixed effects as well as control variables chosen usingLasso.

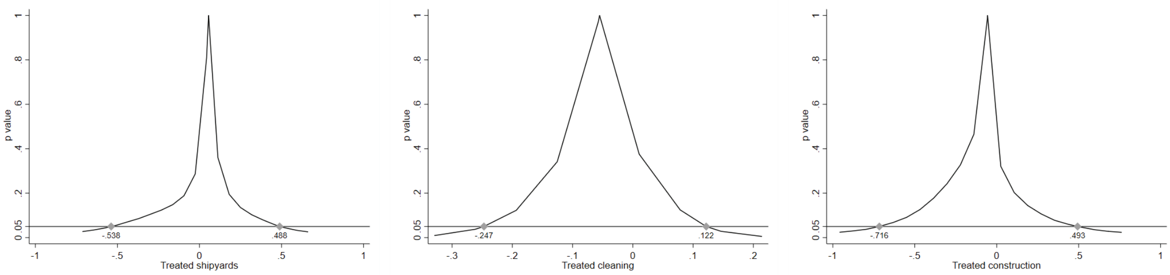

(a) Shipyards (b) Cleaning (c) Construction

Figure C5

Notes: The figures graph the confidence function for the sub-cluster bootstrap. The confidence interval come from bootstrapping sub-clustering on industry and year. This means m that bootstrapping is clustered by the intersections of the clustering variables. The bootstrap is done on the estimate from a regression comparing individuals in the synthetic control industry to those in the treated industry including individual fixed effects, year dummies and controls chosen by Lasso with standard errors clustered at the industry the year before the minimum wage introduction.

D Notes on running Lasso to select controls

While most machine learning techniques aim to classify or predict, the post-double-selection Lasso can be used to address omitted variable bias when many potential controls are available. This can be done using STATA's *pdslasso*(Ahrens et al., 2020).[43]

Although Lasso is a data-driven method and therefore potentially less sensitive to the researcher's choices, it is still the researcher supplying Lasso with a list of control variables and interactions. The variables and interactions included can potentially affect the result. Below I list the variables and interactions included in Lasso. All controls are variables potentially related to wages and occupation. Mainly this makes it possible to control for education, immigrant background and age in a very flexible way using all details that are available in the administrative registers. Together these make a total of 526 controls (all are interacted with $time$ and $time^2$).

D.1 List of controls Lasso:

Below is an exhaustive list of all control variables supplied to Lasso. The data on education is from the Norwegian Education Data Base (NUDB), demographics are from population registers and county is the county of work as recorded in the employee-employer data.

- Education:
  - Highest level of education as registered in the Norwegian Education database (NUDB) using the Norwegian Standard Classification of Education.[44]
    - Highest level of education level (first digit NUS code, from 1 to 8).
    - Education field for highest level of education (second digit nuscode, from 1 to 8)
  - Dummy for missing information on education.
  - Certificate of apprenticeship ("fagbrev")

[43] PDSLASSO also reports results from post-reglarization or CHS-methodology which instead of using Lasso-selected controls in a OLS estimation, construct orthogolaized versions of the dependent variables and the exogenous variable. However, in all regressions run this made no difference to the results and therefore these results are not reported.

[44] The framework and codes are available here https://www.ssb.no/klass/klassifikasjoner/36/.

  - Academic track in high school (”almennfag”/ ”studiespesialisering”)

- Region in Norway (”Fylke”)

- Demographics:

  - Birth year, gender, age

  - Immigrant: Dummy = 1 for individuals not born in Norway. Also included in interaction terms with immigrant gender, education and education level.

  - Years in Norway and years in Norway squared.

  - Country of birth group:

    1. Norway
    2. Nordic countries except Norway
    3. Countries entering the EU in 2004 (Estonia, Latvia, Lithuania, Poland, Czech Republic, Slovakia, Slovenia, Malta Cyprus and Hungary)
    4. Countries entering the EU in 2007 (Bulgaria, Romania)
    5. Northern Europe excluding Nordics
    6. Southern Europe (Portugal, Spain, Italy and Greece)
    7. Rest of Europe
    8. Africa
    9. Asia including Turkey.
    10. U.S, Canada and Australia
    11. Latin America

