

The Effect of a Booming Local Economy in Early Childhood on the Propensity to Vote: Evidence from a Natural Experiment

HENNING FINSERAAS*

Growing up in a booming local economy can influence turnout in adulthood because family income influences the realization of cognitive abilities, investments in human capital and socio-economic status. Exploiting the discovery of oil outside the Norwegian county of Rogaland, this article identifies cohorts that experienced a shock in family income in childhood. This shock enables the effect of economic resources in childhood to be isolated from other characteristics of parents, such as their education level and personality traits. The study uses a differences-in-differences approach and finds that the affected cohorts are about 4 percentage points more likely to vote. The results suggest that potential mechanisms in addition to family income are changes in local public spending and in peers' political behaviour.

Many citizens choose not to exercise their right to vote. The decision to vote is correlated with socio-economic characteristics, thereby creating social inequalities in political participation. These inequalities are often considered a democratic problem since public policies are likely to reflect the interests of those who vote rather than the citizenry at large,¹ and might be viewed as particularly problematic if inequalities are transferred from one generation to the next. It is therefore notable that recent research on political participation suggests that pre-adult contexts and experiences are more strongly associated with the decision to vote than previously thought.² For instance, while the strong correlation between one's level of education and the propensity to vote has usually been interpreted as a direct result of education,³ recent advances suggest instead that education merely reflects parental characteristics and cognitive development prior to pursuing higher education.⁴

Two mechanisms on how parental characteristics affect the propensity to vote dominate in the literature. The conventional view is that parents transmit their own perceptions and values regarding political alienation, the importance of civic and political engagement, and political interest to their children.⁵ The key variable in this story is parental level of education, since parents' values are closely linked to their level of education.⁶ The alternative – yet not

* Institute for Social Research (e-mail: henning.finseraas@samfunnsforskning.no). I would like to thank Lars Gulbrandsen and Atle Hennum Haugsgjerd for help with data collection, and Bernt Aardal, Erling Barth, Johannes Bergh, Robert S. Erikson, Atle Hennum Haugsgjerd, Niklas Jakobsson, Rune Karlsen, Andreas Kotsadam, Axel West Pedersen, Jo Saglie, Øyvind Bugge Solheim, Kevin Young, the audience at the Friday Research Seminar at Trinity College Dublin, and anonymous reviewers for comments and suggestions. Grant number 199836 (Research Council of Norway) is acknowledged. Part of the data used in this article as provided, prepared and made available by the Norwegian Social Science Data Service (NSD). NSD is not responsible for the analyses/interpretation of data presented here. Data replication sets are available at <https://dataverse.harvard.edu/dataverse/BJPolS> and online appendices are available at <http://dx.doi.org/doi:10.1017/S0007123415000277>.

¹ Lijphart 1997.

² Verba, Burns, and Schlozman 2003.

³ E.g., Verba, Schlozman, and Brady 1995, 305; Wolfinger and Rosenstone 1980, 8–9.

⁴ Denny and Doyle 2008; Kam and Palmer 2008.

⁵ Jennings and Niemi 1968; Kam and Palmer 2008; Verba, Burns, and Schlozman 2003.

⁶ Verba, Burns, and Schlozman 2003, 47.

mutually exclusive – story is that parents transmit their political participation through genetic dispositions. For instance, Fowler and Dawes show that two genes, which are associated with pro-social behaviour, predict the propensity to vote, and argue that genes constitute an alternative explanation for why parental characteristics are correlated with voting.⁷

A third possibly important characteristic of parents is their economic resources. The research on the importance of parents' economic resources is surprisingly scant, particularly in light of the strong intergenerational transmission of economic position⁸ and the strong correlation between socio-economic position and turnout. As I discuss in the next section, parental income in childhood is likely to increase one's propensity to vote later in life through several channels. It can influence human capital development, not only by improving the quality of the child's formal education, but also by creating a home environment that facilitates the development of cognitive abilities.⁹ Parental income is also linked to better health conditions¹⁰ and higher earnings¹¹ in adulthood. Thus family income during childhood improves the propensity to vote by (1) improving abilities and thereby reducing the costs of obtaining political information and (2) by improving one's relative social status in adulthood.

The main empirical challenge is to disentangle the effect of family income from other turnout-enhancing characteristics of parents; a particular concern is the strong correlation between education and income. The discovery of oil outside the Norwegian county of Rogaland in 1969 can be considered a natural experiment that enables me to do so. The oil discovery led to a rapid growth in family income in the county, which was unrelated to parents' level of education.¹² Since parents' income growth is unrelated to parents' level of education, it is also plausibly unrelated to time-invariant characteristics of parents that might be important for political socialization and turnout, such as personality traits and genes.

I study the propensity to vote for a small number of Rogaland cohorts born right before the oil discovery using a differences-in-differences (DID) approach. The 'treated' cohorts grew up in families that experienced an exogenous growth in their parents' income during the early years of their childhood. I compare these cohorts with same-age cohorts born in other parts of Norway and to age cohorts born in Rogaland and other parts of Norway eight to ten years before the treated cohorts. These control cohorts either did not experience any oil discovery (those born in other parts of Norway) or experienced the income growth when they were older (those born in Rogaland). Thus I identify the effect on those who experienced the shock in early childhood, at an age when family resources are found to be of particular importance.¹³

I find that the cohorts born in Rogaland County right before the oil discovery are substantively more likely to vote. The estimated difference in the propensity to vote between the Rogaland cohorts and the control cohorts is about 4 percentage points. The empirical model I set up ensures that the estimates of the treatment effect account for effects of the oil discovery that affected everyone in Rogaland County the same, thus the results should be interpreted as the effect of the oil discovery on those born right before the oil boom. The results further suggest that the effect of the oil boom is too large to solely reflect the shock in family income. Additional analyses point to rapid increases in local public spending aimed at small children and

⁷ Fowler and Dawes 2008, 589.

⁸ Bowles and Gintis 2002.

⁹ Haveman and Wolfe 1995.

¹⁰ Willingham 2012.

¹¹ Heckman 2008.

¹² Løken 2010; Løken, Mogstad, and Wiswall 2012.

¹³ Heckman 2008.

the changed political behaviour of peers and parents of children in the ‘treated’ cohorts as possible additional channels from the oil boom to turnout. Thus, although the pre-adult experience that the oil boom represents appears to be important for turnout, the data do not allow me to identify the exact importance of family income as the mechanism.

The results are in line with the recent wave of research that identifies the importance of pre-adult experiences for outcomes in adulthood,¹⁴ and illustrate how economic inequalities can (re-)produce political inequalities. Norway’s small differences in earnings and income levels, low level of poverty and automatic voter registration imply that it is hard to detect the consequences of parental economic resources on turnout in Norway. Effects of parental economic resources are therefore likely to exist in other developed countries as well.

PARENTAL INCOME IN CHILDHOOD AND VOTING

The Socio-Economic Status (SES) model is usually considered the dominant framework for explaining stable differences across individuals in their probability of voting. In the SES framework, education is regarded as the main indicator of socio-economic status, and is argued to provide resources that lower the costs of voting; it also potentially increases the benefits of and motivation for voting.¹⁵ In line with the SES model, socio-economic characteristics are routinely found to be strongly correlated with the propensity to vote.¹⁶

A recent wave of research suggests that the relationship between current socio-economic status and political participation mainly represents characteristics of an individual’s parents and their pre-school experiences.¹⁷ Parents’ level of education has been emphasized as particularly important in this respect, as it is related to one’s own socio-economic status (education, income and occupation) as well as to cognitive skills, exposure to political stimuli at home and political interest.¹⁸

Parents’ level of education is not the only parental SES indicator that is likely to influence one’s propensity to vote. Although parents’ level of income is sometimes mentioned among other potentially important indicators of parents’ SES, it has not played an important role in this framework. Its importance for turnout in adulthood is largely unknown. This negligence is unfortunate, because parental income is associated with several of the same mechanisms that link parental education to one’s propensity to vote, making it a plausible determinant of the probability of voting.

