WORKING PAPERS IN ECONOMICS AND STATISTICS

Report No 11 | 2024

Workplace Peer Effects in Turnout

Magnus Carlsson and Henning Finseraas

Linnæus University

Workplace Peer Effects in Turnout

Magnus Carlsson and Henning Finseraas

Workplace Peer Effects in Turnout

Magnus Carlsson and Henning Finseraas

Series editor: Thomas Giebe

ISBN: 978-91-8082-225-1 (pdf)
DOI: https://doi.org/10.15626/ns.wp.2024.11

Report No 11, Department of Economics and Statistics,
Linnaeus University, Växjö, 2024

Workplace Peer Effects in Turnout

Magnus Carlsson and Henning Finseraas*

October 3, 2024

Abstract

The potential for peer pressure at the workplace is high since social interactions are frequent and we care about our social standing at work. Peer effects in politics at the workplace are important to understand since workplaces are becoming more sorted according to human capital, which implies that workplace peer effects can increase social inequalities in turnout. To quantify peer effects we use population-wide administrative data from Sweden that covers several general elections and allows us to measure the turnout of colleagues. To identify peer effects we use the turnout of colleagues' family members in earlier elections as an instrumental variable, and leverage the richness of the data to assess assumptions, improve interpretation, and study heterogeneity. Our estimates suggest that workplace peer effects contribute to social inequality in turnout.

Keywords: Voter Turnout, Peer Effects, Social Networks, Workplace Dynamics

WORD COUNT: 8,604

*Magnus Carlsson: Linnaeus University, 391 82 Kalmar, Sweden, e-mail: magnus.carlsson@lnu.se. Henning Finseraas: Norwegian University of Science and Technology, P.box 8900 Torgarden, 7491 Trondheim, Norway, e-mail: henning.finseraas@ntnu.no. We thank Sven Hegewald, Filip Kostelka, Matthew Pietryka, and panelists at the EPSA Conference in Cologne, July 2024, and the annual meeting of the APSA, Phildelphia, Septemner 2024 for helpful comments.

Introduction

In Brady et al.'s (1995) influential model of voter turnout, people do not vote because "they can't", "they don't want to", or because "nobody asked". While "they can't" points to individual resources, the latter two refer to engagement and mobilization, factors with a social dimension. Engagement of citizens in politics can be stimulated by membership in politically active groups, and mobilization can take place through networks that encourage citizens to participate in politics. These factors underscore the role of social networks for voter participation.

The workplace constitutes one of the most important social networks we are part of. The majority of the electorate is employed by a firm, we spend many hours each day in contact with our colleagues, and we often care deeply about our social standing at work. Workplace peers can establish social norms about political participation, and peer pressure at work can motivate to vote. The workplace can serve as a political arena that provides political information and knowledge through informal conversations and discussions (Mutz and Mondak 2006). Workplaces can further provide mobilization efforts, e.g., through trade union activity (Leighley and Nagler 2007), and it can shape political preferences (Kitschelt and Rehm 2014).

However, the empirical literature on the role of workplaces and colleagues for voter turnout is surprisingly scarce. Mutz and Mondak (2006) pointed this out more than 20 years ago, but research has still been more concerned with the role of civic (Gerber, Gruber, and Hungerman 2016), neighbourhood (Bratsberg et al. 2021), elite (Pietryka and DeBats 2017), and family (Dahlgaard 2018; Dahlgaard et al. 2022) networks. The main reason for this, we believe, is that few data sets have the information that makes it possible to analyze the influence of the workplace and the political behaviour of colleagues.[1]

We leverage data that allow us to investigate coworker peer effects on voter turnout. Documenting workplace peer effects is not only important for understanding the deter-

[1] Aggeborn and Andersson (2022) and Cox, Fiva, and King (2024) are two notable exceptions. They study how political candidates mobilize turnout at their workplaces, which is a related but different question from ours. Another related paper is Alt et al. (2022), who use a similar identification logic as us to identify the effects of unemployment in one's social network on political attitudes and vote choice.

minants of voter turnout, but can also shed light on the development of political inequalities. Recent work on political inequalities has emphasised political mechanisms to explain inequalities, such as weak political representation (Heath 2018) or a mismatch between demand and supply of political platforms (Evans and Tilley 2017). Labour market developments might be equally important, however, since peer effects can exacerbate inequalities in turnout between demographic groups (DiMaggio and Garip 2012): If there are important peer effects at the workplace, they will exacerbate inequalities if workplace networks are homophilous to individual characteristics. Håkanson, Lindqvist, and Vlachos (2021) show that Swedish firms are becoming more strongly sorted on skills and other characteristics, i.e. colleagues in the same firm are more similar to each other today than in the past. This trend is mainly driven by exogenous changes in the labour market, which means that if it contributes to political inequality, it will be a harder social issue to solve.

To identify coworker peer effects, we rely on an IV approach. The problems of disentangling peer influence from other factors that can produce correlations between the behaviour of group members is well-understood (Manski 1993). In the absence of experimental designs, IV methods are frequently used to estimate causal peer effects. We follow several previous papers (e.g. Nicoletti, Salvanes, and Tominey 2018; Carlsson and Reshid 2022; Buggle et al. 2023) and use the so-called peers of peers approach (Bramoullé, Djebbari, and Fortin 2009) to construct an instrument for turnout among colleagues. The idea is to identify social networks that influence colleagues, but which do not overlap with the social networks of the respective worker, and therefore do not influence the worker. If these networks are exogenous to the worker and have sufficient influence on colleagues, we have a promising candidate for an instrument for workplace turnout.

In our application, we use the share of coworkers' family members who voted in a previous election as an instrument for the average turnout among coworkers. The key assumptions required for a causal interpretation of our estimates are that the influence runs solely from peers of peers (no reflection) and that the peers of peers have no direct influence on the worker (the exclusion restriction). Below we describe in detail how we implement our identification strategy and conduct several tests of the key assumptions.

We use administrative micro-data from Statistics Sweden which spans from 2009 to 2022 and covers the full population of individuals, firms and workplaces in Sweden. The most important features of the data are that they contain information on turnout at the individual level for several elections, firms and workplaces, and family members. The data enable us to identify coworkers and family members of all workers from which we can calculate the explanatory variable (turnout among coworkers) and the instrument (turnout among family members of coworkers).

We find evidence of a positive coworker peer effect. Our baseline estimate reveals that a ten percentage points higher turnout among coworkers causes a 1.5 percentage point increase in the probability of voting, which corresponds to a 1.7 percent increase in turnout compared to the average turnout in the sample. However, we find significant and politically meaningful heterogeneous effects. We find stronger peer effects among individuals with characteristics that predict lower turnout (previous abstention, low turnout among parents, males, low education, unmarried, low income). Stronger peer effects among workers that are less likely to vote imply that turnout differences across demographic groups can be reduced if groups with high and low turnout interact. However, the recent development of an increasing sorting of workers with similar characteristics into firms (Håkanson, Lindqvist, and Vlachos 2021) reduce such interactions and instead lead to a preservation and strengthening of inequalities in turnout.

The rest of the paper is structured as follows. Next, we discuss the literature on peer effects on turnout and argue that colleagues might influence propensity to vote. The following two sections describe the research design and the data, before we discuss and assess the IV assumptions. Then we present the empirical results, with a particular focus on heterogeneity, before the final section concludes.

Peer effects on turnout

The rich literature on why people vote has identified many relevant factors at the individual level. These include the importance of individual resources (Brady, Verba, and

Schlozman 1995), psychological dispositions (Gerber, Green, and Shachar 2003), genetics (Fowler and Dawes 2008), and rational cost-benefit calculations (Downs 1957). While the literature for long also has acknowledged the role of various social dimensions of voting as important – see Lazarsfeld, Berelson, and Gaudet (1944) for a classic study, and more recently Rolfe (2012) and Sinclair (2012) – its impact is still much less understood (Campbell 2013). Most work on social aspects of voting has studied the role of social connections and networks that are explicitly political, such as mobilization through political campaigns (Gerber and Green 2000) or how political candidates leverage their social networks (Cox, Fiva, and King 2024).

The social aspects of turnout are harder to study empirically, which is probably why this literature is less developed. This is particularly true for the effects of daily social interactions. It is highly plausible that the behaviour of people in our social networks influences our behaviour ("peer effects"), but, as outlined by Manski (1993), correlated behaviour might be due to selection to networks (we prefer to interact with people similar to us) and reflection (influence can go both ways). Moreover, we often lack data on social connections and the behaviour of peers.

Still, important work over the last two decades has improved our understanding of peer effects on turnout. Causal peer effects in turnout have been identified within households (Nickerson 2008), neighbourhoods (Finan, Seira, and Simpser 2021), and across Facebook friends (Bond et al. 2012). Typically, two mechanisms are proposed for explaining these peer effects. The first is an information and learning mechanism: Through our social networks we can learn about, e.g., the process of voting, the norms of voting, the platforms of political parties, candidate characteristics, and the potential coalition partners after the election, all of which can mobilize us to vote. The second mechanism is through social norms and peer pressure. Previous work shows or suggests that social pressure from neighbours (Gerber, Green, and Larimer 2008), community members (Funk 2010), and friends (Bond et al. 2012) can influence the propensity to vote, in particular if there is a risk that others can reveal that you did not comply with the social norm of voting (Gerber, Green, and Larimer 2008; Green and Gerber 2010).