D.2 Sensitivity to choice of tuning parameter

Central to the Lasso procedure is the choice of the tuning parameter λ (lambda). There are three main ways of determining λ: 1) K-fold cross-validation, 2)Penalized goodness of fit and 3) Theory driven. For the theory, also called “rigorous” Lasso, the penalization is chosen to dominate the noise of the data-generating process (Ahrens et al., 2020). PDSLASSO by default implements a theory driven choice of lambda which in the robust case means:

$\lambda = 2c\sqrt{N}\Phi^{-1}(1 - \gamma/2p)$ where c and γ can be changed and n in the case with panel data is the number of clusters. Thus increasing γ decreases lambda and thus the penalization. The default choice in PDSLASSO is set at $c = 1.1$ and $\gamma = 0.1/log(n)$ (Ahrens et al., 2020). Alternatively, Belloni et al. (2014) suggest using $c = 1.1$ and $\gamma = .05$. In the case with 89 industry specific clusters the default γ value, .051 comes very close to this.

These choices are consequential for the number of controls retained by Lasso. Ahrens et al. (2020) proposes an alternative (also X-independent) with $\lambda = 2c\sigma p\sqrt{2nlog(2p/\gamma)}$ which according to the authors will lead to a more parsimonious model but also to a larger bias. In order to investigate the sensitivity of these choices I show how the controls selected change with the choice of lambda as γ is set to a very high value (1), the default (2) and alternative sparse model(3) in Figure D1. This results of this exercise suggest no specific pattern in which controls are dropped as less controls are included. Generally, the Lasso procedure results in a flexible way to control for age, birth year, education and immigrant background.

Moreover, Figure D1 shows how the estimate varies with γ. Figure D1 shows how the estimate changes with γ. The results show that small changes in the number of controls (higher *gamma*) does not matter much for the estimated coefficient.

	Gamma=.5	**Deafult(gamma = .05)**	**Alternative (sparse model)**
Construction Estimate Number controls Adjusted R-sq Root MSE	 -.028 (0008) 54 .58 .55	 -.053 (0009) 37 .58 .56	 -.051 (0009) 24 .58 .56
Controls included:	c.time#(byear: 1944 1948 1949 1951 1952 1953 1954 1955 1956 1958 1962 1963 1973 1974 1975 1977 1978 1982, age: 26 27 28 41 44 56 57 58 59 60 61, immigrant.byear: 1952 1956 1957 1962 1966, edlevel: 3), years in norway: 0 26 27, apprentice:0, high school academic: 0 1, edlevel: 2 3, edfield: 7, c.time2#(byear: 1946 1947 1950, age: 53 54 55 57 58, 0.male, ed level:7)	c.time#(byear:1948 1949 1951 1952 1954 1960 1962 1963 1974 1975, age: 26 27 28 41 45 56 58 59 60 61, years in Norway: 0 26 27, high school academic: 0 1, ed level: 3 7, edfield: 9) c.time2#(byear:1946 1947 1950, age: 53 54 55 57 58)	c.time#(byear:1949 1950 1952 1954 1974 1983, age: 27 41 56 59 60 61 , years in Norway: 0 27, high school academic: 0 1 , ed level: 3, field: 3 7,) c.time2#(byear:1948, age: 53 55 57 58)
Cleaning Estimate Number controls Adjusted R-sq OLs Root MSE	 -.038 (003) 68 .58 .60	 -.0 59 (003) 40 .57 .61	 -.071(003) 30 .57 .61
Controls included:	c.time#(byear:1944 1946 1947 1948 1949 1950 1951 1952 1953 1954 1955 1956 1957 1958 1961 1962 1964 1973 1974 1975 1976 1977 1978 1979 1980, age: 26 27 28 39 41 43 56 58 59 60 61, 1.region, 0.male, immigrant.byear:1952 1955 1956 1957 1958 1959 1961 1962 1963 1964, years in Norway: 3 24 26 27 44 , high school academic: 0 1, ed level: 2 3 6 7 ,) c.time2#(1959 1962 1947 1948 1964, age:57, edlevel: 6 7)	c.time#(byear:1947 1949 1950 1954 1955 1956 1958 1962 1974 1975 1976 1977 1978 1979 1980, age:26 27 28 41 43 56 58 59 60 61, immigrant.byear:1962, years in Norway: 3 26 27 44 57 58 61, high school academic: 0 1, ed level: 3) c.time2#(edlevel: 7, high school ac:0)	c.time#(byear: 1947 1948 1954 1955 1958 1974 1976 1977 1978 1979, age: 26 27 28 41 43 56 58 59 60, immigrant.byear: 1962, years in Norway: 3 26 58 61 3, high school academic 1 0 , .ed level: 3) c.time2#(byear:1964, years in Norway:57, high school ac.:0)
Shipyards Estimate Number controls Adjusted R-sq Root MSE	 .007(003) 63 .591 .0559	 -.0019 (003) 41 .59 .559	 -.0022(003) 27 .59 .561
Controls included:	c.time#(byear:1943 1944 1946 1948 1949 1950 1951 1952 1953 1954 1955 1956 1957 1958 1961 1962 1973 1974 1975 1976 1977 1978, age: 26 27 28 39 43 56 59 60 61, region:1 5 , immigrant.male byear.immigrant:1956 1958 1961 1962 1966 1979, years in Norway: 0 5 26 27 44, academic high school 1 0, ed level: 2 3 7) c.time.2#(byear.immigrant: 1959 1960 1961 1965 1966 1943 1944 1947 1948 1985, ed level: 6 7)	c.time#(byear:1948 1949 1950 1951 1952 1953 1954 1955 1956 1957 1958 1961 1962 1973 1974 1975 1976 1977 1978, age: 26 27 28 43 56 59 60 61, male.immigrant, byear.immigrant:1956 1958, years in Norway: 0 26 27, high school academic: 1 0, ed level: 2 3 7) c.time2#(byear:1961, immigrant.byear: 1947 , age.58,)	c.time#(byear:1948 1949 1950 1951 1952 1954 1955 1957 1958 1974 1976 1977 1978, age: 43 56 57 58 59 60 61, years in Norway: 0 , high school academic: 1 0 , ed level: 2 3) c.time.2#(byear:1947, high school.ac:0, age: 57 58)