Theoretically, we expect parents’ ability and willingness to invest in their children’s human capital to increase with family income.¹⁹ Tuition fees are negligible at Norwegian public universities. However, students usually need student loans to cover their expenses while studying, and the opportunity costs of education are large due to the loss of earnings and work-life

¹⁴ Heckman 2008.

¹⁵ E.g., Verba Schlozman and Brady 1995, 305; Wolfinger and Rosenstone 1980, 8–9. Education is argued to improve cognitive abilities, civic skills and civicness, and to place one in ‘mobilizing networks’, all of which makes one more likely to vote.

¹⁶ See, e.g., Blais 2000; Verba, Schlozman, and Brady 1995; Wolfinger and Rosenstone 1980. Strømsnes (2003) presents evidence from Norway.

¹⁷ Kam and Palmer (2008), Berinsky and Lenz (2011), Pelkonen (2012), Tenn (2007), but see also Sondheimer and Green (2010) and Henderson and Chatfield (2011). Pelkonen (2012) is of particular interest since he relies on Norwegian data. Following Aakvik, Salvanes, and Vaage (2010), he identifies exogenous variation in years of education through the gradual phasing in of compulsory schooling from seven to nine years. Pelkonen finds no evidence that the increase in years of education caused by the reform increased turnout.

¹⁸ See, e.g., Verba Schlozman and Brady (1995, ch. 15) and the discussion in Kam and Palmer (2008).

¹⁹ Becker and Tomes 1979.

experience during the years of education. Thus there is an economic risk involved in pursuing higher education, and parents' ability to function as a safety net depends on their income.²⁰

More importantly for this article, parents' investments in their children's human capital are not restricted to investments in the length and quality of formal education, as one's pre-school experiences and environment are very important for a child's attainments.²¹ Cognitive abilities are related to one's level of education, but are conceptually distinct to such a degree that education is argued to be a poor proxy for cognitive skills.²²

Well-off parents can better afford to create a home environment that improves cognitive abilities. For instance, improving the child's access to books will increase demands for 'cognitive complexity',²³ and such demands improve cognitive abilities. Well-off parents can also improve housing conditions, which affect performance at school,²⁴ cognition and the ability to concentrate.²⁵ In addition, the quality of schooling and day care depends on the area of residence, and well-off parents can afford to move to better areas.

There is also empirical evidence of the impact of early cognitive development on turnout. Denny and Doyle find that cognitive abilities measured at the age of eleven are significantly correlated with the probability of voting, and the inclusion of the cognitive ability variable renders the relationship between education and turnout insignificant.²⁶ This finding squares well with the child development and intervention literature, which shows that child cognitive and non-cognitive skills diverge before school age and change little afterwards.²⁷ Thus, irrespective of whether the relationship between formal education and turnout is causal, family income can affect voter turnout by improving cognitive abilities beyond the effect that formal education has on them.

The effects of parental income are also traceable to health problems later in life. Low income is related to less-nutritious diets, greater exposure to lead due to substandard housing and parental stress.²⁸ These health risks have an impact on cognitive abilities, but also affect health later in life. Smith finds that those who report having grown up in a family with a good economic status during childhood have better health in their fifties and sixties, even when controlling for level of education and current economic status.²⁹ Costello et al. study the effects of opening a casino in an Indian reservation in North Carolina that moved 14 per cent of the Indian families out of poverty.³⁰ They find that the income growth caused by the casino had a remarkably strong effect, as it improved the ex-poor children's psychiatric symptoms to the level of the non-poor children. In line with this research, the empirical research on the relationship between good health and turnout finds the expected positive relationship.³¹

Finally, family income and family influence during the child's early years are important for earnings and labour market outcomes in adulthood.³² By improving the child's future socio-economic status, family income can influence the propensity to vote through standard

²⁰ Løken 2010. Consistent with the risk of education framework, Løken, Mogstad, and Wiswall (2012) find causal evidence that family income increases years of education in Norway.

²¹ See, e.g., Haveman and Wolfe 1995; Heckman 2008.

²² Denny and Doyle 2008.

²³ Alwin and Thornton 1984, 787.

²⁴ Goux and Maurin 2005.

²⁵ Willingham 2012.

²⁶ Denny and Doyle 2008.

²⁷ Heckman 2008.

²⁸ Willingham 2012.

²⁹ Smith 2004, 119.

³⁰ Costello et al. 2003.

³¹ Denny and Doyle 2007; Schur and Kruse 2000.

³² Heckman 2008, but see Bleakley and Ferrie 2013.

mechanisms in the SES literature, such as higher political efficacy and political confidence. One might think of these as multiplier effects from a positive family environment in early life on turnout.

The main empirical challenge of identifying the causal effect of economic resources in early life is to disentangle the effect of economic resources from correlated turnout-enhancing parental characteristics, such as parental level of education and personality traits. In the next section I suggest that an oil discovery in Norway is a natural experiment that provides us with such exogenous variation in family income.

THE OIL DISCOVERY AS A NATURAL EXPERIMENT

The first oil discovery in Norway was made in 1969 in the North Sea outside the county of Rogaland in Southwestern Norway (see Figure A.1 in the Appendix for a map of Norway). Stavanger, the main city of Rogaland and Norway's fourth biggest city by population, became the country's 'oil capital': it was the closest city to the production areas, and it had the necessary infrastructure (including harbour, airport and helicopter base) and business-friendly local politicians.³³

The oil discovery had an immediate effect on the economy in the Rogaland area. During the 1970s, employment in the manufacturing industries increased, while it decreased in the rest of Norway due to the oil crisis-induced international recession.³⁴ Thus in the beginning of the Norwegian oil adventure, the Rogaland area benefitted from its closeness to the oil discovery and reaped the revenues of the natural resources.

After the first initial years of oil production, the revenues spread fairly rapidly throughout the rest of the economy due to heavy government taxation and regulation of the oil industry, direct government involvement in oil and gas production, and an explicit strategy that production should be spread across the country.³⁵ Moreover, from the mid-1970s the government ran big budget deficits in anticipation of the rapidly increasing oil revenues, and made heavy investments in infrastructure across the country.³⁶ The centralized wage bargaining system was also important for the rapid spread of revenues, as the wage negotiations in 1974 (effective the next year) were the most extravagant of the post-war period.³⁷

The oil discovery and the subsequent rapid spread of the revenues imply an exogenous shock in family income growth that was geographic and cohort specific. Thus it is possible to construct treatment and comparison groups to identify the effect of the shock, defining treatment as exposure to a booming economy in early childhood. The shock affected the formative first years for some cohorts of young children born in Rogaland right before the oil discovery, but did not affect the formative first years of children born at the same time in other parts of Norway. Nor did the shock influence the formative first years for older children, irrespective of their place of birth. The 'as if' random increase in family income at a young age for a small set of cohorts born in Rogaland

³³ Bråten et al. 2007, 23.

³⁴ Bråten et al. 2007, 26. At the time of the oil discovery, several big oil companies had already established offices in Stavanger to administer the search drilling. However, prior to 1969 the impact of the search activity on the region's economy was small. For instance, only thirty-two houses were built by the municipality to house people moving to Stavanger to conduct search drilling (Nerheim 1992, 101), and in the summer of 1969 several oil companies began reducing their activity in the region because the search results were disappointing (Nerheim 1992, 119–21).

³⁵ Hanisch 1992. At the end of the 1970s, oil and gas production (or preparation for it) had spread to several Norwegian counties (Bråten et al. 2007; Hanisch 1992).

³⁶ Hanisch 1992, 437f.

³⁷ Hanisch 1992, 438ff. The negotiations led to a whopping 25 per cent increase in nominal hourly wages; however, inflation was high in the mid-1970s, at 9.1 per cent in 1974.

makes it possible to construct credible treatment and control groups to identify the effect of the shock in economic resources on turnout later in life. All cohorts were, however, treated in older age (as oil revenues spread throughout the economy), an issue I return to below.