Peers at the workplace can influence turnout through both learning and social norms. By engaging in conversations about politics with our colleagues, the learning and information mechanism will be in play. This is particularly so since we are more likely to meet people with different political views at the workplace than in more intimate networks, which means that we might get new political information and knowledge through these discussions (Mutz and Mondak 2006). The information and learning effect is likely to be stronger if a higher share of the colleagues vote. Social norms and pressure are also likely to be in play when more of the colleagues vote. While colleagues cannot monitor whether other colleagues turn out to vote, abstention might be revealed during casual conversations, thereby causing social embarrassment, which will uphold the social norm.

We should point out, however, that there might also be costs of engagement in or exposure to political conversations at the workplace. For instance, employee well-being might be negatively affected by learning that one has a minority view on contentious political issues (Rosen et al. 2024), and one might learn biased information (Carlson 2019). Perhaps because of these costs, some avoid political discussions altogether (Krupnikov and Ryan 2022), which will limit the peer effect. Although Swedish politics is less polarized than the US, these countervailing mechanisms might operate there too.

Peer effects in turnout are important to understand because they might explain current social inequalities in turnout (DiMaggio and Garip 2012). Swedish firms in all major industries have become more homogeneous in their cognitive and non-cognitive characteristics over the last decades (Håkanson, Lindqvist, and Vlachos 2021). Turnout is positively correlated with such skills, which means that workplace peer effects are likely to amplify social inequalities in turnout. Stronger sorting is an important driver of wage inequality in Sweden (Håkanson, Lindqvist, and Vlachos 2021), which means workplace peer effects also provide a possible mechanism for the correlation between wage inequality and turnout (Solt 2008). Since stronger sorting reflects trends in labour market restructuring, it means that it is difficult to envision social and political interventions that can easily reduce its negative effect on political inequality (which one might perceive as a negative externality of restructuring).

Research Design

Equation (1) provides the starting point for estimating the coworker peer effect in voting:

$$Voted_{i,2022} = \beta \overline{Voted}_{-i,2018} + \epsilon_i \tag{1}$$

In equation (1), subscript i refers to the workers for whom we estimate whether they are influenced by turnout among their coworker peers. Henceforth, we label these workers *index workers*. The variable $Voted_{i,2022}$ measures whether index worker i voted in the general election in 2022, which is the outcome. The explanatory variable of interest is $\overline{Voted}_{-i,2018} = \frac{\sum_{j \in P_i} Voted_{j,2018}}{n_i}$, which measures turnout in 2018 among index worker i's coworkers, excluding the index worker, indicated by notation $-i$. This exclusion means that the explanatory variable of interest generally takes different values for index workers at the same workplace.[2] In the expression for $\overline{Voted}_{-i,2018}$, P_i denotes the set of i's workplace peers, n_i is the number of coworker peers, and j indexes the coworker peers. The coefficient of main interest is β, which measures the coworker peer effect.

There are several endogeneity problems that most likely will result in a biased estimate of β in equation (1). One endogeneity problem is the reflection problem (Manski 1993; Moffitt 2001) we mentioned above. In our case, the reflection problem means that two workers at the same workplace may simultaneously affect the voting behavior of each other, which makes it hard to determine if the first worker impacts the second or vice versa. A second endogeneity problem is if workers self-select into workplaces. Certain types of workplaces may attract workers who share similar characteristics, and this could make it challenging to distinguish workplace peer effects from correlations that reflect self-selection into firms. One could also imagine endogeneity problems because individuals at a workplace share the same local labor market, as contextual factors may impact voting behavior. For instance, local changes in the the risk of unemployment will be correlated at the workplace, and can influence political behaviour (Alt et al. 2022; Finseraas, Røed,

[2] It will take two different values if all workers at the workplace belong to the same peer group, but as we explain below, we use worker characteristics to construct homogenous peer groups, which means that there could be more than two values at each workplace.

and Schøne 2017).

We address these endogeneity problems with an IV approach that is sometimes labeled the peers of peers approach in the literature. That is, for each index worker, we use the share of coworkers' family members (siblings and cousins) that voted in the general election in 2010 ($\overline{Voted}_{F,-i,2010}$; F=family) as an instrument for the average turnout of the coworkers in 2018 (i.e.,$\overline{Voted}_{-i,2018}$). Specifically, the instrument is obtained by first calculating the fraction that voted among siblings and cousins of each coworker and then taking the workplace average of these numbers. In the calculation of the instrument for an index worker, the family members of the index worker herself are excluded. Again, the index worker exclusion implies that the instrument generally takes different values for index workers at the same workplace.

The IV approach relies on the existence of peers belonging to different networks that are not related in a direct way (Bramoullé, Djebbari, and Fortin 2009; Carlsson and Reshid 2022; De Giorgi, Pellizzari, and Redaelli 2010; Nicoletti, Salvanes, and Tominey 2018). Thus, we assume that the naturally occurring social networks at workplaces and within families are distinct. The main identifying assumption for the IV to be valid is that the index worker must not interact directly with the family members of their coworkers. This exclusion restriction means that the impact of the voting behavior of a family member of a coworker on the index worker's turnout must operate only via the coworker. In other words, if we assume that individuals 1 and 2 are coworker peers and individuals 2 and 3 are cousins, then the assumption is that individual 3 has no direct influence on individual 1. Instead, any influence must operate indirectly through its effect on individual 2, who is a coworker of individual 1 and a cousin of individual 3.

Following previous research on peer effects (e.g. Cornelissen, Dustmann, and Schönberg 2017), we adopt the concept of homogeneous peer groups (HPGs). This means that peers should not only belong to the same workplace or family but should also share similar characteristics such as gender and educational level, which we use to construct HPGs. Thus, some workplaces will have several HPGs, while others will have only one. The rationale for using HPGs is the notion that social interactions are more likely to occur between individuals with similar characteristics. Peer effects are thus more plausible within HPGs.

We address the reflection problem by constructing the explanatory variable and the IV based on the voting behavior of coworkers and their family peers in prior elections. Specifically, we measure turnout of index workers in 2022, turnout of the coworkers in 2018, and turnout of the coworkers' family members in 2010.[3] The temporal order of the variables reduces the most obvious reflection problems, but there could still be a reflection problem for index workers who were entitled to vote in elections prior to 2010. For example, if these index workers influenced coworkers in the 2002 election, which in turn influenced their family members in the 2006 election. To address this, we control for whether the an index worker voted in the 2010 general election (and whether the person was below 18 years old in 2010 and thus ineligible to vote).

This setup using three elections and our definition of homogeneous coworkers means that, e.g., a highly educated male index worker with observed turnout in 2022 is paired with i) highly educated male coworkers at the same workplace with turnout measured in 2018 and ii) highly educated male family members (siblings and cousins) of these coworkers with turnout measured in 2010. Thus, there are four cells of individuals at workplaces and in families defined by the two dimensions education and gender.

A remaining endogeneity problem is that even if there are no social interactions between an index worker and the coworkers' family members, we cannot rule out common factors in the family peer groups to which index workers and coworkers belong. For example, index workers and the family members of coworkers could live in similar neighbourhoods and/or share similar background characteristics. If these factors also affect turnout and are not accounted for, the ignorability assumption is violated, which will bias the first stage and the reduced form estimates. Moreover, the exclusion restriction is violated if the effect of the instrument runs through such variables.

We account for such common factors in several ways. First, the included control variable for whether the index worker voted in the 2010 general election, i.e., lagged

[3]Turnout from the 2014 election is not digitized and not available from Statistics Sweden.

turnout, will control for many unobservable factors of the index worker that is shared by the coworkers' families. Second, whenever we use the IV, we include as a control variable the average turnout of the index worker's family members (excluding the index worker), akin to a family fixed effect. We call this variable the "individual IV" (Nicoletti, Salvanes, and Tominey 2018) because it is closely related to the instrument. Controlling for the individual IV will reduce potential biases caused by unobserved characteristics at the family level.[4] Third, we add electoral district fixed effects for the 2018 electoral district of the index worker. There are approximately 6,000 electoral districts in Sweden, which typically correspond to small homogeneous neighborhoods. The electoral district fixed effects address potential bias from selection of index workers and coworker' family members into the same neighborhood. Fourth, using the year of the instrument (2010), we control for average characteristics of the coworker's family members that correlates with turnout (age, income, marital status, number of children, and employment status). Finally, because there are systematic differences in turnout between the HPGs we include fixed effects for which type of HPG an index worker belongs to.

We do the estimations using two-stage least squares (2SLS), and the second-stage IV-regression takes the following form:

$$\begin{aligned} Voted_{i,2022} = \beta\widehat{\overline{Voted}}_{-i,2018} + \gamma\overline{Voted}_{F,-i,2010} + \delta Voted_{i,2010} + \\ + \boldsymbol{\theta}\mathbf{ElecDistr}_{i,2018} + \boldsymbol{\mu}\overline{\mathbf{CoworkerFamilyX}}_{i,2010} + \boldsymbol{\rho}\mathbf{TypeOfHPG}_{i,2018} + \epsilon_i \end{aligned} \tag{2}$$

The outcome $Voted_{i,2022}$ is a binary indicator of whether the index worker voted in 2022. The explanatory variable of main interest is $\widehat{\overline{Voted}}_{-i,2018}$, which is the predicted value of 2018 turnout among coworkers (excluding the index worker) obtained from the first-stage regression, and $\overline{Voted}_{F,-i,2010}$ is the individual IV. The remaining variables control for the issues discussed above, i.e., potential common factors in the family peer groups to which index workers and coworkers belong: $Voted_{i,2010}$ is the turnout of the index worker in 2010, $\mathbf{ElecDistr}_{i,2018}$ are fixed effect for electoral district of the index

[4]It will also reduce the so-called exclusion bias in studies of peer effects (Caeyers and Fafchamps 2023). The exclusion bias is caused by a correlation between the turnout of the index worker and the average of the peer group in small peer groups, which is created mechanically when the index worker is excluded in the calculation of the average of the peer group.

worker, $\overline{\textbf{CoworkerFamilyX}}_{i,2010}$ is a vector of the means of coworkers' family average characteristics, and $\textbf{TypeOfHPG}_{i,2018}$ are four dummies for the type of peer group.