Table D1: Selected controls with different choices of gammas

Notes: The table shows the results from using three different γ-values: 1) 0.5 (High), 2) 0.05 (default) and 3) Alternative sparse model (low). Estimate refers to the coefficient on treated from regression 3 with log yearly wage as the outcome variable. The adjusted R^2 and and the root-mean-square error are from an OLS-regression including the chosen control variables (in addition individual fixed-effects and year-dummies). Furthermore, the controls chosen for different values of γ are listed where *c.time* refers to a linear and *c.time*2 refers to a squared time-trend.

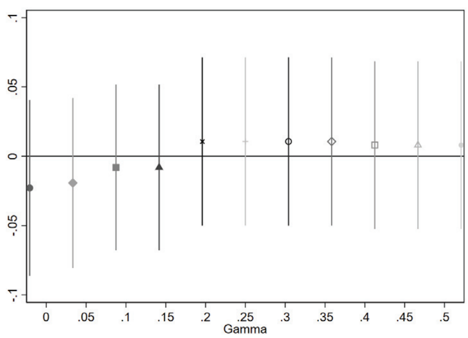

(a) Shipyards

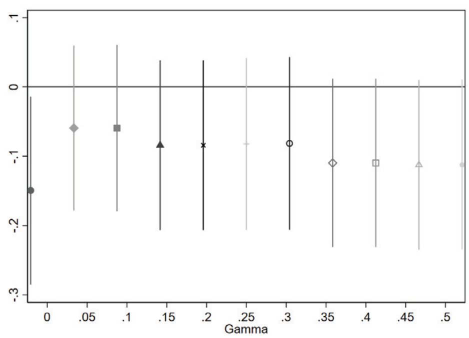

(b) Cleaning

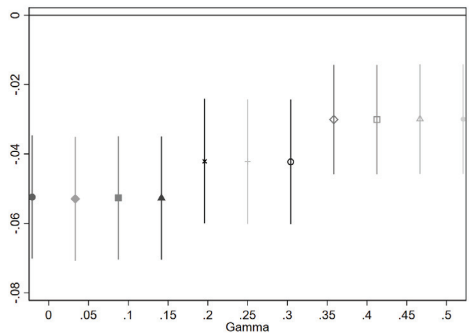

(c) Construction

Figure D1: Plot of how the estimated effect varies with the choice of gamma in post-double-selection Lasso

Notes: All plots show the coefficient from the OLS regression on log yearly wage on controls chosen using post-double selection Lasso with individual and year fixed effects and standard errors clustered at the industry level, varying the choice of γ. A smaller gamma gives a larger λ which is a larger penalization leading to fewer selected controls. Thus, the larger the gamma, the more control variables.