Previous empirical research on the consequences of the oil discovery has identified substantive effects on several social outcomes.³⁸ Through access to rich administrative data covering the entire Norwegian population, this research demonstrates that the growth rate of family income for the 1967–69 cohorts of Rogaland was significantly higher in all income quartiles for the years 1970–76, thus affecting most families in the county. In particular, annual family income from age one until age eleven was about 26,500 NOK (US \$4,569 in real 1999 income), and was higher for those born in Rogaland in 1967–69 compared to the same cohorts born in other counties.³⁹ Since the average annual family income for these cohorts in this age bracket was around 252,000 NOK (US \$43,450), the income shock was substantial. Importantly, the shock in family income was unrelated to parents' level of education, as parents typically had already made their investments in their own human capital at the time of the discovery.

Although all income quartiles experienced a growth in family income, relatively well-off families did benefit more than relatively poor families.⁴⁰ To what extent income inequality increased is difficult to assess due to lack of data, but Gerdrup presents some descriptive statistics on regional inequality since 1967.⁴¹ He shows that inequality increased quite rapidly in Rogaland from 1967 to 1975; however, this was also the case in other counties. From 1978 to 1981, however, inequality increased rapidly in Rogaland without the same marked development elsewhere. Gerdrup further shows that the increasing inequality in Rogaland from 1978 was due to a marked drop in the middle-income group's share of income – that is, the income growth of the rich outpaced that of the rest of the population.

The oil discovery had effects on social outcomes that were important for turnout, such as years of education and the probability of dropping out of high school.⁴² The effects are not restricted to formal education. Løken, Mogstad and Wiswall also identify effects on adult IQ using data from military records (military service is mandatory in Norway). The effects of the oil discovery on these outcomes make it highly plausible that the Rogaland 1967–69 cohorts will also be more likely to vote in elections.

ECONOMIC BOOMS AND VOTING

Thus far I have emphasized how the oil discovery increased family income in childhood and why family income might influence an individual's propensity to vote in adulthood. However, the shock in family income was not the only potentially relevant effect of the oil discovery. In addition to its effect on families' economic resources, shocks in the local economy might also have turnout-enhancing effects on those growing up through changes in local public spending or through behavioural changes in the local electorate.

The central Norwegian government had an explicit goal that as much of the oil revenues as possible should accrue to the central government, to be distributed across the country.⁴³ Thus the shock in private incomes was not to the same extent a shock in local public revenues. Still, increases in private incomes spill over to public revenues through local income taxes,

³⁸ Løken 2010; Løken, Mogstad, and Wiswall 2012.

³⁹ Løken, Mogstad, and Wiswall 2012, 14–15.

⁴⁰ Løken, Mogstad, and Wiswall 2012, 22.

⁴¹ Gerdrup 1996.

⁴² Løken 2010; Løken, Mogstad, and Wiswall 2012.

⁴³ Ryggvik 2014.

giving the local authorities the ability to increase public spending. The positive effects of the oil shock on local public spending imply that any turnout-enhancing effects of the oil discovery might be a mix of its positive effects on private as well as public resources. As I explain below, I identify the effect of the oil discovery by comparing the 1967–69 Rogaland cohorts to earlier cohorts from other counties as well as from Rogaland. This empirical strategy implies that changes in public spending that affect the 1967–69 cohorts in particular are most relevant, since the research design accounts for the effects of spending that affected everyone equally.

One of the key responsibilities of local authorities is the provision of schooling up until higher education. If education spending increased rapidly after the oil discovery, the quality of the education of the 1967–69 cohorts would be better than for earlier cohorts and those living elsewhere. Such an increase in the quality of education might have produced some of the beneficial effects of the oil discovery that Løken and Løken, Mogstad and Wiswall identify.⁴⁴

The 1970s was also a period of rapid expansion of public child care in Norway.⁴⁵ The expansion was most extensive in the late 1970s, but it started in the beginning of the 1970s. Havnes and Mogstad show that the expansion of child care had positive long-term effects on education levels and labour market participation, presumably because public child care replaced informal child care of lower quality.⁴⁶ It is possible that the increase in local public revenues caused an early expansion of public child care in Rogaland, affecting in particular the 1967–69 cohorts.

Local economic shocks might also change the political behaviour of the electorate, which in the next step might affect those growing up as they learn from their parents and peers. The effect of a boom in the economy on current voters is, however, theoretically ambiguous.⁴⁷ Some argue that economic strain causes a withdrawal from politics due to increases in alienation and distraction from politics. If the effects of economic shocks are symmetric, then this argument suggests that a boom will increase turnout. Increases in public revenues might increase the salience of elections, and more voters might develop attachments to political parties. Others argue that economic strain mobilizes voters by creating discontent, which increases efforts to hold politicians accountable. Moreover, personal wage growth might lower the propensity to vote, as higher hourly wages increase the opportunity costs of voting. Given symmetry, the mobilization argument suggests lower turnout due to the boom.⁴⁸

While the early empirical literature tends to support the withdrawal argument,⁴⁹ two recently published articles on US data find results consistent with the mobilization argument.⁵⁰ Charles and Stephens' article is particularly relevant, since they directly model the effects of local economic shocks.⁵¹ After correcting for endogeneity, they find negative effects of shocks to per capita earnings on turnout, and interpret it as a result of less time devoted to gathering political information. A negative effect of income shocks on parents' and peers' turnout does not rule out the possibility that income shocks can have positive effects on those growing up during the boom, but suggests that socialization into political behaviour works through parental or public economic resources rather than by observing parental and peers' political behaviour.

⁴⁴ Løken (2010) and Løken, Mogstad, and Wiswall (2012) argue against the possibility of an increase in the quality of education by reference to the central government's redistributive agenda.

⁴⁵ Havnes and Mogstad 2011.

⁴⁶ Havnes and Mogstad 2011.

⁴⁷ See Rosenstone 1982.

⁴⁸ There is also a literature on inequality and turnout, which suggests that if inequalities increased due to the shock then turnout will decline (see, e.g., Anderson and Beramendi 2008).

⁴⁹ Rosenstone 1982.

⁵⁰ Burden and Wichowsky 2014; Charles and Stephens 2013.

⁵¹ Charles and Stephens 2013.

The Dependent Variable

The dependent variable is a binary variable representing whether the respondent voted in the previous election. When possible, I use information on whether the respondent voted in the previous national election instead of the previous (national or local) election, to minimize the potential influence of variation across municipalities in the costs and/or benefits of voting. In most of the surveys, the respondent is asked a ‘yes’ or ‘no’ question of whether she voted in the previous national election, while in some surveys ‘did not vote’ is one of the response options for the question of what party the respondent voted for in the previous national election.⁵²

Treatment and Control Groups

The treatment group is those born in the years 1967–69 in Rogaland. The expectation is that these cohorts will be more likely to vote due to the many beneficial oil-boom effects identified for these cohorts.⁵³ Survey data are the only available data on turnout at the individual (or cohort) level, so a major challenge is to gather a large enough sample from the 1967–69 cohorts from Rogaland on which to conduct meaningful statistical analyses. This challenge is complicated by the fact that most surveys do not ask about county of birth, only about county of current residence. To get a reasonably sized treatment group I merge data from sixteen public opinion survey datasets, with a total of 172 respondents from the 1967–69 cohorts of Rogaland. The sixteen datasets are briefly described in the online appendix. The sample size is still small, however, and results should thus be treated with caution (see discussion in footnote 62 below).