Data and sample

Population-wide micro-data

We use micro-data from Statistics Sweden's registers, which span the years 2009-2022, covering the full population of individuals above 16 years old and all firms and workplaces in Sweden. The most important features of the registers are that they i) contain information on turnout at the individual level for the full population for several elections, ii) contain identifiers of individuals, parents and children, firms, and workplaces, so they can be linked together and across years. Thus, we can link individuals to workplaces, identify the index workers' coworkers, identify the family members of index workers and coworkers, and observe turnout and other characteristics of each individual. In the appendix, we provide additional details of the main registers we use.

Constructing the sample

To implement the peers of peers method, we need to identify coworkers and family members (siblings and cousins) of each index coworker. The former is required to construct the explanatory variable and the latter to construct the instrument. We identify coworkers by identifying workers who share the same unique workplace identifier. To identify the family members of coworkers, we use the family identifiers to construct a dataset representing a family tree for the full population of individuals. From the family tree we can identify siblings and cousins of each individual, defined by sharing parents or grandparents.[5]

The nature of workplace peer effects on turnout, as well as data limitations, motivate a few restrictions on the sample. We make all sample restrictions based on predetermined variables (typically measured in 2017, i.e., before the 2018 election when treatment takes place), to avoid post-treatment bias. The most obvious restriction is excluding individuals who are not employed. We exclude those without a workplace in 2017, and those who have

[5] Thus, we require that siblings share at least one parent and cousins at least one grandparent.

a very low income in 2017, which typically means that they work few hours and thus have a weak connection to the workplace. Next, we exclude workers for whom the individual IV is unobserved, i.e., workers who do not have any siblings or cousins. This restriction means that many immigrants are excluded since the individual IV is often unobserved for immigrants (because there is no information on parents and grandparents in the registers). To keep the sample homogeneous we therefore exclude all first-generation immigrants.

Moreover, we exclude workers who have a sibling or cousin at the workplace. Including these workers would violate the identifying assumption of non-overlapping networks. Next, we exclude workplaces with only one worker, in which case there cannot be any peer influence. We also exclude workplaces with more than 50 coworkers in 2017. We do this for two reasons: i) The size of the relational dataset between index workers, coworkers, and family members grows exponentially with the size of the workplace, making the dataset unmanageable. ii) Co-worker turnout will to a lesser extent reflect political interactions at large workplaces. Since the upper limit of 50 workers is somewhat arbitrary we examine robustness with different thresholds in the empirical analysis. Finally, we limit the sample to workers who are between 18 and 40 years old in 2018 when treatment takes place. The upper age restriction is motivated by previous literature which shows that voter habits are strongly correlated with age (Cravens 2020).[6] The final sample consists of 390,317 index workers at 84,666 workplaces.

Descriptive statistics

Summary statistics for turnout and the demographic variables of the index workers can be found in Table A1. The most important to notice is that turnout is very high in our sample, at about 90 percent.[7] This reflect the general high level of turnout in Sweden, but also that we condition on employment in 2017 and exclude immigrants. As explained, we need to make these sample restrictions to get internally valid estimates, but it means

[6] The lower age restriction implies that voters below 26 years of age in 2018 were ineligible to vote in 2010 and are represented by a missing dummy for $Voted_{i,2010}$. Table A4 shows that we get similar results if we exclude these workers from the analysis.

[7] One might worry that using a linear probability model on such a skewed dependent variable produces biased estimates. However, we get almost identical marginal effects as those reported below if we use a probit specification.

that external validity might be limited to other high turnout samples.

Table A2 documents how turnout is influenced by previous voting and demographic variables. The estimates reveal that workers who voted in 2010, have parents that voted in 2010, are females, are younger than median, have college education, are married, have above median income are more likely to vote. These patterns coincide with what is commonly found in the literature (Smets and Van Ham 2013). As it turns out, these patterns are important when we later interpret our main findings.

Assessing the IV assumptions

We need several assumptions to be satisfied for our estimate of the parameter of interest in equation (2) to reflect the unbiased causal peer effect (see, e.g., Angrist and Pischke 2009). First, the instrument (turnout among coworkers' families) must be correlated with the treatment variable (turnout among coworkers), i.e., there must be a first-stage, and to avoid bias in the second-stage estimate, the first-stage correlation needs to be strong (the *relevance* assumption). Second, the *monotonicity* assumption must be satisfied, which states that a change in the instrument either has no impact on the treatment variable or changes it in the same direction for all index workers. Third, the instrument should only affect the outcome (turnout of index workers) through the instrument (*exclusion restriction*). Fourth, the instrument and the outcome should not be jointly correlated with unobserved variables (*ignorability assumption*).[8] In our case, these assumptions are conditional on the included covariates. We conduct several direct and indirect tests to assess whether the assumptions are plausible.

First-stage and monotonicity

The first-stage assumption requires that the turnout among coworkers' family members can predict turnout among coworkers. We test this by regressing turnout among coworkers on turnout among the coworkers' family members. Moreover, if monotonicity holds, a

[8]One additional assumption is the *stable unit treatment value assumption*, which says that index workers' potential outcomes are unrelated to the co-worker turnout of other index workers. This assumption will be violated if peers-of-peers' networks overlap. We cannot test this assumption, but as explained above we have made sample restrictions to make it more plausible.

testable implication is that the first-stage should not be negative for any subgroup of index workers. Following Bhuller et al. (2020), we examine this implication by estimating the first-stage in sample splits - in our case for gender and education, which are the variables defining the HPGs. We find positive and significant first-stage estimates for all subgroups, also for the sample of college-educated where almost everyone votes (see Table A3).

Exogeneity tests using predetermined characteristics

While the exclusion restriction and ignorability assumptions cannot be directly tested, it is possible to conduct indirect tests. For example, if these assumptions hold, we should not find that the β parameter is significant if we re-estimate equation (2) but replace the dependent variable with different predetermined characteristics of the index workers. For this exercise, we use such characteristics which are often included in models of turnout (Smets and Van Ham 2013) – age, age squared, income, marital status, and number of children – but not used to construct the HPGs (i.e., education and gender). The estimates are presented in Table 1. Reassuringly, for all outcomes, we find coefficient estimates that are trivial in size and statistically insignificant.

Table 1: Exogeneity tests on predetermined individual characteristics.

	Age	Age-squared	Income	Married	Number of children
WP Turnout	0.236	16.638	203.429	0.036	-0.024
	(0.316)	(15.226)	(112.813)	(0.029)	(0.079)
Controls	Yes	Yes	Yes	Yes	Yes
HPG FE	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	1487	1487	1487	1487	1487
N	390970	390970	390970	390970	390970
Mean of Y	22.35	531.40	3179.73	0.20	0.95

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** p<.01; ** p<.05.

Results

Main estimates of the coworker peer effect

Table 2 presents the main estimates of the coworker peer effect, which are visualized in Figure 1. The dependent variable is turnout measured in 2022 and the explanatory variable is turnout among coworkers in 2018. The first column presents the OLS estimate of the peer effect. The coefficient estimate of 0.047 means that ten percentage points difference in the turnout among coworkers increases the likelihood that an index worker votes by approximately half a percentage point. Since the standard deviation of turnout among coworkers is 20 percentage points, this means that a standard deviation increase in turnout among coworkers increases the probability of voting of an index worker by one percentage point. However, the OLS estimate is likely to be biased for the reasons discussed above. In particular, we suspect that reflection and selection effects will cause a positive bias in the estimate.

Table 2: Main results.

	OLS	First stage	Reduced form	Second stage
WP Turnout	0.047***			0.152***
	(0.003)			(0.025)
Instrument		0.145***	0.022***	
		(0.004)	(0.004)	
Controls	Yes	Yes	Yes	Yes
HPG FE	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Kleibergen-Paap F				1487
N	390970	390970	390970	390970
Mean of Y	0.90	0.91	0.90	0.90

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** p<.01; ** p<.05.

Next, we turn to the estimate of the peer effect using the IV approach explained above and start by again verifying the existence of a strong first-stage and showing the reduced form estimate. The second column in the table shows that a ten percentage point difference in turnout among coworkers' family members is associated with 1.45 percentage points higher turnout among coworkers. This correlation is precisely estimated, which means that weak instrument bias in the second stage is not an issue.[9] Column 3 presents the reduced form estimate, i.e., the impact of turnout among coworkers' family members on turnout among index workers, which is positive and highly significant. The magnitude of the coefficient is moderate in size: A ten percentage points higher turnout among coworkers' family members twelve years prior leading to 0.22 percentage points higher turnout among index workers.

Figure 1: Plots visualizing the main results.

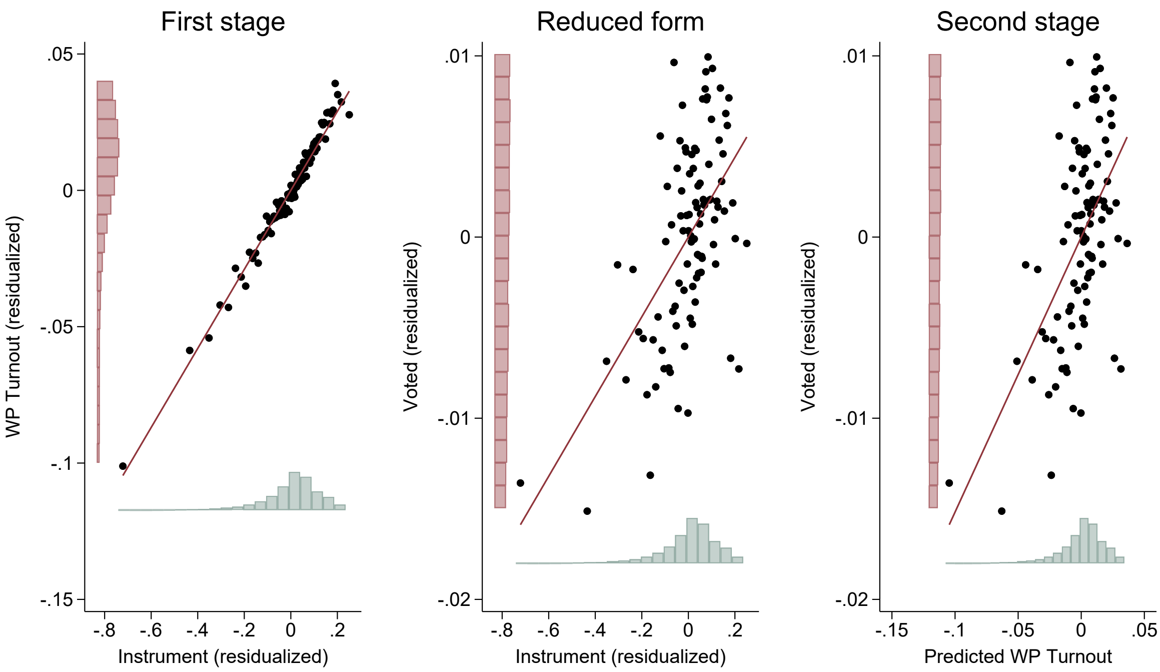

Note: The figure displays the variation driving the estimates in Table 2. Variables on the Y and X axes are residualized using the covariates in Equation . Each dot represents one percent of the observations, while the regression line is based on all observations. Bars on the axes are histograms of the respective variables.

Finally, in column 4 we present the main IV estimate corresponding to parameter β in equation (2) above. The IV assumptions imply that the reduced form estimate solely reflects the influence of coworkers, which means that the second stage estimate is the ratio of the reduced form and the first stage. The second stage estimate of 0.152 is thus

[9]The Anderson-Rubin 95% confidence interval for the second stage estimate is very close to the conventional confidence interval, which indicates no weak instrument bias (see e.g. Lal et al. 2024).

our estimate of the causal coworker peer effect: A ten percentage points higher turnout among coworkers causes an increase in the probability of voting among index workers of 1.52 percentage points. In the Appendix we show that this estimate is robust to adding various types of additional control variables (see Table A6).

To get a better sense of the importance and plausibility of the estimate, it is informative to compare it to other estimates of peer effects. The most straightforward comparison to make is to the neighbourhood peer effect in Finan et al. (2021). They find that a one standard deviation difference in neighbourhood turnout is associated with a peer effect in turnout of about 1 percentage point. In our case, a standard deviation difference in workplace turnout (0.2) amounts to a peer effect that is about three times larger ($.152 * .2$). Since most of us have more interaction with colleagues than neighbours, we believe that a larger workplace peer effect is plausible. Moreover, Nickerson (2008) estimates a within-family peer effect of about six percentage points. It is arguably less straightforward to compare this estimate to ours, but adding one additional voter to ones' workplace peer group will in most cases have a much smaller effect than this. We only get close to a six percentage point peer effect if the peer group goes from consisting of two non-voting colleagues to three colleagues and the additional colleague is voting (e.g. $.152 * .33 = .05$). We find it plausible that the workplace peer effect is larger than the neighbourhood peer effect, but smaller than the within-family peer effect. Still, valid concerns can be raised against the estimate, which is what we turn to next.

Comparing the OLS and IV estimates

The causal estimate is more than three times larger than the OLS estimate (0.152 and 0.047), which is somewhat surprising since we argue above that the OLS estimate is likely to be positively biased. In this section, we explore factors that can result in a larger IV than OLS estimate (see e.g. Lal et al. 2024).

First, our second stage estimate is the local average treatment effect (LATE) for the complier group. With heterogeneous treatment effects, LATE might differ from the average treatment effect (ATE) in the population. OLS estimates the ATE, but probably a biased estimate due to the endogeneity problems discussed above. Thus, the peer effect we estimate may be larger than the OLS estimate because of heterogeneous treatment effects.

Second, and related, the heterogeneity causing a larger IV estimate might be because the treatment effect varies according to the (unobserved) underlying probability of being treated (resistance to treatment). The LATE will differ from the ATE if those with high treatment effects have more weight in the LATE estimate.

Third, IV can address attenuation bias, due to noise in or miss-classification of the explanatory variable, which will result in a downward bias using OLS.

Finally, while the three explanations just discussed all assume a valid instrument, the IV estimate will also be biased if the instrument is invalid. We addressed some of the IV assumptions above, such as weak instrument bias and monotonicity, and below we conduct a set of placebo exercises (see section 6.3) to further examine instrument validity.

Heterogeneous effects based on observables

We first examine treatment heterogeneity on observables. If the peer effect is stronger among index workers with certain characteristics, and those in the complier group have these characteristics to a larger extent, then the LATE will be larger than the OLS estimate (see e.g. Angrist and Pischke 2009). To investigate this, we first explore heterogeneous effects with respect to observed characteristics of the index workers and then profile the complier group with respect to these characteristics.

We estimate heterogeneous peer effects for index workers with different observed characteristics by adding interaction variables between the treatment variable and the respective characteristics. For example, for gender we multiply a female and male dummy indicator with turnout among coworkers and add the resulting interaction variables and the female dummy to the regression. We leave out the variable measuring turnout among coworkers, which eases comparing the estimates within subgroups with the overall estimate: the main coefficient in Table 2 is then the (weighted) average of the coefficients of the two interaction effects. We use the same approach for the other characteristics of the index workers. In this analysis, we examine the predetermined variables of the index workers (age, income, married, and number of children, education, gender, and turnout in 2010) and parental turnout in 2010. The characteristics are analyzed using separate regressions.

The estimates of heterogeneous effects for the various subgroups are presented in Table 3. We find strong heterogeneous effects with a consistent pattern such that the peer effect is stronger in subgroups with lower turnout. Specifically, the peer effect is stronger among those index workers who did not vote in the 2010 general election, have parents that did not vote in the 2010 election, are male, younger, do not have a college degree, are not married, and have income below the median.

Table 3: Heterogeneous effects.

WP Turnout	
A) Index worker turnout 2010	
x Voted 2010	0.018
	(0.026)
x Not Voted 2010	0.271***
	(0.031)
B) Parental turnout 2010	
x Parents voted 2010	0.139***
	(0.025)
x Not parents voted 2010	0.400***
	(0.055)
C) Gender	
x Female	0.138***
	(0.042)
x Male	0.203***
	(0.029)
D) Age	
x Age above mean	0.134***
	(0.029)
x Age below mean	0.213***
	(0.028)
E) Education	
x College	0.049
	(0.171)
x Not College	0.185***
	(0.025)
F) Marital status	
x Married	0.055
	(0.036)
x Not Married	0.206***
	(0.026)
G) Income	
x Income above mean	0.036
	(0.030)
x Income below mean	0.268***
	(0.027)

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** p<.01; ** p<.05.

The stronger peer effect among workers with lower propensities to vote, means that

Figure 2: Predicted turnout across workplace turnout by parents' turnout.

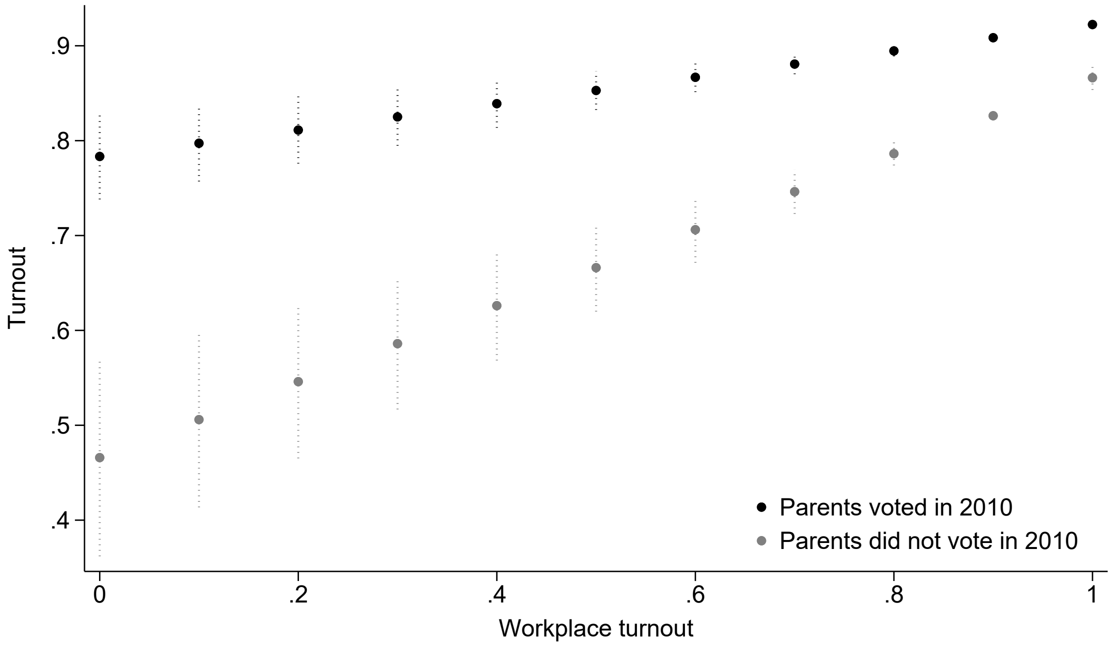

Note: Predictions are based on second-stage estimates from Table 3. Workplace turnout is the predicted workplace turnout based on the first stage. Bullets are point estimates while dotted lines refer to 95 percent confidence intervals.

the increasing sorting of workers across firms (Håkanson, Lindqvist, and Vlachos 2021) might increase social inequalities in turnout. Figure 2 illustrates why: The figure is based on the second stage estimates in Panel B, Table 3, and shows the predicted level of turnout for workers with parents that did not vote in 2010 (i.e. workers with a lower propensity to vote) versus parents that did vote in 2010 (higher propensity to vote) across workplace turnout (predicted from the first stage). Stronger sorting of workers across workplaces implies a movement upward on the curve for high propensity voters and a movement downwards for low propensity voters, thereby increasing the turnout gap between these two groups.