Six of these datasets do not include information on county of birth, only municipality of current residence. For these surveys, county of birth is identified by a question on whether the respondent has lived in the municipality of residence her whole life.⁵⁴ Those who have not lived in their current municipality their whole life are excluded from the analysis. There are two obvious problems with this approach. First, the restriction on municipality of birth is too strict, since respondents who moved across municipalities within Rogaland County are excluded. Secondly, the decision to migrate from one’s municipality of birth might depend on family income in childhood. This could be through its effect on education, as more educated people are more mobile. Thus when identifying the treatment group using the mobility question, I condition on a potentially endogenous variable, which might introduce bias to the estimated treatment effect. Fortunately, I can assess the severity of the ‘migration bias’ by restricting the analysis to surveys in which the respondents are asked about their county or municipality of birth, and examine whether the coefficient changes when I restrict the sample to these surveys. All analyses are thus conducted on two samples: one that includes all surveys, and a smaller sample that is restricted to surveys that directly identify the county of birth.

⁵² It is well known that turnout is over-reported in surveys, in part due to social desirability bias. Previous research on Norwegian election data suggests that unstable voters are those most likely to give incorrect information on previous voting behaviour (Waldahl and Aardal 2000). If treatment led to more stable voting, then social desirability bias will be correlated with treatment status and, somewhat paradoxically, the estimated treatment effect will be biased downwards.

⁵³ Løken 2010; Løken, Mogstad, and Wiswall 2012.

⁵⁴ One survey, GGS-Norway, asks whether the respondent has lived in the ‘area’ her whole life. Two surveys, the local election study from 2003 and from 2007, do not include year of birth, only current age. These surveys were conducted at the end of the year, implying that year of survey minus current age will give the correct year of birth for most respondents. Furthermore, measurement error should not be correlated with the treatment group indicator. Results are not sensitive to the exclusion of these respondents.

The analytical strategy is to compare the turnout of those born in the years 1967–69 in Rogaland to turnout of a control group. I define the control group as respondents born in the years 1967–69 in the ten Norwegian counties that constitute the control group in Løken, Møgstad and Wiswall, in addition to those born in the years 1957–59 in Rogaland and in the same ten counties.⁵⁵ In an additional analysis I also include the intermediate cohorts. The ten counties constituting the control area are chosen since they have long driving distances to Rogaland, but at the same time are similar with respect to the family background characteristics of the population.⁵⁶ Table A.1 in the Appendix presents mean differences between Rogaland and the control area on a set of family background characteristics for children born from 1967–69. I compare the 1967–69 cohorts to the 1957–59 cohorts since the latter is not old enough to make generational differences an important problem, and the 1957–59 cohorts in Rogaland were beyond their most formative years when the oil discovery happened. Thus the comparison will produce the effect of growing up in a booming economy in early childhood on voting.

Empirical Strategy

To estimate the effect of the oil discovery I set up a DID model. More specifically, I estimate linear probability models of the following form, where i indexes the voters, j indexes the county, c indexes the cohort, t indexes the 1967–69 cohorts and s indexes the survey:

$$TURNOUT_i = \beta ROGALAND_j * TCOHORTS_t + \alpha_j + \pi_c + \gamma_s + \mu_i,$$

where $ROGALAND_j * TCOHORTS_t$ is a binary indicator equal to 1 if the respondent is born in Rogaland in 1967–69 and 0 otherwise (including if the respondent is born in Rogaland, 1957–59), α_j are county fixed effects, π_c are cohort fixed effects, γ_s are survey fixed effects, while μ_i is the error term, assumed to be normally distributed. I include county fixed effects to allow the counties in the control area to be dissimilar. The survey fixed effects are included to account for unobserved differences across surveys that affect all respondents in the particular survey similarly, such as variation across surveys in the length of time since the previous election. Finally, the cohort fixed effects are included to pick up cohort (generational) differences in turnout that are not related to the oil discovery.⁵⁷ In sensitivity analyses I address the sensitivity of the results to functional form assumptions and to changes in what counties constitute the control area.

The set-up is a standard DID set-up in a situation with geography cohort-based treatment assignment in which area fixed effects and cohort fixed effects are included. Note that the ‘main effects’ of $ROGALAND$ ⁵⁸ and $TCOHORT$ ⁵⁹ are absorbed by the area and cohort fixed effects. Models without area and cohort fixed effects are presented in the online appendix. β is the key estimate of interest, which tells us how the generational difference in voting differs in Rogaland compared to in the control area (hence the label DID). By comparing the 1967–69 cohorts from Rogaland to control cohorts within Rogaland, we ensure that we do not simply pick up an effect of the oil discovery that affected everyone in Rogaland, and by comparing the generational difference in Rogaland to that in a control area we ensure that we do not simply pick up a

⁵⁵ Løken, Møgstad, and Wiswall 2012. The ten counties are Sør-Trøndelag, Nord-Trøndelag, Hedmark, Oppland, Vestfold, Aust-Agder, Telemark, Sogn og Fjordane, Møre og Romsdal and Buskerud. Norway consists of nineteen counties, see Figure A.1.

⁵⁶ See Løken, Møgstad, and Wiswall 2012.

⁵⁷ E.g., Gallego 2009.

⁵⁸ Dummy equal to 1 if the respondent is born in Rogaland.

⁵⁹ Dummy equal to 1 if the respondent is born in the treated cohorts 1967–69, irrespective of county.

generational effect that affected everyone from the 1967–69 cohorts. Although those in the control group were affected by the oil revenues when the revenue spread throughout the economy, they were affected when they were older and not in early childhood. Thus a substantive and significant β is convincing evidence that the oil discovery had a particular effect on those born in Rogaland right before the oil discovery.

The DID-estimate is the causal effect of the oil discovery on turnout under the assumption that the oil discovery is the sole reason for the generational difference in voting between the two groups of cohorts (Rogaland and the control area). In the causal inference terminology, I estimate the intention-to-treat (ITT) effect, that is, the effect of being eligible for the treatment caused by the oil boom. As discussed above, treatment might consist of both treatment in the form of parental economic resources and treatment in the form of changes in public spending or local political behaviour, but the family income growth that the 1967–69 cohorts from Rogaland experienced in their early childhood is likely to be a particularly important part of the treatment. With access to the respondents' family income in childhood – which I do not have since it is never reported in the surveys – one could have estimated the direct effect of family income on turnout, the so-called treatment effect on the treated (TOT). The ITT can be very different from the TOT, since the ITT does not discriminate between those born in Rogaland 1967–69 who did not benefit from the oil discovery and those who did. It has been shown, however, that 'most families in Rogaland did experience effects of the oil shock'.⁶⁰ I return to this issue below when interpreting the results.

Some readers might view the model specification as overly parsimonious, given the myriads of factors that might influence the decision to vote. However, since I want to estimate the causal effect of belonging to the group of people experiencing the shock, I do not want to control for variables that are potentially influenced by the oil discovery. Family income during childhood is potentially important for most post-childhood variables, including socio-economic status in adulthood and attitudinal variables such as whether the respondent considers it a duty to vote in national elections. Thus, these variables could potentially be dependent variables in my analysis and, to use Angrist and Pischke's terminology, they are 'bad controls' when interested in the effect of the oil discovery on turnout, and should not be included in the empirical model.⁶¹

EMPIRICAL RESULTS

Main Results

The main results are presented in Table 1. In the first column, I present the results when including the surveys that indirectly identify the place of birth, while in the second column I restrict the sample to surveys that directly identify the county of birth.

Our key interest is to what degree the generational difference is smaller in Rogaland, which is the expectation based on the positive effects of the oil discovery. The positive and significant coefficient for the DID estimate, $ROGALAND_j * TCOHORTS_{jt}$, supports the expectation. The difference between the 1957–59 and 1967–69 cohorts in Rogaland deviates at about 4 percentage points from the average difference in the control counties.⁶² As evident, the estimates are similar in the two samples.

⁶⁰ Løken 2010, 122.

⁶¹ Angrist and Pischke 2009.