To examine if the heterogeneous effects can help explaining why the IV estimate is larger than the OLS estimate, the next question is if the index workers in the complier group to a larger extent have characteristics associated with stronger peer effects. To examine this, we profile the characteristics of the compliers (Marbach and Hangartner 2020). It is not straightforward to profile the compliers since we have treatment and instrument variables that are continuous, and we need to condition on covariates. Therefore, to facilitate profiling, we discretize residualized versions of the treatment and instrument variables: We run regressions with the continuous treatment and the instrument as dependent variables with the full set of covariates as the independent variables, obtain the residuals, and then discretize them around their medians. With these variables we can profile the compliers using the approach in Marbach and Hangartner (2020). We acknowledge that the compliers when using these variables might be different from the compliers in Table 2. However, the ratio of the OLS to the IV estimate is similar when we analyze these variables, see Table A5, which suggests that the complier groups are similar.

Table 4 presents the means of individual-level characteristics for the whole sample and the compliers. Most notably, we find that the compliers have lower turnout in previous elections and lower level of education. Both these characteristics are associated with stronger peer effects (see Table 3). The compliers also tend to have the other characteristics associated with stronger peer effects. While this will contribute to the IV estimate being larger than the OLS, the differences in complier characteristics are not dramatic, which means that it is unlikely that these differerences explain the whole difference between the two estimates.

Table 4: Complier characteristics.

	Mean whole sample	Mean compliers
Voted 2010	0.61	0.56
Voted 2018	0.91	0.89
Parents voted 2010	0.93	0.92
Female	0.43	0.41
Age	34.35	33.92
College education	0.31	0.25
Married	0.20	0.19
Income	3179	3184

Note: Mean characteristics in the total sample and for compliers.

Marginal treatment effects

Next, we explore heterogeneity in the treatment effect along the unobserved gain from treatment (Heckman and Vytlacil 2007; Brinch, Mogstad, and Wiswall 2017), which can explain why the IV is larger than the OLS estimate. To do so we estimate Marginal treatment effects (MTE:s), which will differ from the average treatment effect if the impact of treatment varies across workers based on unobserved traits. A concrete example would be a HPG of workers who are less interested in politics and therefore possibly more sensitive to peer mobilization (Rolfe 2012). As a result, average turnout in this example HPG (i.e, the treatment) is expected to more strongly depend on turnout among their family members (i.e., the instrument). If so, index workers in this HPG will be over-represented in the complier group. But we would also expect a stronger response to treatment for index workers in this HPG, because they are more sensitive to peer mobilization. This illustrates a situation where the magnitude of the treatment effect is expected to be larger among compliers, i.e., a heterogeneous treatment effect based on the unobserved gain from treatment (co-worker turnout).

In the supplementary appendix we present the details and results of the MTE analysis. In our example with interest in politics as the unobserved trait, we shuld expect a positive slope of the MTE curve: Those with the lowest political interest should have larger treatment effects. However, this is not what we find, as the slope is negative but never significant. Thus, we find no evidence of differences in the MTE along the unobserved dimension, which makes this an unlikely candidate to explain the larger IV estimate compared to the OLS estimate.

Attenuation bias

While we have high-quality measures of voter turnout, we lack information on the mechanisms we propose as drivers of the peer effects. Ideally, we would like to have measures of e.g. the degree of social interactions with colleagues and the extent to which political topics are discussed during breaks. The implication is that although coworkers' average turnout will be strongly related to e.g. the extent of political discussions, there might be measurement error in the mapping between the two concepts (see online appendix for further discussion).

The presence of noise in the explanatory variable tends to bias the OLS estimate toward zero. In contrast, the IV estimate will contain less bias and tend to be larger than the OLS estimate, and the difference will grow with the extent to which the instrument contains less noise. In the online appendix we present an analysis of to what extent this bias correction can explain that the IV estimate is larger than the OLS estimate. We do so by comparing the difference between the OLS and IV estimates across smaller and larger peer groups. We find that the differences between the OLS and IV estimates are larger for smaller peer groups, which indicates that attenuation bias contributes to the larger IV estimates, since smaller peer groups have a more noisy estimate of the peer effect.

Placebo tests

Finally, we conduct a series of placebo, or negative control, tests (Eggers, Tuñón, and Dafoe 2023; Shi, Miao, and Tchetgen 2020) to further examine the validity of our identification strategy. These tests involve the construction of either negative control (placebo) exposures or negative control (placebo) outcomes. We conduct both types of tests.

Negative control exposure tests involve replacing the true exposure variable with a placebo. The negative control exposure variable should be theoretically unrelated to the treatment (coworker's turnout), but susceptible to similar biases. If the negative control exposure does not significantly influence the outcome, it strengthens the conclusion that the observed effect is likely due to the treatment itself, and not due to unobserved confounding factors.

We conduct four different negative control exposures, labeled A to D, in which we re-estimate the coworker peer effect after having randomly assigned the index worker to another HPG - a placebo HPG with workers that are assumed to have no contact with the index worker. The primary aim of this exercise is to demonstrate that the observed peer effect is unlikely to be influenced by unobserved factors unrelated to social interactions.

Specifically, if we detect a significant effect when using placebo coworkers from a placebo HPG, it suggests that unobserved factors may be driving the estimated effect.

In placebo test A, we randomly assign each index worker to any other HPG in the sample. Thus, the placebo coworkers share the same gender and education as the index worker, but are not expected to have any social interactions with the index worker. Then we re-estimate the baseline regression using the turnout among these placebo coworkers as a negative control exposure. Reassuringly, the association between the negative control exposure and turnout of the index worker is small and insignificant, see column 1 in Table 5.

We follow the same procedure for placebo test B, but with the added condition that the actual HPG and the placebo HPG must be located in the same municipality (column 2). In placebo test C (column 3), we maintain the municipality criterion but also require that the actual HPG and the placebo HPG belong to firms in the same industry (at the 1-digit level). Finally, in placebo test D (column 4), we require that the placebo HPG is in the same industry and of similar size as the original HPG (above or below the median HPG size). Across the columns, the placebo peer effects are much smaller than the true peer effect and always insignificant.

Table 5: Results from placebo regressions.

	Placebo A	Placebo B	Placebo C	Placebo D	Turnout in 2009
WP Turnout	0.014	-0.035	0.027	-0.011	0.071
	(0.022)	(0.023)	(0.025)	(0.027)	(0.046)
Controls	Yes	Yes	Yes	Yes	Yes
HPG FE	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	837	769	702	753	982
N	390958	387412	357051	338762	255076
Mean of Y	0.90	0.90	0.90	0.90	0.39

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** p<.01; ** p<.05.

Next we turn to the negative control outcome method, which implies using a different outcome variable that cannot have been influenced by the treatment but is susceptible to the same confounding factors. In our application, a good candidate for a negative control outcome would be turnout prior to the 2010-2022 period. While we do not have turnout for Swedish parliamentary elections before 2010, we have information on turnout in the 2009 European election. In the last column of Table 5, we repeat the main regression except that we use turnout in the 2009 European election instead of turnout in the 2022 national election. The estimated coefficient is not statistically significant and much smaller than our main estimate. Still, the estimate is not zero, which suggests that the correlation between the instrument and the error term in equation might not be exactly zero, as assumed by the exclusion restriction. We examine sensitivity to the violation of the exclusion restriction next.

Sensitivity to violation of the exclusion restriction

Conley, Hansen, and Rossi (2012) suggest that when valid concerns can be raised against the exclusion restriction one should conduct a sensitivity check of the consequences of violations. This involves making assumptions about how severely the exclusion restriction might be violated, for instance by specifying a range of plausible correlations between the instrument and the error term.

We follow Azar, Marinescu, and Steinbaum (2022) and use the reduced form estimate (see Table 2) as the upper range of the bias (i.e. the direct effect of the instrument is completely explained by an unobserved channel) and 0 bias (i.e. the exclusion restriction holds perfectly). Using Clark and Matta's (2018) implementation of Conley et al.'s (2012) method, we estimate bounds on the treatment effect under this range of potential violations. The bounds refer to the lowest and highest estimate of the treatment effect that all the respective confidence intervals contain. We find that the interval ranges from .048 to .102, thus, while it is smaller than our main estimate (.152), the interval excludes zero. Thus, we find that the peer effect is positive also when we allow quite severe violations of the exclusion restriction.

Conclusion

In this paper, we present evidence of a positive coworker peer effect. Our baseline estimate is that a ten percentage points higher turnout among coworkers causes an increase in the probability of voting of 1.52 percentage points. In comparison to peer effects from other social networks, it is about three times larger than the neighbourhood peer effect in Finan et al. (2021), but substantively smaller than peer effects within the family (Nickerson 2008). Given that the degree of social interactions with colleagues is less than within the family, but (on average) more than with neighbours, we believe that the estimate has face validity. We interpret the estimate in causal terms, and present a range of tests and discussions to underpin this interpretation, but will stress that our study is observational and arguably rests on strong assumptions that need to be kept in mind when given this interpretation.

Our result suggests that workplace peer effects can contribute to social inequalities in turnout. We generally find stronger peer effects among individuals with characteristics that predict lower turnout, which in light of the increasing sorting of workers with similar characteristics into firms (Håkanson, Lindqvist, and Vlachos 2021), could strengthen inequalities in turnout. This is so because low-propensity voters benefit from social interactions with colleagues who are likely to vote. If developments in the labour market reduce the amount of social interactions across groups of voters with different latent propensities to vote, peer effects at the workplace will increase turnout inequalities. For the peer effects to work in the opposite direction, groups with low turnout must interact with coworkers from other backgrounds where turnout is higher. Our results are consistent with other work on workplace political mobilization, for instance by trade unions (Leighley and Nagler 2007), but this work has relied on inferior, aggregated, data.

Our results come with several caveats. We study a high turnout group (employed, native Swedes), which means that external validity might be questioned. We suspect that workplace peer effects might be larger in countries with lower turnout, as ceiling effects are smaller. We also lack measures of social interactions and norms at the workplace, which means that we do not have empirical support for the potential mechanisms that we propose as important. Qualitative work on political conversations at the workplace, for instance on sorting into groups that discuss politics, will also be helpful in this regard. Moreover, experimental work on the consequences of exposure to political conversations for political engagement is lacking. Future work should improve on these weaknesses of current research.

References

Aggeborn, Linuz, and Henrik Andersson. 2022. "Workplace Networks And Political Selection." Working paper https://preprints.apsanet.org/engage/apsa/article-details/6284d3176b12b649b07623e9.

Alt, James E., Amalie Jensen, Horacio Larreguy, David D. Lassen, and John Marshall. 2022. "Diffusing Political Concerns: How Unemployment Information Passed Between Social Ties Influences Danish Voters." *Journal of Politics* 84(1): 383–404.

Andresen, Martin E. 2018. "Exploring marginal treatment effects: Flexible estimation using Stata." *Stata Journal* 18(1): 118–158.

Angrist, Joshua D., and Jörn-Steffen Pischke. 2009. *Mostly Harmless Econometrics: An Empiricist's Companion.* Princeton, NJ: Princeton University Press.

Azar, José, Ioana Marinescu, and Marshall Steinbaum. 2022. "Labor Market Concentration." *Journal of Human Resources* 57(S): S167–S199.

Bhuller, Manudeep, Gordon B. Dahl, Katrine V. Løken, and Magne Mogstad. 2020. "Incarceration, Recidivism, and Employment." *Journal of Political Economy* 128(4): 1269–1324.

Bond, Robert M., Christopher J. Fariss, Jason J. Jones, Adam D.I. Kramer, Cameron Marlow, Jaime E. Settle, and James H. Fowler. 2012. "A 61-million-person experiment in social influence and political mobilization." *Nature* 489(7415): 295–298.

Brady, Henry, Sidney Verba, and Kay Lehman Schlozman. 1995. "Beyond SES: A Resource Model of Political Participation." *American Political Science Review* 89(2): 271–294.

Bramoullé, Yann, Habiba Djebbari, and Bernard Fortin. 2009. "Identification of peer effects through social networks." *Journal of Econometrics* 150(1): 41–55.

Bratsberg, Bernt, Jeremy Ferwerda, Henning Finseraas, and Andreas Kotsadam. 2021. "How settlement locations and local networks influence immigrant political integration." *American Journal of Political Science* 65(3): 551–565.

Brinch, Christian N, Magne Mogstad, and Matthew Wiswall. 2017. "Beyond LATE with a discrete instrument." *Journal of Political Economy* 125(4): 985–1039.

Buggle, Johannes, Thierry Mayer, Seyhun Orcan Sakalli, and Mathias Thoenig. 2023. "The Refugee's Dilemma: Evidence from Jewish Migration out of Nazi Germany." *The Quarterly Journal of Economics* 138(2): 1273–1345.

Caeyers, Bet, and Marcel Fafchamps. 2023. "Exclusion Bias in the Estimation of Peer Effects." Working paper https://web.stanford.edu/~fafchamp/ExclusionBias.pdf.

Campbell, David E. 2013. "Social Networks and Political Participation." *Annual Review of Political Science* 16: 33–48.

Carlson, Taylor N. 2019. "Through the Grapevine: Informational Consequences of Interpersonal Political Communication." *American Political Science Review* 113(2): 325–339.

Carlsson, Magnus, and Abdulaziz Abrar Reshid. 2022. "Co-worker peer effects on parental leave take-up." *Scandinavian Journal of Economics* 124(4): 930–957.

Clarke, Damian, and Benjamín Matta. 2018. "Practical considerations for questionable IVs." *Stata Journal* 18(3): 663–691.

Conley, Timothy G., Christian B. Hansen, and Peter E. Rossi. 2012. "Plausibly Exogenous." *Review of Economics and Statistics* 94(1): 260–272.

Cornelissen, Thomas, Christian Dustmann, and Uta Schönberg. 2017. "Peer Effects in the Workplace." *American Economic Review* 107(2): 425–456.

Cox, Gary W., Jon H. Fiva, and Max-Emil M. King. 2024. "Bound by Borders: Voter Mobilization through Social Networks." *British Journal of Political Science* .

Cravens, Matthew D. 2020. "Measuring the strength of voter turnout habits." *Electoral Studies* 64: 102117.

Dahlgaard, Jens Olav. 2018. "Trickle-Up Political Socialization: The Causal Effect on Turnout of Parenting a Newly Enfranchised Voter." *American Political Science Review* 112(3): 698–705.

Dahlgaard, Jens Olav, Yosef Bhatti, Jonas Hedegaard Hansen, and Kasper M Hansen. 2022. "Living together, voting together: Voters moving in together before an election have higher turnout." *British Journal of Political Science* 52(2): 631–648.

De Giorgi, Giacomo, Michele Pellizzari, and Silvia Redaelli. 2010. "Identification of Social Interactions Through Partially Overlapping Peer Groups." *American Economic Journal: Applied Economics* 2(2): 241–275.

DiMaggio, Paul, and Filiz Garip. 2012. "Network Effects and Social Inequality." *Annual Review of Sociology* 38: 93–118.

Downs, Anthony. 1957. "An Economic Theory of Political Action in a Democracy." *Journal of Political Economy* 65(2): 135–150.

Eggers, Andrew C, Guadalupe Tuñón, and Allan Dafoe. 2023. "Placebo Tests for Causal Inference." *American Journal of Political Science* .

Evans, Geoffrey, and James Tilley. 2017. *The New Politics of Class: The Political Exclusion of the British Working Class*. Oxford University Press.

Finan, Frederico, Enrique Seira, and Alberto Simpser. 2021. "Voting with one's neighbors: Evidence from migration within Mexico." *Journal of Public Economics* 202: 104495.

Finseraas, Henning, Marianne Røed, and Pål Schøne. 2017. "Labor market competition with immigrants and political polarization." *Quarterly Journal of Political Science* 12(3): 347–373.

Fowler, James H., and Christopher T. Dawes. 2008. "Two Genes Predict Voter Turnout." *Journal of Politics* 70(3): 579–594.

Funk, Patricia. 2010. "Social Incentives and Voter Turnout: Evidence From the Swiss Mail Ballot System." *Journal of the European Economic Association* 8(5): 1077–1103.

Gerber, Alan S., and Donald P. Green. 2000. "The Effects of Canvassing, Telephone Calls, and Direct Mail on Voter Turnout: A Field Experiment." *American Political Science Review* 94(3): 653–663.

Gerber, Alan S., Donald P. Green, and Christopher W. Larimer. 2008. "Social Pressure and Voter Turnout: Evidence From a Large-Scale Field Experiment." *American Political Science Review* 102(1): 33–48.

Gerber, Alan S., Donald P. Green, and Ron Shachar. 2003. "Voting May Be Habit-Forming: Evidence From a Randomized Field Experiment." *American Journal of Political Science* 47(3): 540–550.

Gerber, Alan S., Jonathan Gruber, and Daniel M. Hungerman. 2016. "Does Church Attendance Cause People to Vote? Using Blue Laws' Repeal to Estimate the Effect of Religiosity on Voter Turnout." *British Journal of Political Science* 46(3): 481–500.

Green, Donald P., and Alan S. Gerber. 2010. "Introduction to Social Pressure and Voting: New Experimental Evidence." *Political Behavior* 32: 331–336.

Håkanson, Christina, Erik Lindqvist, and Jonas Vlachos. 2021. "Firms and skills the evolution of worker sorting." *Journal of Human Resources* 56(2): 512–538.

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

Heckman, James J., and Edward J. Vytlacil. 2007. "Econometric Evaluation of Social Programs, Part I: Causal Models, Structural Models and Econometric Policy Evaluation." *Handbook of Econometrics* 6: 4779–4874.