⁶² The sample size is small, and the results should of course be treated with caution. One reviewer suggests conducting a power analysis to assess the sufficiency of the sample size. Power calculations are usually conducted prior to the study, and even those advocating retrospective power calculations do not consider them

TABLE 1 *Linear Probability Model*

	All	Birthsample
ROGALAND × TCOHORTS	0.04** (0.017)	0.04** (0.017)
Survey FE	Yes	Yes
County FE	Yes	Yes
Cohort FE	Yes	Yes
Observations	2,336	1,787
No. treated	172	119

Note: the dependent variable is whether the respondent voted in the previous election. Robust standard errors adjusted for cohort-county clustering. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

I would like to stress again that I do not simply compare voters in Rogaland to voters in the control area, but rely on the generational difference within Rogaland to estimate $ROGALAND_j * TCOHORTS_t$. The DID design implies that the effects of the oil discovery that affected everyone in Rogaland the same – for instance changes in the incentives to participate in politics in order to redistribute the oil revenues or to keep a larger portion of the oil revenues in local rather than state hands – are accounted for by the set-up of the empirical model. Thus, as discussed above, in order for something other than the oil discovery to explain the DID estimate, this unknown factor has to have had a particular effect on the generational difference in Rogaland (or in the control area). I return to this issue below.

One can also find indications of an effect of the oil discovery on turnout from aggregated turnout data. What we can learn about the effect of the oil discovery from aggregated turnout data is, however, limited since aggregated data are not broken down by respondents' county of birth. Aggregated turnout in Rogaland is a mix of those born in Rogaland and those who have moved to Rogaland, and many of those born in Rogaland have moved to other counties, including counties in the control area. Thus we cannot learn much about the effect of growing up during the oil boom from examining aggregated turnout. However, it is fair to assume that the majority of those in the 1967–69 cohorts still lived in their county of birth when they voted for the first time, since they were young and the majority still live at home with their parents at that age. Thus unless geographic mobility is extremely high, one might expect the trend in turnout in Rogaland to deviate from the trend in the control counties at the time when these cohorts entered the electorate. The 1967 cohort were eligible to vote in the 1985 election, while the 1968 and 1969 cohorts could vote for the first time in the 1989 election. Although I find no indication that Rogaland municipalities deviate from the trend in turnout when the 1967 cohort entered the electorate, Figure 1 shows that the average difference in turnout between 1989 and 1985 was significantly higher (1.68 percentage points, $t = 5.38$, $N = 258$) in the Rogaland municipalities compared to the municipalities in the control group.⁶³

necessary if the estimated treatment effect is significant. That said, one might consider my study a 'pilot study' and use the estimated treatment effect, the estimated standard deviation of the outcome (turnout) and the ratio of treated to control units to calculate the sample size needed in order to have a high power in a replication study. Doing so suggests that a sample size of about 4,500 is needed to achieve a power of 0.80, i.e., an exact replication study needs 4,500 respondents in order to have an 80 per cent chance of detecting a true treatment effect of 4 percentage points.

⁶³ The dependent variable is the first difference in turnout at the municipality level.




Fig. 1. Difference in turnout between 1989 and 1985 at the municipality level in control counties and Rogaland County

Note: difference in mean is 1.7 percentage points ($t = 5.38$); $N = 258$.

Sensitivity Checks

In this section I present a set of sensitivity checks for the sample with the surveys that indirectly identify the place of birth; the same checks for the restricted sample are included in the Appendix.

In the main results I estimate linear probability models (LPM) even though the dependent variable is a binary indicator. In practice, there are two well-known problems with the LPM approach. First, it does not constrain predictions to be within the 0–1 interval. Secondly, conventional standard errors are known to be wrong due to inherent heteroscedasticity in the LPM. I am not very concerned with the first issue, since I am interested in the marginal effects rather than the ‘structural parameters’ (the predicted probabilities). Thus a low number of predictions outside the 0–1 interval is not a serious concern because it will not affect β much.⁶⁴ As for the second concern, I consistently rely on standard errors adjusted for heteroscedasticity to account for the inherent heteroscedasticity in the LPMs.⁶⁵ Nonetheless, Table 2 presents average marginal effects after logit and probit regressions to show that the conclusions are robust to non-linear model specifications.

While turnout varies across individuals, the key independent variable varies only across counties and cohorts. This implies that the standard error of β might be biased downwards. In Table 1 I rely on the conventional fix of adjusting the standard errors for clustering at the county-cohort level, which should produce accurate standard errors since the number of clusters (sixty-six) is sufficiently high.⁶⁶ One alternative approach is to aggregate the data to the county-cohort level and estimate the model using Weighted Least Squares with the number of observations at the county-cohort level as weights. Column 3 in Table 2 shows that this approach produces coefficient estimates and standard errors that are similar to those in Table 1.

Next I examined how sensitive the DID estimate is to the composition of the control area by re-estimating the main model ten times and leaving out one county each time in a rotating fashion. The DID estimate is 3 or 4 percentage points in all specifications and is always

⁶⁴ Wooldridge 2002, 455.

⁶⁵ See Angrist and Pischke (2009, 94ff.) for a defense of the LPM to estimate causal effects.

⁶⁶ E.g., Angrist and Pischke 2009, 293ff.

TABLE 2 *Alternative Model Specifications*

	Marginal effects after logit	Marginal effects after probit	Weighted least squares	Cohort 'placebo'
ROGALAND × TCOHORTS	0.03** (0.012)	0.03** (0.012)	0.04** (0.018)	0.01 (0.021)
County FE	Yes	Yes	Yes	Yes
Survey FE	Yes	Yes	Yes	Yes
Cohort FE	Yes	Yes	Yes	Yes
Observations	2,336	2,336	66	2,425
No. treated	172	172	6	196

Note: the dependent variable is whether the respondent voted in the previous election. Robust standard errors adjusted for cohort-county clustering in Columns 1, 2 and 4; robust standard errors in Column 3. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

significant at the 10 per cent level (results available upon request). Therefore the conclusions are not driven by small changes in the control group.

In the main analysis I consider the 1967–69 cohorts as treated, since they experienced the highest level of treatment intensity and because I have extensive information on how these cohorts were affected by the oil boom. However, there is a degree of arbitrariness in the decision to consider only these cohorts as treated, as earlier cohorts were treated as well, but at an older age. As an alternative, I consider the 1960–66 cohorts as partly treated, and rely on the 1950s cohorts as the control cohorts. The results are presented in the online appendix. The treatment coefficient for the 1967–69 cohorts increases to almost 6 percentage points in this specification, which is perhaps not that surprising since average age in the control group is now higher at the time of the boom. There is no significant treatment effect for the 1960–66 cohorts – that is, the treatment effect is concentrated among those with a very early exposure to the boom.

Finally, I conduct two 'placebo' analyses. First I pretend that the oil discovery happened ten years earlier and define the 1957–59 cohorts in Rogaland as treated, and use the 1947–49 cohorts as the control cohorts. The 1957–59 cohorts were not born immediately before the oil discovery, thus we should not expect the 1957–59 cohorts to be more likely to vote than the 1947–49 cohorts if the mechanism linking the 1967–69 cohorts to higher turnout is the booming local economy in early childhood. Reassuringly, the placebo DID estimate is much smaller than the true DID and is statistically insignificant (Table 2, Column 4).

In the second placebo analysis I re-estimate the model in Table 1 ten times and in rotating fashion pretend that the oil discovery happened in the other counties.⁶⁷ Neither of these mock treatment groups experienced an exogenous income shock in this period, so these regressions should produce insignificant treatment coefficients. The placebo DID estimates from these regressions are presented in Figure 2. Reassuringly, all confidence intervals include zero.

Interpretation

The empirical analysis has provided convincing evidence that the 1967–69 cohorts from Rogaland are more likely to vote in national elections. The evidence presented here – together with that in Løken and Løken, Mogstad and Wiswall – makes the oil discovery a very plausible

⁶⁷ I consistently exclude the data from Rogaland when running these placebo regressions.




Fig. 2. Mock-treatment effects from placebo regressions

Note: dots represent coefficient estimates, while the lines represent the 95 per cent confidence intervals.

reason why these birth cohorts are more likely to vote.⁶⁸ In this section I discuss how to interpret this result, and assess the possible mechanisms for the finding.