Kitschelt, Herbert, and Philipp Rehm. 2014. "Occupations as a Site of Political Preference Formation." *Comparative Political Studies* 47(12): 1670–1706.

Krupnikov, Yanna, and John Barry Ryan. 2022. *The Other Divide: Polarization and Disengagement in American Politics.* Cambridge University Press.

Lal, Apoorva, Mackenzie Lockhart, Yiqing Xu, and Ziwen Zu. 2024. "How Much Should We Trust Instrumental Variable Estimates in Political Science? Practical Advice Based on 67 Replicated Studies." *Political Analysis* pp. 1–20.

Lazarsfeld, Paul F., Bernard Berelson, and Hazel Gaudet. 1944. *The People's Choice: How the Voter Makes Up His Mind in a Presidential Campaign.* Columbia University Press.

Leighley, Jan E., and Jonathan Nagler. 2007. "Unions, Voter Turnout, and Class Bias in the US Electorate, 1964–2004." *Journal of Politics* 69(2): 430–441.

Manski, Charles F. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60(3): 531–542.

Marbach, Moritz, and Dominik Hangartner. 2020. "Profiling Compliers and Noncompliers for Instrumental-Variable Analysis." *Political Analysis* 28(3): 435–444.

Moffitt, Robert A. 2001. "Policy interventions, low-level equilibria, and social interactions." In *Social Dynamics*, ed. Steven N. Durlauf, and H. Peyton Young. MIT Press, Cambridge, US pp. 45–82.

Mutz, Diana, and Jeffery Mondak. 2006. "The Workplace as a Context for Cross-Cutting Political Discourse." *Journal of Politics* 68(1): 140–155.

Nickerson, David W. 2008. "Is Voting Contagious? Evidence From Two Field Experiments." *American Political Science Review* 102(1): 49–57.

Nicoletti, Cheti, Kjell G. Salvanes, and Emma Tominey. 2018. "The Family Peer Effect on Mothers' Labor Supply." *American Economic Journal: Applied Economics* 10(3): 206–234.

Pietryka, Matthew T., and Donald A. DeBats. 2017. "It's Not Just What You Have, But Who You Know: Networks, Social Proximity to Elites, and Voting in State and Local Elections." *American Political Science Review* 111(2): 360–378.

Rolfe, Meredith. 2012. *Voter Turnout: A Social Theory of Political Participation*. Cambridge University Press.

Rosen, Christopher, Joel Koopman, Allison Gabriel, Young Eun Lee, Maira Ezerins, and Philip Roth. 2024. "Hidden Consequences of Political Discourse at Work: How and Why Ambient Political Conversations Impact Employee Outcomes." *Journal of Applied Psychology* .

Shi, Xu, Wang Miao, and Eric Tchetgen Tchetgen. 2020. "A Selective Review of Negative Control methods in epidemiology." *Current Epidemiology Reports* 7: 190–202.

Sinclair, Betsy. 2012. *The Social Citizen: Peer Networks and Political Behavior*. University of Chicago Press.

Smets, Kaat, and Carolien Van Ham. 2013. "The embarrassment of riches? A meta-analysis of individual-level research on voter turnout." *Electoral studies* 32(2): 344–359.

Solt, Frederick. 2008. "Economic Inequality and Democratic Political Engagement." *American Journal of Political Science* 52(1): 48–60.

Online appendix

Contents

A Additional information on the register data

All register data used in this study are managed and maintained by Statistics Sweden and made accessible through MONA (Microdata Online Access), which is Statistics Sweden's tool for delivering microdata for research purposes. The register of the total population (Registret över totalbefolkningen (RTB)) contains information about all individuals registered in Sweden and serves as the basis for our sample. Another key register is register-based labor market statistics (Registerbaserad arbetsmarknadsstatistik (RAMS)), which contains employment information and establishes a unique ID between individuals, their workplaces, and firms.

Another key register is the voter turnout register, which is available for the elections in 2010, 2018, and 2022. This register contains information on, e.g., whether a person was eligible to vote, turnout, and electoral district. Another important register is the longitudinal integrated database for health insurance and labor market studies (Longitudinell integrationsdatabas för sjukförsäkrings- och arbetsmarknadsstudier (LISA)) which contains extensive information on demographics and labor market outcomes such as birth date, gender, marital status, immigrant status, earnings, and level of education.

Finally, the multi-generational register (Flergenerationsregistret) contains the information we need to identify family members; for each individual, it lists parents and children so that we can derive siblings (those who share a parent) and cousins (by first deriving grandparents and then identifying which individuals share a grandparent).

B Descriptive statistics

Table A1: Descriptive statistics. N=390,970.

	(1) Mean	(2) SD
Voted 2022	0.904	0.294
WP Turnout 2018	0.912	0.196
Instrument	0.864	0.158
Index worker voted 2010	0.606	0.489
Index worker voted 2010 missing	0.269	0.443
Income	3,180	1,488
Female	0.431	0.495
Number of children	0.951	1.064
College education	0.314	0.464
Married	0.198	0.398
Age 2010	22.35	5.650
WP relatives income	2,236	1,101
WP relatives employed	0.675	0.211
WP relatives age2010	39.26	6.160
WP relatives married	0.359	0.223
WP relatives n. children	1.071	0.548
Parents voted 2010	0.929	0.257

C Worker characteristics and turnout

Table A2: OLS regression of worker characteristics and turnout.

Voted in 2010	0.078***
	(0.001)
Parents Voted in 2010	0.080***
	(0.002)
Female	0.033***
	(0.001)
Age above median	-0.024***
	(0.001)
College Education	0.052***
	(0.001)
Married	0.020***
	(0.001)
Income above median	0.036***
	(0.001)
Constant	0.741***
	(0.002)
Observations	390,970

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** $p<.01$; ** $p<.05$.

D First stage estimates by sub-groups

Table A3: First-stage results for subgroups of index workers.

	Sample Men	Sample Women	Sample Low edu.	Sample High edu.
Instrument	0.165***	0.117***	0.169***	0.028***
	(0.005)	(0.005)	(0.005)	(0.005)
Controls	Yes	Yes	Yes	Yes
HPG FE	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
N	222464	168493	268176	122703
Mean of Y	0.89	0.94	0.88	0.98

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** p<.01; ** p<.05.

E Results when excluding workers who were ineligible to vote in 2010

Table A4: Results when excluding workers who were ineligible to vote in 2010.

	OLS	First stage	Reduced form	Second stage
WP Turnout	0.045***			0.142***
	(0.004)			(0.030)
Instrument		0.140***	0.020***	
		(0.004)	(0.004)	
Controls	Yes	Yes	Yes	Yes
HPG FE	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Kleibergen-Paap F				1487
N	278131	278131	278131	278131
Mean of Y	0.91	0.92	0.91	0.91

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** $p<.01$; ** $p<.05$.

F Results using binary measures

Table A5: Results when using binary measures of WP Turnout and the instrument.

	OLS	First stage	Reduced form	Second stage
WP Turnout (Binary)	0.016***			0.051***
	(0.001)			(0.010)
Instrument (Binary)		0.097***	0.005***	
		(0.002)	(0.001)	
Controls	Yes	Yes	Yes	Yes
HPG FE	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes
Kleibergen-Paap F				1487
N	390970	390970	390970	390970
Mean of Y	0.90	0.50	0.90	0.90

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** $p<.01$; ** $p<.05$.

G Adding additional control variables

If the instrument is exogenous, the estimate of the main coefficient should not change much when we add additional covariates to main model. In Table A6, we present results when adding various covariates at the individual, HPG, and firm level. First, we add the covariates at the index worker level we use as outcomes in Table 1. For comparison the first column of Table A6 repeats the main specification. The estimates presented in the second column confirm that the IV estimate barely moves when we expand the main specification with these controls. In column 3, we add HPG averages of the predetermined characteristics used in Table 1 as well as the number of workers in the HPG. In column 4, we include predetermined firm characteristics. These are the type of firm (dummies for private, state, municipal, regional, cooperative, and non-profit), total wages at the firm, number of males and females employed, level of education among males and females, turnover, financial result (earnings before interest and taxes), and equity. Finally, in column 5, we include all the covariates in column 2-4 simultaneously. Across specifications, we find that the coefficient is very similar to the main estimate in column 1.

Table A6: Results when adding firm and HPG controls.

	Main estimate	Adding ind X	Adding HPG controls	Adding firm controls	Adding ind X & HPG controls & firm controls
WP Turnout	0.152***	0.149***	0.148***	0.169***	0.164***
	(0.025)	(0.024)	(0.025)	(0.027)	(0.027)
Main controls	Yes	Yes	Yes	Yes	Yes
Ind X	No	Yes	No	No	No
Firm controls	No	No	No	Yes	Yes
HPG controls	No	No	Yes	No	Yes
HPG FE	Yes	Yes	Yes	Yes	Yes
District FE	Yes	Yes	Yes	Yes	Yes
Kleibergen-Paap F	1487	1487	1478	1252	1241
N	390970	390970	390970	305296	305296
Mean of Y	0.90	0.90	0.90	0.90	0.90

Note: Standard errors adjusted for clustering on HPGs in parentheses.
*** p<.01; ** p<.05.

H Marginal Treatment Effects

To estimate MTE:s, we follow Bhuller et al. (2020) and Andresen (2018). We use the same discrete version of the treatment as in the complier analysis, but here we need the continuous instrument. The MTE is defined at the index worker level as the difference in potential outcomes as treated and non-treated, conditional on an unobserved continuous random variable (e.g., interest in politics) and the observed fixed factors (e.g., education). The MTE is estimated by first obtaining the propensity score for being treated, which serves as an estimate of the unobserved gain of treatment, and then exploring how the estimated treatment effect varies along the estimated propensity score.[10]

Figure A1: Common support and marginal treatment effects.

(a) Common support (b) Marginal treatment effects

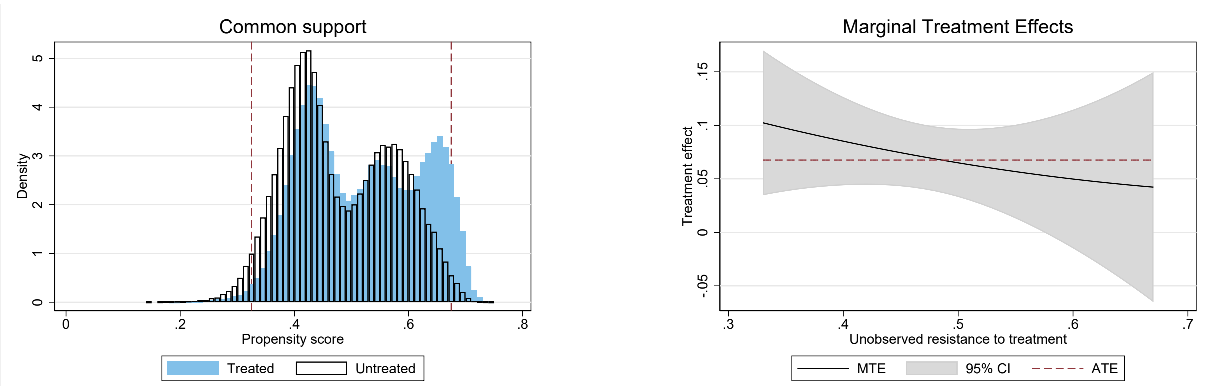

Note: Figure A shows the distributions of the propensity score for treated and untreated workers. The red dotted lines in this figure indicates the cut-offs for trimming the tales of the distributions. We see that while the distributions do not perfectly overlap, we have common support across this range. *Figure B* shows the estimates MTE:s across the inverse of the propensity score, denoted "unobserved resistance to treatment". The black line estimates the slope of the MTE:s, while the shaded area is the 95 percent confidence interval around the point estimates. The dotted red line is the estimated ATE.

Figure A1 A graphs the propensity score distributions for the treated and untreated index workers. The vertical dashed lines show the upper and the lower points of the propensity score with common support.[11] Figure A1 B plots MTE:s estimates by the

[10]In addition to exogeneity of the instrument and monotonicity, the assumption needed to allow identification of MTE:s is separability (Andresen 2018). We estimate the MTE:s based on quadratic specification but the estimates are very similar if we instead use a linear or cubic specification.

[11]Following Bhuller et al. (2020) we trim 1 percent of the sample with overlap.

unobserved variable (e.g., interest in politics). The label on the horizontal axis, "unobserved resistance to treatment", is the terminology used in the literature on MTE:s. There it refers to a situation where participants are randomized into a treatment (e.g., a new medicine) but some participants do not take it because they do not think they benefit from it. In our application, it is natural to think about unobserved resistance to treatment as an unobserved trait such as interest in politics that is correlated with both the magnitude of the treatment effect and turnout among the family members of the coworkers. The dotted red line is the estimated ATE when using the binary treatment variable. We find MTE estimates which are positive over the whole range of the propensity score and the slope of the MTE curve is not significantly different from zero, i.e. we keep the null hypothesis of equal LATE across the unobserved gains.[12] Thus, we find no evidence of differences in the MTE along the unobserved dimension, which makes this an unlikely candidate for explaining the larger IV estimate compared to the OLS estimate.

[12] In our example with interest in politics as the unobserved trait, we expect a positive slope of the MTE curve: Those with the lowest political interest (highest resistance to treatment in the figure) should have larger treatment effects. We see, however, a negative slope, but the slope is not significant.

I Attenuation bias

As noted in the paper, there might be measurement error in the mapping between the turnout among colleagues and the degree of political conversations at the workplace. For instance, while both the extent of political conversations at the workplace and coworker turnout are naturally conceptualized along continuous dimensions that vary in intensity across individuals and HPG:s, the average turnout among coworkers can only be conceptualized as continuous as long as the HPG is large in size. This is evident if we consider the extreme case with a HPG with only two workers. Here the average turnout among coworkers will be a dichotomous variable taking the value zero or one, depending on whether the colleague voted. Thus, for smaller HPG:s the coworker turnout will show a weaker correlation with the latent variables measuring the mechanisms that drive the peer effect. Thus, the average turnout among coworkers is a noisier measure of peer influence in smaller HPG:s.

To investigate to what extent the treatment variable is noisier than the instrument, we plot their distributions in Figure A2. The vertical axis shows the number of observations, while the horizontal axis shows voter turnout in percent. In both graphs, there is a spike at 1 (where everyone in the HPG voted) which is omitted for visual clarity.[13] Clearly, the distribution of the instrument (Figure 2B) is smoother than the distribution of the treatment variable (Figure 2A). This is because the treatment is on average calculated based on 5 coworkers while the instrument is on average calculated using 22 family members of coworkers (because many coworkers have more than one sibling and cousin). Therefore, since coworkers who are more likely to have political conversations at the workplace also will have such conversations with their family members, the instrument, which is based on more individuals, will have a stronger correlation with this mechanism. This implies that the OLS estimate of the peer effect will have an attenuation bias that is reduced in the IV regression.

[13]This spike contains 74 percent and 34 percent of the observations in the left and right graph.

Figure A2: Distribution of the treatment variable and the IV

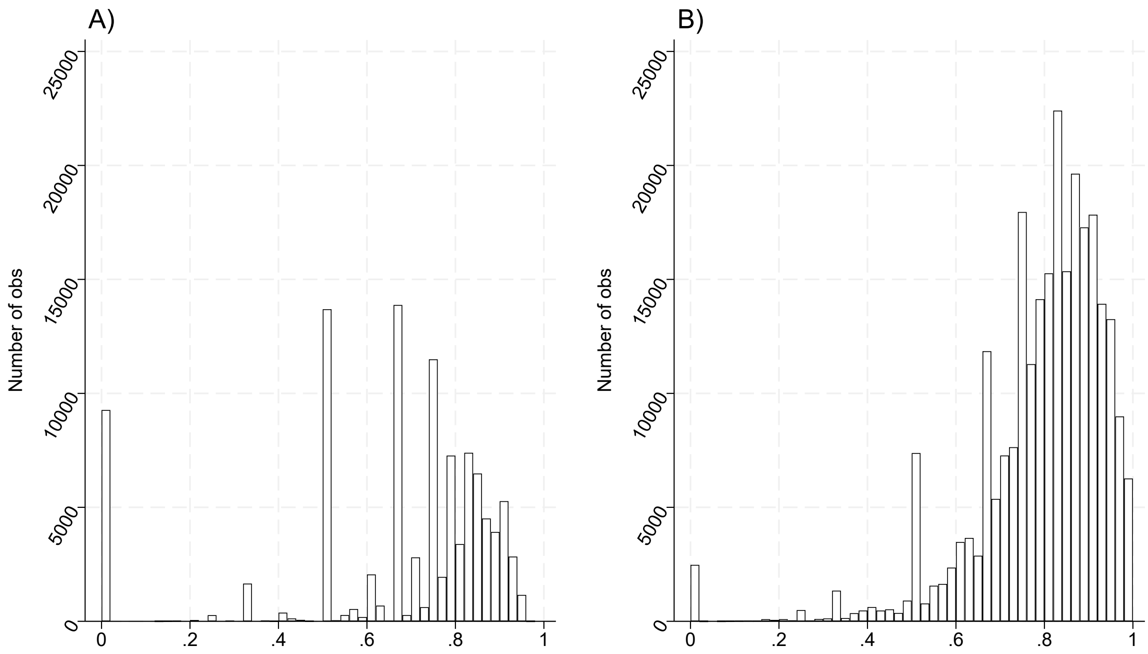

Note: Figure A shows the distribution of co-worker turnout (the treatment), while Figure B shows the distribution of turnout among coworkers family members (the instrument). Both variables spike at the value 1, which is removed for visual clarity.

The next question is to what extent noise in the treatment variable matters empirically for explaining the difference between the OLS and IV estimate. One way of examining this is to compare the OLS and IV estimates for smaller and larger HPG:s. If attenuation bias is important, we expect that the difference between the OLS and IV estimates are larger for smaller HPG:s, since smaller HPG:s have a more noisy estimate of the peer effect. To examine this we create two dummy indicators for small and large HPG:s.[14] We interact these dummies with turnout among coworkers and run a regression similar to those in Table 3. The results in Table A7 show that the IV estimates for small and large HPG:s are virtually the same, while it is especially for small HPG:s that the IV estimate is larger than the OLS estimate. This pattern is consistent with the treatment variable containing more noise for smaller HPG:s, and that the difference in the OLS and IV estimates are partly due to attenuation bias in the OLS estimate.

[14]We split the sample depending on whether the HPG has less or more than 10 workers. Other cutoffs yield similar results.

Table A7: Small and large HPG:s.

	OLS	IV
WP Turnout		
x Small HPG	0.045***	0.151***
	(0.003)	(0.024)
x Large HPG	0.121***	0.161***
	(0.020)	(0.044)

Note: Standard errors adjusted for clustering on HPGs in parentheses. *** p<.01; ** p<.05.

ISBN: 978-91-8082-225-1 (pdf)

Linnæus University