There is a long-standing debate regarding whether economic resources (and education) have absolute or relative effects on turnout.⁶⁹ Many of the turnout-enhancing mechanisms of the oil boom that I point to above are about absolute resources, which can improve cognitive abilities. However, if people compare themselves to others of the same age in other parts of the country (for example, if the reference point of comparison is national), then the estimated effect is also partly a relative effect. Those who were young in Rogaland at the time of the discovery improved their relative standing compared to those who were young in other parts of the country. In addition, identification is based on place of birth; as people have moved in and out of Rogaland, there will be a relative effect (although probably quite small) even if the reference point is local. Thus the estimated effect should not be interpreted as solely reflecting an absolute effect on turnout.

I have made the argument that the most important mechanism linking the oil discovery to political participation is the effect of the oil discovery on family income in childhood. However, Tables 1 and 2 estimate the effect of the oil discovery, not the effect of family income. As mentioned above, the effect of the oil discovery – the ITT effect – can be viewed as the effect of being *eligible* for a shock in family income, which will be very different from the *effect* of experiencing an increase in family income if few families in the treated Rogaland cohorts experienced an increase in family income as a result of the oil discovery. If few were affected, the ITT estimated above is likely to reflect some other turnout-enhancing effect of the oil discovery than the effect of family income. Løken shows, however, that most families were affected by the shock.⁷⁰

⁶⁸ Løken 2010; Løken, Mogstad, and Wiswall 2012.

⁶⁹ Campbell 2009; Persson forthcoming.

⁷⁰ Løken 2010, 122–3. However, all families did not receive the same amount of treatment. Since inequalities increased (see discussion above), some families probably benefitted more from the oil discovery than others, i.e., treatment intensity varies.

The detailed presentation of the shock in family income on the 1967–69 cohorts in Løken, Mogstad and Wiswall, together with census data from Statistics Norway, allow for a crude and approximate estimate of the relationship between family income and turnout even though I do not have access to data on the family income of the respondents.⁷¹ To do so, I follow the logic of Angrist's two-sample IV approach.⁷² This approach exploits the fact that the coefficient for the endogenous variable in the second stage of an IV regression (in this case, the unknown coefficient for family income on turnout) is equal to the ratio of two numbers that I do have. The numerator of the ratio is the reduced-form coefficient, in this case the coefficient representing the turnout difference between the Rogaland 1967–69 cohorts and the control group, while the denominator is the first-stage coefficient, which is the coefficient representing the difference in family income between the Rogaland 1967–69 cohorts and the control group. The two-sample IV relies on the same assumptions as the standard IV estimator, the most important of which are the absence of weak instruments and the exclusion restriction.⁷³

Løken, Mogstad and Wiswall show that annual family income between the ages of two and twelve was 26,500 NOK (US \$4,569) higher for the 1967–69 cohorts in Rogaland compared to the control counties.⁷⁴ Unfortunately, they do not report the corresponding difference for the 1957–59 cohorts. Using the 1970 census data from Statistics Norway, I estimate the difference for these cohorts to be approximately 11,400 NOK (US \$2,021). Thus the first-stage coefficient is about 1.51 in 10,000 NOK (US \$1,724). The relationship between birth cohort and family income is very strong, which implies that there is no bias from weak instruments.⁷⁵ In Table 1 the turnout difference between the treatment and control groups – the reduced-form coefficient – is estimated to be about 0.04. The effect of a 10,000 NOK increase in annual family income on the probability of voting is thus $0.04/1.51 = 0.026$ – that is, a 10,000 NOK increase in family income in childhood increases the estimated probability of voting by about 2.6 percentage points. It is important to understand that this estimate should be interpreted as a Local Average Treatment Effect (LATE),⁷⁶ in this case an estimate of the effect of family income on turnout for those who would not have experienced the increase in family income absent the oil discovery, and not a general causal effect of family income on turnout.

Even with the LATE interpretation in mind, and the fact that annual differences in family income sum up to large amounts over time, the crude (and admittedly somewhat speculative) estimate of the income effect is too large to be a plausible estimate of the effect of family income. The effect of family income appears too large even to be considered as an upper-bound estimate of the effect of family income on turnout. The large effect suggests that the oil discovery had an impact on turnout through additional channels other than family income. Thus the estimate in Table 1 appears to be the sum of the impact of several turnout-enhancing effects of a booming economy in early childhood.

I suggested above that a growth in local government revenue that benefitted the 1967–69 cohorts in particular might be one additional channel of how the oil shock might have influenced the turnout of these cohorts. A rapid growth in spending on education might have such turnout-enhancing effects if there is a causal effect of education on turnout. However, Figure A.2 in the Appendix shows that the trend in per capita current expenditure on education in Rogaland was similar to (or even weaker than) the control area in the period 1972–82.

⁷¹ Løken, Mogstad, and Wiswall 2012.

⁷² Angrist 1990. See also Angrist and Pischke (2009, 147ff.) and Dunning (2012, 143ff.).

⁷³ Angrist and Pischke 2009.

⁷⁴ Løken, Mogstad, and Wiswall 2012, 15–16.

⁷⁵ Løken, Mogstad, and Wiswall 2012.

⁷⁶ See, e.g., Angrist and Pischke 2009, 155.

Only after 1982 did Rogaland have a rapid growth in education spending. Education spending after 1982 should have a limited effect on the 1967–69 cohorts. A perhaps more precise measure of the quality of schooling is the number of students per teacher. Using data from historical sources on the number of pupils and teachers in primary schools,⁷⁷ I calculate the student-teacher ratios in 1970 and 1981. In Table A.2 I rank the counties by the changes in this ratio in the years 1970–81. Table A.2 shows that the number of pupils per teacher increased in most counties, including in Rogaland, but Rogaland does not stand out in comparison to the control counties. Thus the improved quality of local schools does not appear to drive the effect on turnout.

Improvements in public child care might be another mechanism linking the oil discovery to turnout, under the assumption of turnout-enhancing effects of replacing lower-quality informal child care with public child care.⁷⁸ Table A.3 documents quite large increases in public child care spending in most counties from 1967 to 1974,⁷⁹ which covers the period the 1967–69 cohorts were exposed to public child care. Moreover, the increase in Rogaland was higher than in most of the control counties. Thus, the public child care expansion appears to be one additional channel of how the oil discovery might have influenced turnout. Although it is impossible to assess the relative importance of the increase in public child care spending to the increase in family income, it seems highly implausible that child care spending is equally important to parental resources. However, the relatively large increase in child care spending in Rogaland makes it plausible that public resources have increased in other areas of importance for the 1967–69 cohorts as well.

A second possibility is that the oil discovery changed the political behaviour of voters in Rogaland. If peers and parents in Rogaland became more likely to vote due to the oil boom, those growing up might learn to vote from them. This is a big if, however, since recent research suggests that income growth depresses turnout.⁸⁰ Changes in aggregate turnout between 1973 and 1969 did in fact significantly deviate in Rogaland from changes in the control area (Figure 3), and average turnout across Rogaland municipalities was higher than in the control area in the 1977 and 1981 elections, despite being slightly lower in 1969 (see Table A.4). Thus in contrast to the ‘discontent as mobilization’ hypothesis,⁸¹ turnout appears to have increased among peers. However, since I also base identification on within-Rogaland differences, the effect on peers should have had a particular influence on the 1967–69 cohorts and/or be especially prominent among the parents of the 1967–69 cohorts in order to explain the results above. The former is hard to assess, but to explain the results, the peer effect must have a very strong cumulative effect (that is, a few more years of exposure to peers’ political behaviour makes a big difference) or be particularly important at a very young age (since the treated cohorts were exposed at a young age whereas the control cohorts were exposed when they were older). The effect on parents is also hard to assess, but the Norwegian National Elections Study has a panel feature in which about 1,000 respondents are followed in two subsequent elections. Unfortunately, in the 1969–73 panel dataset, respondents were asked about their region of residence instead of county of residence, and regions include several counties. For what it is worth, I find an insignificant coefficient for being a parent in the Rogaland region on change in turnout from 1969 to 1973, but the sample size is small ($N = 101$), which makes it difficult to draw strong conclusions. The size of the coefficient is quite large, which suggests a possible effect on parents.⁸²

⁷⁷ SSB 1971, 1983.

⁷⁸ Havnes and Mogstad 2011.

⁷⁹ SSB 1969, 1976.

⁸⁰ Charles and Stephens 2013.

⁸¹ Burden and Wichowsky 2014; Charles and Stephens 2013.

⁸² The dependent variable is voting in 1973 minus voting in 1969, i.e., a categorical variable in which -1 implies that the respondent did not vote in 1973, but did vote in 1969 (demobilized); 0 implies the same




Fig. 3. Difference in turnout between 1973 and 1969 at the municipality level in control counties and Rogaland County

Note: difference in mean is 1.6 percentage points ($t = 2.78$); $N = 260$. Difference in means is 1.1 ($t = 2.93$) if I exclude the two extreme observations.

Another plausible political behaviour mechanism is turnout contagion among Rogaland children. The decision to vote is influenced by whether people in your social network are voting.⁸³ Since people from the same cohort tend to share more networks – they go to school together, they are friends and they marry – multiplier effects from social interaction might increase the effect of the oil discovery beyond the direct effect of family income. People in the treatment group who barely benefitted from the oil discovery interacted with people who benefitted a lot, and if the latter transmitted some of the turnout effect to the former, the difference between the treatment group and the control group will increase. People who benefitted a lot from the oil discovery of course interacted with people in the control group as well, which pulls the treatment effect downwards, but contagion is likely to be larger within the treatment group since one tends to share more networks within one's own cohort. This type of contagion can be viewed as a violation of the assumption of independence between units,⁸⁴ because potential outcomes depend not only on an individual's treatment, but also on the intensity of treatment within his or her social network. I find this type of contagion to be very plausible, and it can explain why the estimated effect of family income in the IV model is large.

A final caveat regards treatment heterogeneity. Since the development of children's cognitive abilities depends on their 'initial' level of cognitive abilities,⁸⁵ it is conceivable that the effect of family income is heterogeneous across parents' socio-economic status. Indeed, Løken, Mogstad

behaviour in both elections; while 1 implies that the respondent did vote in 1973, but did not vote in 1969 (mobilized). I regress this variable on whether the respondent has children (ordinary least squares). The sample is restricted to the Rogaland region. The coefficient for being a parent is 0.05, $t = 0.99$, $N = 101$.

⁸³ Fowler 2005; Nickerson 2008; Sinclair, McConnell, and Green 2012.

⁸⁴ E.g., Sinclair, McConnell, and Green 2012.

⁸⁵ Heckman 2008.

and Wiswall find stronger treatment effects of family income at low levels of income.⁸⁶ Without access to family income I cannot explore whether there is important treatment heterogeneity across the socio-economic status of parents. Norwegian authorities are currently working on digitalizing the voter census, implying that in the near future there will be administrative, population-based data on turnout at the individual level. When this data are available for research, these and other related issues can be examined further.

CONCLUSION

What happens in early childhood is important for a range of different outcomes in adulthood. The development of cognitive abilities, health and labour market outcomes are all influenced by family resources in childhood.⁸⁷ These effects of economic resources make it plausible that political participation in adulthood will be influenced as well. However, theoretical models of turnout have emphasized the importance of other characteristics of parents, such as their level of education and personal values, and empirical evidence on the importance of economic resources in childhood is lacking.

In an effort to shed light on the importance of family resources, I exploit the fact that the discovery of oil outside the county of Rogaland led to a boom in the local labour market and a rapid earnings growth for families in the county. The shock was unrelated to other characteristics of parents, which is useful for estimating the importance of family income. I establish that a group of cohorts that experienced the oil boom in their early childhood is about 4 percentage points more likely to vote, compared to a similarly aged group of voters that did not experience the oil boom in their early childhood.

While the analysis makes it highly plausible that the oil boom explains why these cohorts are more likely to vote, the article is less successful in identifying the exact influence of family income in this process. Using secondary source information on the size of the income shock, the results suggest that a 10,000 NOK (US \$1,724) increase in family income due to the oil discovery increases the probability of voting by more than 2 percentage points. The sheer size of the estimated effect of family income in childhood suggests that the oil boom had an impact on turnout beyond the direct effect of family income. Further analyses point to increases in public spending on child care and increased voting among peers and parents as additional mechanisms for how the oil boom affected these specific cohorts.

Although several mechanisms are clearly in operation, the results and the nature of the oil boom make it implausible that family income is not one of the most important mechanisms in play. Therefore I argue that the results provide empirical support to the SES model, but they also restrict and nuance it. Verba, Schlozman and Brady emphasize two paths for the intergenerational transmission of political participation.⁸⁸ One path runs through political stimuli: well-educated parents are more likely to be politically active and to discuss politics at home. The other path is socio-economic: parents' education level influences the socio-economic status of their children. Both of these paths emphasize parents' level of education. The research design in this article holds the effect of parental level of education constant because parents' investments in education were largely made at the time of the oil discovery. Furthermore, the research design accounts for time-invariant characteristics of the parents, such as their personality, basic value system and genetic transmission. The findings therefore suggest that the

⁸⁶ Løken, Mogstad, and Wiswall 2012.

⁸⁷ Heckman 2008.

⁸⁸ Verba, Schlozman, and Brady 1995, 437ff.

socio-economic path is very important for the propensity to vote, but is perhaps less linked to parents' level of education than previously suggested. In agreement with Kam and Palmer, I suggest that researchers within the SES school should place more emphasis on directly establishing the role of parents and pre-adult experiences.⁸⁹

REFERENCES

- Aakvik, Arild, Kjetil G. Salvanes, and Kjetil Vaage. 2010. Measuring Heterogeneity in the Returns to Education Using an Education Reform. *European Economic Review* 54:483–500.
- Alwin, Duane F., and Arland Thornton. 1984. Family Origins and the Schooling Process: Early Versus Late Influence of Parental Characteristics. *American Sociological Review* 49:784–802.
- Anderson, Christopher J., and Pablo Beramendi. 2008. Income, Inequality, and Electoral Participation. In *Democracy, Inequality, and Representation*, edited by Pablo Beramendi and Christopher J. Anderson, 278–311. New York: Russell Sage Foundation.
- Angrist, Joshua D. 1990. Lifetime Earnings and the Vietnam Era Draft Lottery: Evidence from Social Security Administrative Records. *American Economic Review* 80:313–36.
- Angrist, Joshua D., and Jorn-Steffen Pischke. 2009. *Mostly Harmless Econometrics*. Princeton, NJ: Princeton University Press.
- Becker, Gary S., and Nigel Tomes. 1979. An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility. *Journal of Political Economy* 87:1153–89.
- Berinsky, Adam J., and Gabriel S. Lenz. 2011. Education and Political Participation: Exploring the Causal Link. *Political Behavior* 33:357–73.
- Blais, Andre. 2000. *To Vote or Not to Vote: The Merits and Limits of Rational Choice Theory*. Pittsburgh, PA: University of Pittsburgh Press.
- Bleakley, Hoyt, and Joseph P. Ferrie. 2013. Shocking Behavior: Random Wealth in Antebellum Georgia and Human Capital Across Generations, Working Paper No. 19348. Cambridge, MA: NBER.
- Bowles, Samuel, and Herbert Gintis. 2002. The Inheritance of Inequality. *Journal of Economic Perspectives* 16:3–30.
- Bråten, Magne, Gudmund Hernes, Jon M. Hippe, and Jørgen Svalund. 2007. Forutsetninger for olje- og gassvirksomhet i Nord-Norge [Preconditions for oil and gas production in Northern Norway], Working Paper No. 19. Oslo: FAFO.
- Burden, Barry C., and Amber Wichowsky. 2014. Economic Discontent as a Mobilizer: Unemployment and Voter Turnout. *Journal of Politics* 76:887–98.
- Campbell, David E. 2009. Civic Engagement and Education: An Empirical Test of the Sorting Model. *American Journal of Political Science* 53:771–86.
- Charles, Kerwin Kofi, and Melvin Stephens. 2013. Employment, Wages, and Voter Turnout. *American Economic Journal: Applied Economics* 5:111–43.
- Costello, E. Jane, Scott N. Compton, Gordon Keeler, and Adrian Angold. 2003. Relationships Between Poverty and Psychopathology: A Natural Experiment. *Journal of the American Medical Association* 290:2023–9.
- Denny, Kevin J., and Orla M. Doyle. 2007. '...Take up Thy Bed, and Vote'. Measuring the Relationship Between Voting Behaviour and Indicators of Health. *European Journal of Public Health* 17:400–1.
- Denny, Kevin, and Orla Doyle. 2008. Political Interest, Cognitive Ability and Personality: Determinants of Voter Turnout in Britain. *British Journal of Political Science* 38:291–310.
- Dunning, Thad. 2012. *Natural Experiments in the Social Sciences: A Design-Based Approach*. Cambridge: Cambridge University Press.
- Fowler, James. 2005. Turnout in a Small World. In *Social Logic of Politics*, edited by Alan Zuckerman, 269–87. Philadelphia, PA: Temple University Press.
- Fowler, James H., and Christopher T. Dawes. 2008. Two Genes Predict Voter Turnout. *Journal of Politics* 70:579–94.

⁸⁹ Kam and Palmer 2008.

- Gallego, Aina. 2009. Where Else Does Turnout Decline Come From? Education, Age, Generation and Period Effects in Three European Countries. *Scandinavian Political Studies* 32:23–44.
- Gerdrup, Karsten. 1996. *Inntektsfordeling og økonomisk vekst i norske fylker: En empirisk studie basert på data for perioden 1967-93* [Income distribution and economic growth in Norwegian counties: An empirical study of 1967-93]. Oslo: Statistics Norway.
- Goux, Dominique, and Eric Maurin. 2005. The Effect of Overcrowded Housing on Children's Performance at School. *Journal of Public Economics* 89:797–819.
- Hanisch, Tore J. 1992. Oljen i norsk økonomi [Oil and the Norwegian economy]. In *Norsk oljehistorie: Fra vantrø til overmot?* [Norwegian oil history: From disbelief to overconfidence?], edited by Tore J. Hanisch and Gunnar Nerheim, 405–57. Oslo: Leseselskapet.
- Haveman, Robert, and Barbara Wolfe. 1995. The Determinants of Children's Attainments: A Review of Methods and Findings. *Journal of Economic Literature* 33:1829–878.
- Havnes, Tarjei, and Magne Mogstad. 2011. No Child Left Behind: Subsidized Child Care and Children's Long-Run Outcomes. *American Economic Journal: Economic Policy* 3:97–129.
- Heckman, James J. 2008. Role of Income and Family Influence on Child Outcomes. *Annals of the New York Academy of Sciences* 1136:307–23.
- Henderson, John, and Sara Chatfield. 2011. Who Matches? Propensity Scores and Bias in the Causal Effects of Education on Participation. *Journal of Politics* 73:646–58.
- Jennings, M. Kent, and Richard G. Niemi. 1968. The Transmission of Political Values from Parent to Child. *American Political Science Review* 62:169–84.
- Kam, Cindy D., and Carl L. Palmer. 2008. Reconsidering the Effects of Education on Political Participation. *Journal of Politics* 70:612–31.
- Lijphart, Arend. 1997. Unequal Participation: Democracy's Unresolved Dilemma. *American Political Science Review* 91:1–14.
- Løken, Katrine V. 2010. Family Income and Children's Education: Using the Norwegian Oil Boom as a Natural Experiment. *Labour Economics* 17:118–29.
- Løken, Katrine V., Magne Mogstad, and Matthew Wiswall. 2012. What Linear Estimators Miss: The Effects of Family Income on Child Outcomes. *American Economic Journal: Applied Economics* 4:1–35.
- Nerheim, Gunnar. 1992. Den første leteboringen i Nordsjøen og på norsk sokkel [The first search drills for oil in the North Sea]. In *Norsk oljehistorie: Fra vantrø til overmot?* [Norwegian oil history: From disbelief to overconfidence?], edited by Tore J. Hanisch and Gunnar Nerheim, 75–126. Oslo: Leseselskapet.
- Nickerson, David W. 2008. Is Voting Contagious? Evidence from Two Field Experiments. *American Political Science Review* 102:49–57.
- Pelkonen, Panu. 2012. Length of Compulsory Education and Voter Turnout – Evidence from a Staged Reform. *Public Choice* 150:51–75.
- Persson, Mikael. forthcoming. Education and Political Participation. *British Journal of Political Science*.
- Rosenstone, Steven. 1982. Economic Adversity and Voter Turnout. *American Journal of Political Science* 26:25–46.
- Ryggvik, Helge. 2014. Norsk oljehistorie [Norwegian oil history]. Store norske leksikon [The big Norwegian encyclopaedia]. Available from https://snl.no/Norsk_oljehistorie, accessed 6 October 2014.
- Schur, Lisa A., and Douglas L. Kruse. 2000. What Determines Voter Turnout? Lessons from Citizens with Disabilities. *Social Science Quarterly* 81:571–87.
- Sinclair, Betsy, Margaret McConnell, and Donald P. Green. 2012. Detecting Spillover in Social Networks: Design and Analysis of Multilevel Experiments. *American Journal of Political Science* 56:1055–69.
- Smith, James P. 2004. Unraveling the SES-Health Connection. *Population and Development Review* 30:108–32.
- Sondheimer, Rachel M., and Donald P. Green. 2010. Using Experiments to Estimate the Effects of Education on Voter Turnout. *American Journal of Political Science* 54:174–89.
- SSB. 1969. *Barneomsorg* [Child care]. Oslo: Norges offisielle statistikk.

- . 1971. *Undervisningsstatistikk Folke og framhaldsskoler* [Education statistics: Primary education]. Oslo: Norges offisielle statistikk.
- . 1976. *Barneomsorg* [Child care]. Oslo: Norges offisielle statistikk.
- . 1983. *Utdanningsstatistikk: Oversikt* [Education statistics: An overview]. Oslo: Norges offisielle statistikk.
- Strømsnes, Kristin. 2003. *Folkets makt: Medborgerskap, demokrati, deltakelse* [The power of the people: Citizenship, democracy, participation]. Oslo: Gyldendal.
- Tenn, Steven. 2007. The Effect of Education on Voter Turnout. *Political Analysis* 15:446–64.
- Verba, Sidney, Kay L. Schlozman, and Henry E. Brady. 1995. *Voice and Equality: Civic Voluntarism in American Politics*. Cambridge, MA: Harvard University Press.
- Verba, Sidney, Nancy Burns, and Kay L. Schlozman. 2003. Unequal at the Starting Line: Creating Participatory Inequalities Across Generations and Among Groups. *The American Sociologist* 34:45–69.
- Waldahl, Ragnar, and Bernt Aardal. 2000. The Accuracy of Recalled Previous Voting: Evidence from Norwegian Election Study Panels. *Scandinavian Political Studies* 23:373–89.
- Willingham, Daniel T. 2012. Why Does Family Wealth Affect Learning? *American Educator* 36:33–9.
- Wolfinger, Raymond E., and Steven J. Rosenstone. 1980. *Who votes?* New Haven, CT: Yale University Press.
- Wooldridge, Jeffrey M. 2002. *Econometric Analysis of Cross Section and Panel Data*. Cambridge, MA: MIT Press.